Journal of Drug Discovery and Health Sciences

journal home page : https://jddhs.com/index.php/jddhs/index

Research Article

Artificial Intelligence in Regulatory Affairs: Transforming Drug Development, Regulatory Decision-Making, and Global Compliance

Soma Sekhar Pulamarasetti [1], Siddela Johnson[2], Omkar Rai[1], Kovvada Vandana[2] Bongu Ramya Priya[2]*

1Department of Pharmacy, B. R. Nahata College of Pharmacy, Mandsaur University, Mandsaur, Madhya Pradesh, India (PIN: 458001)
2AU College of Pharmaceutical Sciences (AUCOPS), Andhra University, Visakhapatnam, Andhra Pradesh, India

ARTICLE INFO

Article history:
Received: 01 July, 2026
Revised: 15 July, 2026
Accepted: 12 August, 2026
Published: 25 August, 2026

Keywords:
Artificial Intelligence; Drug Development; Machine Learning; Pharmacovigilance; Clinical Trials as Topic; Government Regulation

DOI:
10.21590/jddhs.03.03.05

ABSTRACT

Background: Artificial intelligence (AI) and machine learning (ML) are increasingly being used throughout the drug development lifecycle, prompting regulatory agencies to develop new frameworks for their use in evidence generation, clinical trial execution, and post-market surveillance.

Objective: To critically appraise current AI/ML applications across drug development and regulatory decision making, compare the regulatory approaches adopted by the US Food and Drug Administration (FDA) and the European Medicines Agency (EMA) and to explore the challenges of global regulatory harmonization and compliance.

Data sources: Literature and guidance documents were searched in PubMed/MEDLINE, Scopus- and Web of Science-indexed journals, Nature Portfolio publications and primary regulatory guidance from FDA and EMA, published predominantly between 2021 and 2026.

Review Methods: We used a comprehensive narrative review format to accommodate the broad scope of the topic, from applied AI/ML literature to regulatory guidance documents and health-policy commentary. Evidence assessed for source authority, considering agency positions, and highlighting convergence and divergence; not chronological summary.

Key Findings: The FDA has transitioned from discussion papers to draft and final guidance to the use of AI in the regulatory decision-making process, device software lifecycle management and predetermined change control, based on its own experience with over 500 submissions of AI components from 2016-2023. The EMA reflection paper sets out a human-centred, risk-based approach covering the whole lifecycle of a medicinal product. AI applications have demonstrated credible retrospective performance improvements in clinical trial recruitment, risk prediction, and pharmacovigilance, but prospective validation and demonstrated impact on trial success or patient safety outcomes remain scarce. The regulators' own use of generative AI tools has led to concerns over reliability due to hallucinations, continuing to underscore the need for oversight by humans.

Conclusion: AI/ML is revolutionizing regulatory affairs in drug development but the current landscape of regulatory structures remains in an early guidance stage of maturity, and global harmonization, algorithmic transparency and prospective outcome validation continue to be the major open challenges to translate the technical promise of AI into reliable regulatory practice.

***Corresponding Author:** Soma Sekhar Pulamarasetti

Address: Assistant Professor, Department of Pharmacy, B. R. Nahata College of Pharmacy, Mandsaur University, Mandsaur, Madhya Pradesh, India (PIN: 458001)

Email ✉**:** dr.somasekhar@gmail.com

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Introduction

Regulatory affairs covers the criteria for the development and monitoring of the safety, efficacy and quality of a pharmaceutical or biological product throughout its life cycle, from first-in-human studies to post-market surveillance. Machine learning and artificial intelligence have come into this domain from two directions simultaneously: sponsors are increasingly using AI/ML to generate the data and analyzes contained in regulatory dossiers, and regulatory agencies themselves are adopting AI tools internally to assist review, and therefore must establish standards for the evaluation of AI-generated evidence submitted by others. The FDA says submissions that include AI elements have grown exponentially, from only a few in 2016 to more than 500 in 2023, a pace that has outstripped the agency's original regulatory framework, which was largely geared toward devices.

Despite its rapid uptake, regulatory science related to AI in drug development is still in its infancy compared with the technical literature. Existing reviews often examine either the technical AI/ML applications (in drug discovery, trial design or pharmacovigilance) or the regulatory guidance documents in isolation, rather than looking at how the two interact, specifically, how the reliability, transparency and validation limitations documented in the technical literature map onto the credibility and risk-based frameworks that agencies like the FDA and EMA have proposed. This gap makes it difficult for regulatory affairs professionals to ascertain which AI applications are ready for use today in support of submissions and which are still hampered by unresolved regulatory or methodological uncertainty.

This review summarizes evidence of AI/ML applications across the drug development lifecycle, with a focus on clinical trial design and pharmacovigilance, and critically discusses the regulatory structures developed by the FDA and EMA to regulate their use, including points of convergence and divergence between the two agencies. It also addresses issues like algorithmic bias, transparency and the challenges of global compliance for regulatory affairs professionals working across jurisdictions. The aim is not a further descriptive catalog of individual applications, but rather a comparative, critically appraised account of where AI's regulatory-facing promise is currently matched by evidence and governance, and where it is not.

AI/ML Applications Across the Drug Development Life Cycle

The adoption of AI/ML methods across the drug development lifecycle is uneven, with the strongest evidence base in early discovery and the weakest in confirmatory, outcome-validated clinical use. In discovery and preclinical development, deep learning and graph neural network approaches enable therapeutic target identification and molecular design, with an increasing number of AI-nominated candidates moving into clinical-stage development. Reported performance is still largely based on retrospective benchmarking rather than prospective confirmation. For example, in chemistry, manufacturing and controls (CMC) and quality domains, AI/ML is used for process optimization and anomaly detection, an area of application where regulators are increasingly looking for documentation within a product's overall AI governance strategy, and not an area that is exempt from oversight simply because it does not have direct patient-facing claims.

The regulatory-facing center of gravity for AI/ML in drug development, and the focus of the remainder of this review, is in clinical trial design and conduct and in pharmacovigilance, as these are the domains in which AI outputs most directly inform the evidence submitted to, or generated by, regulatory agencies. According to the FDA's description of its experience with submissions, the use of AI includes discovery, non-clinical testing, clinical trial design, and post-market safety analysis. The current most intense regulatory scrutiny is on clinical and post-market applications, as they directly impact patient-related decisions.

AI in Clinical Trial Design, Recruitment & Risk Prediction

Clinical trial recruitment and retention remain persistent bottlenecks in drug development, and a 2024 scoping review of AI applications in this domain found that, while AI-based tools, principally natural language processing applied to electronic health records and eligibility criteria, showed promise for accelerating candidate identification and improving investigator ranking of patient eligibility, the underlying evidence base remains methodologically heterogeneous, and the authors concluded that AI's effectiveness for recruitment requires further validation using standardised outcome measures before it can be considered an established practice rather than an emerging one.[7] A separate demonstration of large language model-based patient-to-trial matching, benchmarked directly against eligibility criteria, illustrated the technical feasibility of automating a task that has historically depended on labour-intensive manual chart review, with plausible downstream benefit for both trial timelines and enrolment diversity.[8]

Beyond recruitment, a comprehensive scoping review of 142 studies published between 2013 and 2024 characterised AI applications in clinical trial risk assessment across safety, efficacy, and operational domains, employing traditional machine learning, graph neural networks, transformers, and causal machine learning for tasks including adverse drug event prediction, treatment-effect estimation, and phase-transition prediction.[6] Large language model-based approaches have

grown rapidly within this space, rising from a negligible share of studies before 2022 to a meaningful proportion by 2023, yet the review identified selection bias in training cohorts, a scarcity of prospective studies, and variable data quality as recurring limitations that constrain confidence in the high discriminative performance figures reported in individual studies.[6] This pattern, strong retrospective benchmark performance alongside a marked scarcity of prospective, regulatory-grade validation, recurs across nearly every AI application area discussed in this review and is the central methodological caveat that regulatory affairs professionals should apply when appraising sponsor-submitted AI-derived evidence.

Another asymmetry is highlighted by a comparative reading of the recruitment and risk-prediction literatures: recruitment-focused AI tools are primarily evaluated on process metrics (speed of screening, yield of candidates) whereas risk-prediction tools are evaluated on discriminative accuracy against retrospective outcomes. Neither literature yet shows, in an adequately powered prospective trial, that AI-assisted trial design improves, in a measurable way, the ultimate regulatory outcome of interest, that is, trial success or a favorable benefit-risk determination. This is not a critique of any one study, but a structural gap in the current evidence base which regulatory guidance (discussed in Sections 4 and 5) is only beginning to fill.

Regulatory Systems for AI in Drug Development: US FDA

The FDA's regulatory approach to AI in drug and biological product development has unfolded through a series of discussion papers and guidance documents, rather than a single comprehensive rule, in keeping with the agency's stated preference for iterative, evidence-informed policy development in a fast-moving technical field. The first step was a discussion paper published in May 2023 to solicit public feedback on the use of AI/ML in drug development. The agency states that this paper attracted more than 800 external comments and was informed by an expert workshop convened with the Duke-Margolis Institute for Health Policy in December 2022.[3] This process led to the FDA's first dedicated draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," issued in January 2025 under docket FDA-2024-D-4689. The guidance provides non-binding recommendations on the use of AI to generate data or information intended to support regulatory determinations of safety, effectiveness, or quality.[1]

A critical appraisal of this draft guidance identified its structured, risk-based credibility framework as a genuine strength, since it ties the rigour of validation expected for an AI model to the model's "influence" on the regulatory decision and the decision's underlying risk, rather than applying a uniform standard regardless of context.[4] The same appraisal, however, identified several areas requiring refinement: the guidance's present focus on regulatory decision-making leaves discovery-stage and operational AI applications comparatively unaddressed, its bias-mitigation expectations remain underspecified relative to the algorithmic bias literature, its explainability expectations are not tiered to model risk in the way its credibility framework already ties validation rigour to risk, and its alignment with international regulatory harmonisation efforts remains preliminary.[4] This last point is particularly consequential for global drug development, since sponsors conducting multi-region trials must reconcile FDA expectations with those of other major regulators, a challenge examined further in Section 5.

In parallel with drug- and biologic-focused guidance, the FDA has continued to develop lifecycle-oriented frameworks for AI-enabled device software, including the predetermined change control plan (PCCP), finalised in December 2024, which allows manufacturers to pre-specify anticipated future model modifications and the protocol used to validate them, avoiding a new marketing submission for every incremental retraining cycle.[12] A March 2024 cross-centre coordination document, revised in February 2025, describes how the FDA's medical product centres, CBER, CDER, CDRH, and the Office of Combination Products, are aligning their approaches to AI oversight, reflecting recognition that AI applications increasingly span the traditional boundaries between drug, biologic, and device regulation.[2] As of December 2024, the agency reported having authorised more than 1,000 AI-based medical devices, a volume that stands in useful contrast to the still-nascent, guidance-stage status of AI oversight for non-device drug and biological product submissions.[4]

Regulatory Systems for AI in Drug Development: EMA and Global Harmonization

The European Medicines Agency adopted a materially different tack than the FDA issuing a single reflection paper covering the whole life cycle of the medicinal product rather than a series of narrower, stage-specific guidance documents. In September 2024, after a public consultation in 2023, the EMA's reflection paper on the use of AI in the medicinal product lifecycle was adopted jointly by the Committee for Medicinal Products for Human Use and the Committee for Medicinal Products for Veterinary Use, and sets out a risk-based, human-centred approach to the use of AI in the medicinal product lifecycle from drug discovery and non-clinical development through to clinical trials, manufacturing, and post-authorisation.[5] The paper addresses data acquisition and bias explicitly and requires documented traceability of data sources and processing activities consistent with existing Good Practice (GxP) requirements, and exploratory analysis to identify, and

account for, bias in training data.[5]

Both the EMA and the FDA have converged on risk-based tiering, with validation and transparency expectations scaled to the potential effect of the AI application on patient safety or on the benefit-risk determination, rather than the application of a one-size-fits-all standard regardless of context. A major difference is structural, with the EMA publishing one document covering the lifecycle of the AI device, compared to the increasing but fragmented set of discussion papers and stage-specific guidances from the FDA (covering, at the time of this review, regulatory decision-making evidence, device software lifecycle management, and cross-centre coordination independently). This could create interpretive burden on sponsors who want a single coherent position on AI governance that applies to both jurisdictions. The EMA guidance also notes that in some low-risk contexts, less interpretable "black-box" models can still be acceptable where alternative, more transparent models do less well, a pragmatic allowance that has not yet been explicitly reflected in FDA guidance, which comparatively places more weight on explainability as a general expectation.

Such differences have immediate practical consequence for regulatory affairs professionals managing multi-region development programs: an AI-derived analysis or model validated to FDA credibility-framework expectations is not automatically compliant with EMA's GxP-aligned traceability and bias-documentation expectations, and vise versa. AI governance documentation should be designed from the outset to satisfy the more demanding elements of each applicable framework, rather than retrofitted for a second jurisdiction after initial development.

AI for Pharmacovigilance and Post-Marketing Safety Surveillance

Pharmacovigilance – the systematic collection, detection and assessment of adverse events linked to marketed products – has been a hotbed for the application of AI, since signal detection and case processing are data-intensive and laborious tasks that are amenable to automation. A narrative review of AI applications to pharmacovigilance signal management reported consistent gains in efficiency for automated signal detection, duplicate-case identification, and real-world evidence analysis, and noted the availability of predictive models for anticipating adverse drug reactions and drug-drug interactions ahead of traditional reporting-based detection.[9] A complementary review of machine and deep learning methods for extracting adverse drug events from unstructured clinical text found that transformer-based models, especially BERT-derived architectures, consistently outperformed earlier machine learning approaches on named-entity recognition and relation-classification benchmark tasks that are central to automated case processing.[10]

Both reviews converge on a similar appraisal to that seen in the clinical trial literature: technical benchmark performance is credible and improving, but the reviews explicitly caution that ongoing validation, monitoring, and improvement are essential before AI-driven pharmacovigilance systems can be relied upon as dependable, standalone replacements for expert causality assessment, rather than as decision-support tools operating alongside it.[9] This caution is directly material to regulatory affairs practice, since pharmacovigilance signal outputs feed into regulatory reporting obligations, periodic safety update reports, and, in serious cases, urgent safety restrictions; a false-negative signal missed by an under-validated AI system carries a materially different risk profile than a false positive requiring additional human review, and current guidance from both the FDA and EMA places the burden of demonstrating this asymmetric risk has been adequately managed on the sponsor or marketing authorisation holder deploying the AI tool.

AI Use in Regulatory Agencies Themselves

A separate, though related, trend is the use of AI tools internally by regulatory agencies to support their own review and oversight functions, which has its own credibility implications for regulatory affairs practice. In June 2025, the FDA announced the launch of an agency-wide generative AI tool publicly referred to as Elsa and, in December 2025, the rollout of expanded agentic AI capabilities across agency staff, with the support of an internal CDER AI Council that was created in 2024 to coordinate the center's AI activities. Experience reported with Elsa, however, has raised reliability concerns consistent with known limitations of large language models, such as the generation of false citations and factual hallucinations, prompting the agency's own AI chief to publicly acknowledge that the tool shares this limitation with other contemporary large language models.

This internal experience is instructive for external stakeholders for two reasons. First, it shows that even a well-resourced regulatory agency using AI under presumably favorable conditions, internal data access, defined use cases, and technical oversight, has encountered the same hallucination and reliability limitations documented throughout the sponsor-facing AI literature reviewed in Sections 3 and 6, reinforcing that these are general properties of current-generation large language models, not artifacts of any single deployment context. Second, it explains why the FDA and EMA guidance continues to insist on an obligation of human supervision and expert review of AI-supported outputs, whether these are generated by sponsors or, in the future, by the agencies themselves, rather than treating AI validation as a one-off credibility exercise that obviates the need for downstream human judgment.

Table 1: Regulatory Milestones in AI Governance for Drug Development (2021–2025)

Date	*Agency*	*Document / Action*	*Significance*
Oct 2021	FDA / Health Canada / MHRA	Good Machine Learning Practice: Guiding Principles	First jointly issued cross-agency AI/ML development principles[13]
May 2023	FDA (CDER)	AI/ML for Drug Development Discussion Paper	Foundational public feedback exercise; 800+ comments received[3]
Sep 2024	EMA (CHMP/CVMP)	Reflection Paper on AI in the Medicinal Product Lifecycle	Single lifecycle-spanning, risk-based governance document[5]
Mar 2024 (rev. Feb 2025)	FDA (CBER/CDER/CDRH/ OCP)	Cross-Centre AI Coordination Framework	Harmonised agency-wide approach spanning device and non-device products[2]
Dec 2024	FDA (CDRH)	Final PCCP Guidance for AI-Enabled Devices	Enables prespecified, lower-burden AI model updates[12]
Jan 2025	FDA (CDER/CBER)	Draft Guidance: AI to Support Regulatory Decision-Making	First dedicated drug/biologic AI evidence framework[1]
Jun 2025	FDA (internal)	Agency-wide deployment of generative AI tool ("Elsa")	Regulator's own AI adoption; surfaced hallucination concerns

Algorithmic Bias, Transparency and Credibility Evaluation

Algorithmic bias, which occurs when identifiable subgroups of patients are systematically disadvantaged by the training data, labeling conventions or model design, is a common theme across all the AI applications discussed above: from clinical trial recruitment tools trained on historically non-representative patient populations, to pharmacovigilance systems whose adverse-event training corpora may not represent the reporting patterns of certain demographic groups. Section 4's critical appraisal of FDA draft guidance particularly highlighted bias mitigation as an area that could be strengthened relative to current agency expectations, noting that the guidance's risk-based credibility framework is well suited to validation rigor, but does not yet specify similarly explicit, tiered requirements for bias auditing and subgroup performance reporting.[4] The EMA's reflection paper takes a somewhat more prescriptive stance, explicitly requiring documented identification and mitigation of bias in training data and traceable documentation of data provenance, though implementation guidance for how sponsors should operationalize this requirement in practice is still in its infancy.[5]

Transparency and explainability requirements highlight another point of regulation tension; while fully interpretable models are easier to audit for bias, and easier for reviewers to interrogate, they often underperform less interpretable deep learning or transformer based alternatives on the predictive tasks discussed in Sections 3 and 6. The pragmatic resolution to this tension has been the EMA's overt acceptance of black-box models in certain low-risk contexts; the FDA's ongoing broad emphasis on explainability, without explicit articulation of risk-tiered exceptions, is comparatively a more conservative position. Neither approach has been tested against a large enough body of post-market AI-related regulatory actions to determine which produces better real-world outcomes. This is an empirical question that only will be answered as more AI-supported submissions accumulate post-market experience.

Global Compliance Challenges: Harmonisation and Data Governance

Multinational drug development programmes must currently reconcile the FDA's evolving, stage-specific guidance architecture with the EMA's single lifecycle-spanning reflection paper, and, prospectively, with the AI governance positions of other major regulators such as the UK's Medicines and Healthcare products Regulatory Agency, none of which yet share a common credibility-assessment vocabulary or a harmonised bias-documentation standard. The critical appraisal of FDA guidance discussed above explicitly identified promoting global harmonisation with other regulatory agencies as a priority area for refinement, indicating that even the issuing agency's closest external reviewers regard current cross-jurisdictional alignment as incomplete.[4] In the absence of an International Council for Harmonisation (ICH)-level AI-specific guideline, comparable to existing ICH quality, safety, and efficacy guidelines, sponsors are left to interpret overlapping but non-identical national and regional expectations, increasing both compliance cost and the risk of AI governance documentation that satisfies one jurisdiction's reviewers but not another's.

Data governance considerations compound this challenge. AI models trained or validated using data from one region may be subject to that region's specific privacy, data-residency, and consent requirements, which can constrain a sponsor's ability to use the same trained model, or the

same underlying training dataset, across jurisdictions without additional validation or re-consent processes. This is a structural rather than incidental barrier to AI-enabled drug development at global scale, and it is compounded by the algorithmic bias concern discussed in Section 8: a model trained predominantly on data from one geographic and demographic population carries a materially elevated risk of degraded performance, and potentially of biased output, when applied to populations in a different region without local validation, an outcome that both the FDA's risk-based credibility framework and the EMA's bias-documentation requirements are designed, in principle, to catch, but that neither framework has yet been tested against at scale in a truly global, multi-region AI-supported submission.

Future Perspectives

Three developments are likely to shape the near-term evolution of AI in regulatory affairs. First, the FDA's currently fragmented set of discussion papers and stage-specific guidances is likely to consolidate over time into a more unified drug and biological product AI framework, plausibly incorporating the risk-tiered bias and explainability provisions that current external critical appraisal has identified as missing. Second, growing regulator-side AI adoption, exemplified by the FDA's own generative AI deployment, is likely to generate a parallel body of internal validation and reliability experience that may itself come to inform future sponsor-facing guidance, particularly regarding acceptable hallucination and error-rate thresholds for AI systems supporting regulatory-grade conclusions. Third, continued absence of an ICH-level AI-specific harmonisation guideline is likely to remain the single largest structural barrier to efficient global AI-enabled drug development until such a guideline, or an equivalent multilateral coordination mechanism, is established.

Research Gaps

- Prospective validation demonstrating that AI-assisted trial design or recruitment measurably improves ultimate trial success or benefit-risk outcomes, rather than only process-level or retrospective discriminative metrics, remains largely absent.
- Comparative, empirical evaluation of the FDA's explainability-forward and the EMA's risk-tiered black-box-permissive approaches against real-world post-market AI-related regulatory actions has not yet been conducted.
- Standardised, cross-jurisdictionally recognised methods for AI bias auditing and subgroup performance reporting in regulatory submissions are not yet established.
- Publicly available data on regulator-side generative AI error and hallucination rates in actual review workflows remain limited, constraining external assessment of associated risk.
- No ICH-level or equivalent multilateral AI-specific harmonisation guideline currently exists, leaving sponsors to reconcile divergent national and regional AI governance expectations without a common reference standard.

Clinical Implications

For regulatory affairs professionals and sponsors, the current evidence and guidance landscape supports incorporating AI tools into drug development workflows where their outputs feed into well-scoped, human-reviewed tasks, such as candidate trial-eligibility screening, pharmacovigilance signal triage, or process-level trial optimisation, while treating AI outputs intended to independently support core safety or efficacy conclusions with the heightened validation rigour that both the FDA's risk-based credibility framework and the EMA's lifecycle-spanning reflection paper require. Given the documented gap between retrospective benchmark performance and prospective, outcome-level validation across nearly every AI application area reviewed here, sponsors should expect regulatory reviewers to request the underlying validation evidence supporting any AI-derived data submitted, and should design AI governance documentation to satisfy the more demanding elements of each jurisdiction in which a product will be developed rather than a single reference standard.

Limitations of Current Evidence

This review is a narrative synthesis rather than a formal systematic review, and does not apply independent risk-of-bias or GRADE-equivalent assessment to the regulatory guidance documents discussed, which are policy rather than primary research outputs and are evaluated here on the basis of their stated scope, requirements, and available external critical appraisal rather than empirical outcome data. Regulatory guidance in this area is evolving rapidly; draft guidance documents referenced in this review, including the FDA's January 2025 draft guidance, remain subject to revision or finalisation, and this review's account of the FDA-EMA comparison reflects the state of published guidance at the time of writing rather than a settled, final regulatory landscape. The underlying technical AI/ML literature is itself weighted toward retrospective benchmarking and single-institution or single-region datasets, a limitation inherited from the primary literature discussed throughout this review rather than introduced by the review process itself.

Conclusion

Artificial intelligence is reshaping regulatory affairs across the drug development lifecycle, with the most direct regulatory-facing evidence base concentrated in clinical trial design, recruitment, risk prediction,

and pharmacovigilance. The FDA and EMA have each developed risk-based governance frameworks for AI in drug development, sharing a common emphasis on tiering validation rigour to potential patient impact, while differing structurally in scope and in their treatment of model explainability. Across nearly every application area reviewed, strong retrospective technical performance has outpaced prospective, regulatory-grade validation, and algorithmic bias, transparency, and cross-jurisdictional harmonisation remain substantially unresolved. Regulatory affairs practice in this domain should therefore be paced to the maturity of prospective validation and harmonisation evidence specific to each application and jurisdiction, rather than to the pace of technical publication or regulator-side AI adoption alone.

Data Availability Statement

No new data were generated or analysed in this review. All information discussed is drawn from previously published, publicly available sources cited in the reference list.

References

1. US Food and Drug Administration. Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products (Draft Guidance). Docket No. FDA-2024-D-4689. Silver Spring, MD: FDA; January 2025.
2. US Food and Drug Administration. Artificial Intelligence and Medical Products: How CBER, CDER, CDRH, and OCP are Working Together. Silver Spring, MD: FDA; March 2024 (revised February 2025).
3. US Food and Drug Administration. Using Artificial Intelligence & Machine Learning in the Development of Drug & Biological Products: Discussion Paper and Request for Feedback. Silver Spring, MD: FDA; May 2023 (revised February 2025).
4. Niazi SK. A critical review of the FDA's draft guidance on artificial intelligence in drug and biological product regulation. J Chem. 2026. doi:10.1155/joch/5202999
5. European Medicines Agency. Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle. EMA/CHMP/CVMP/83833/2023. Amsterdam: EMA; September 2024.
6. Teodoro D, Naderi N, Yazdani A, Zhang B, Bornet A. A scoping review of artificial intelligence applications in clinical trial risk assessment. NPJ Digit Med. 2025;8. doi:10.1038/s41746-025-01886-7
7. Lu X, Yang C, Liang L, Hu G, Zhong Z, Jiang Z. Artificial intelligence for optimizing recruitment and retention in clinical trials: a scoping review. J Am Med Inform Assoc. 2024;31(11):2749-2759. doi:10.1093/jamia/ocae243
8. Jin Q, Wang Z, Floudas CS, Chen F, Gong C, Bracken-Clarke D, et al. Matching patients to clinical trials with large language models. Nat Commun. 2024;15:9074. doi:10.1038/s41467-024-53081-z
9. Algarvio RC, Conceição J, Rodrigues PP, Ribeiro I, Ferreira-da-Silva R. Artificial intelligence in pharmacovigilance: a narrative review and practical experience with an expert-defined Bayesian network tool. Int J Clin Pharm. 2025;47(4):932-944. doi:10.1007/s11096-025-01975-3
10. Li Y, Tao W, Li Z, Sun Z, Li F, Fenton S, et al. Artificial intelligence-powered pharmacovigilance: a review of machine and deep learning in clinical text-based adverse drug event detection for benchmark datasets. J Biomed Inform. 2024. doi:10.1016/j.jbi.2024.104621
11. Naik K, Goyal RK, Foschini L, Chak CW, Thielscher C, Zhu H, et al. Current status and future directions: the application of artificial intelligence/machine learning for precision medicine. Clin Pharmacol Ther. 2024;115(4):673-686. doi:10.1002/cpt.3152
12. US Food and Drug Administration. Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions (Final Guidance). Silver Spring, MD: FDA; December 2024.
13. US Food and Drug Administration, Health Canada, UK Medicines and Healthcare products Regulatory Agency. Good Machine Learning Practice for Medical Device Development: Guiding Principles. October 2021.

HOW TO CITE THIS ARTICLE: Pulamarasetti, S. S., Johnson, S., Rai, O., Vandana, K., Priya, B. R. P. Artificial Intelligence in Regulatory Affairs: Transforming Drug Development, Regulatory Decision-Making, and Global Compliance. J. of Drug Disc. and Health Sci. 2026;3(3):20-26. **DOI:** 10.21590/jddhs.03.03.05