

Artificial Intelligence in Regulatory Affairs: Current Applications, Regulatory Challenges, and Future Perspectives in the Pharmaceutical Industry

Fairoosa Thaikkat¹, Afra P², Hanna Asaraf³

^{1,2,3}M.Pharm Pharmaceutical Regulatory affairs Poonthavanam, Kizhattur, Perinthalmanna, Malappuram District, Kerala - 679325

Abstract

Artificial intelligence (AI) is progressively reshaping pharmaceutical regulatory affairs, extending its influence from drug discovery into dossier preparation, electronic submission management, regulatory intelligence, pharmacovigilance, clinical trial operations, and Chemistry, Manufacturing, and Controls (CMC) documentation. This narrative review synthesizes current peer-reviewed literature and official guidance issued by the United States Food and Drug Administration (FDA), the European Medicines Agency (EMA), the UK Medicines and Healthcare products Regulatory Agency (MHRA), Health Canada, Japan's Pharmaceuticals and Medical Devices Agency (PMDA), China's National Medical Products Administration (NMPA), the International Council for Harmonization (ICH), and the World Health Organization (WHO) to characterise the present state, regulatory challenges, and future trajectory of AI adoption in regulatory science. AI-enabled natural language processing (NLP) and large language models (LLMs) are increasingly used to automate Common Technical Document (CTD) authoring, structure electronic Common Technical Document (eCTD) submissions, monitor global regulatory intelligence, and accelerate pharmacovigilance signal detection, while machine learning supports predictive quality-by-design (QbD) frameworks in CMC development. Regulatory agencies have responded with an expanding but fragmented body of guidance, including the FDA's 2025 draft guidance on AI-based regulatory decision-making, the EMA's 2024 reflection paper on AI across the medicinal product lifecycle, and jointly issued Good Machine Learning Practice principles. Despite substantial efficiency gains, critical challenges persist regarding model explainability, algorithmic bias, data provenance, validation of "black-box" architectures, cross-jurisdictional harmonisation, and safeguarding sensitive clinical and manufacturing data against cybersecurity threats. This review critically evaluates the adequacy of existing frameworks, identifies gaps in explainability standards, workforce readiness, and global regulatory convergence, and offers a comparative analysis of national and regional approaches through structured tables. It concludes that realising AI's transformative potential in regulatory affairs will require harmonised, risk-based, and lifecycle-oriented governance, sustained human oversight, and closer collaboration between regulators, industry, and academia. The review provides a scholarly reference point for regulatory professionals, postgraduate researchers, and policymakers navigating the evolving intersection of AI and pharmaceutical regulation.

Keywords: Artificial intelligence; Regulatory affairs; Machine learning; Pharmacovigilance; eCTD; Regulatory intelligence; Quality by design; Drug regulation

1. Introduction

Regulatory affairs constitutes the connective framework between pharmaceutical innovation and public health protection, governing every stage of a medicinal product's lifecycle from preclinical development through post-marketing surveillance. Historically, this function has been characterized by document-intensive, manually executed processes: compilation of the Common Technical Document (CTD), literature-based safety signal review, and country-by-country tracking of shifting regulatory requirements. As the volume, heterogeneity, and pace of biomedical data have expanded, traditional regulatory workflows have increasingly strained under the burden of manual documentation, delayed compliance monitoring, and fragmented global intelligence gathering [1].

Artificial intelligence (AI) — encompassing machine learning (ML), deep learning (DL), natural language processing (NLP), and, more recently, large language models (LLMs) and generative AI — has emerged as a candidate solution to these structural inefficiencies [2,3]. Regulatory agencies have themselves acknowledged this shift: the WHO's 2021 ethical guidance on AI in health explicitly frames AI as central to the ongoing digital transformation of health systems, while cautioning that its data-driven nature introduces new categories of risk that must be actively mitigated [4]. Within pharmaceutical regulatory affairs specifically, AI is now being piloted or deployed across dossier preparation, electronic submission management, regulatory intelligence and horizon scanning, pharmacovigilance signal detection, clinical trial operations, and Chemistry, Manufacturing, and Controls (CMC) documentation [2,3,5].

This growing footprint has prompted a wave of regulatory guidance. The United States Food and Drug Administration (FDA) issued its first discussion paper on AI/ML in drug development in 2023, followed in January 2025 by its inaugural draft guidance specifically addressing the use of AI to support regulatory decision-making for drug and biological products [6,7]. The European Medicines Agency (EMA) finalised its reflection paper on AI across the medicinal product lifecycle in September 2024 [8]. Parallel initiatives have emerged from the UK's Medicines and Healthcare products Regulatory Agency (MHRA), Health Canada, Japan's Pharmaceuticals and Medical Devices Agency (PMDA), and China's National Medical Products Administration (NMPA) [1,9]. Yet these frameworks remain heterogeneous in scope, terminology, and stringency, creating a fragmented global regulatory landscape that pharmaceutical sponsors — particularly those pursuing multi-jurisdictional submissions — must navigate with considerable care.

This narrative review aims to (i) trace the evolution of AI in healthcare and pharmaceuticals as a foundation for understanding its regulatory applications; (ii) systematically examine current AI applications across dossier preparation, eCTD submission, regulatory intelligence, pharmacovigilance, clinical trial regulatory processes, and CMC documentation; (iii) comparatively analyse global regulatory perspectives from the FDA, EMA, MHRA, Health Canada, PMDA, and NMPA; (iv) critically evaluate the ethical, data privacy, and cybersecurity challenges accompanying AI adoption; and (v) identify knowledge gaps and articulate future research and policy directions. The review is intended as a scholarly resource for regulatory affairs professionals, postgraduate researchers, and policymakers engaging with this rapidly evolving domain.

2. Evolution of Artificial Intelligence in Healthcare and Pharmaceuticals

The conceptual origins of AI in biomedicine trace back to early expert systems of the 1970s, such as rule-based diagnostic aids, which demonstrated the feasibility of encoding clinical reasoning into computational logic. In pharmaceutical science, the earliest applications of computational intelligence emerged through computer-aided drug design (CADD) and quantitative structure–activity relationship (QSAR) modelling, which used statistical correlations between chemical structure and biological activity to guide early-stage compound selection. Molecular docking and virtual screening matured through the 1980s and 1990s alongside growing computational power, laying essential groundwork for the machine-learning-driven drug discovery pipelines of the following decades.

The 2010s marked an inflection point, with deep learning architectures — convolutional neural networks, recurrent networks, and later transformer-based models — enabling substantially more accurate prediction of molecular properties, target–ligand interactions, and toxicity profiles [10,11]. A widely cited demonstration of this capability was the discovery of the antibacterial compound halicin through a deep-learning model trained to predict antibacterial activity, illustrating AI's capacity to identify structurally novel candidates that eluded conventional screening approaches [12]. Concurrently, AI began permeating clinical medicine broadly, with applications spanning diagnostic imaging, risk stratification, and clinical decision support, prompting influential commentary on the "convergence of human and artificial intelligence" in medicine [13].

The most consequential recent development for regulatory affairs specifically has been the emergence of generative AI and large language models (LLMs). Unlike earlier predictive ML models confined to narrow, well-defined tasks, LLMs such as GPT-family and related transformer architectures can generate, summarise, and restructure unstructured natural-language text — a capability directly applicable to the document-intensive nature of regulatory submissions, labelling, and safety narratives [3,14]. The FDA has reported a marked and continuing increase in AI-related components within regulatory submissions since 2016, spanning oncology, neurology, and other therapeutic areas, reflecting the breadth of this technological diffusion across the drug development pipeline [7]. This trajectory — from narrow statistical modelling, to deep learning for molecular prediction, to generative AI for document synthesis and decision support — forms the essential backdrop against which contemporary AI applications in regulatory affairs must be understood.

3. AI Applications in Regulatory Affairs: An Overview

AI applications in regulatory affairs can be broadly organised into six interrelated functional domains: (i) dossier preparation and regulatory document management; (ii) electronic Common Technical Document (eCTD) submission management; (iii) regulatory intelligence and compliance monitoring; (iv) pharmacovigilance and signal detection; (v) clinical trial regulatory operations; and (vi) CMC documentation and quality systems. Across these domains, three underlying AI capabilities recur: natural language processing for extracting, classifying, and structuring unstructured regulatory text; predictive machine learning for forecasting outcomes such as approval probability or safety signals; and, increasingly, generative AI for first-draft authoring of regulatory content [2,3,15,16].

Patil and colleagues characterise AI's regulatory affairs contribution as combining human and artificial intelligence to "unlock new possibilities in dossier filing, data extraction, compliance management, and maintenance" of regulatory currency, while cautioning that traditional regulatory processes remain comparatively slow to digitalise relative to other pharmaceutical functions [2]. Similarly, Boopathi and

Raja identify risk assessment, regulatory compliance monitoring, and streamlined approval workflows as the principal near-term value drivers, while flagging ethical considerations, data quality, transparency, and workforce training as the most significant implementation barriers [1]. These recurring benefit–barrier pairings are examined in depth in the domain-specific sections that follow and are synthesised comparatively in Table 2.

4. AI in Dossier Preparation and Regulatory Document Management

Dossier preparation — compiling the Common Technical Document (CTD) modules that document quality, nonclinical, and clinical evidence — has traditionally required extensive manual synthesis across data sources, disciplines, and formatting conventions. NLP-based tools are increasingly applied to convert unstructured text and tabular data into the standardised summaries and structured tables required for CTD modules, reducing manual authoring effort while improving internal consistency [17]. Shi and colleagues demonstrated the application of fine-tuned transformer models (BERT) for automatic semantic labelling of Absorption, Distribution, Metabolism, and Excretion (ADME) content within FDA drug labelling, illustrating how domain-adapted NLP can support product-specific guidance assessment and labelling consistency checks [18].

Beyond labelling, generative AI is being explored for automated extraction and structuring of Summary of Product Characteristics (SmPC) content into the Identification of Medicinal Products (IDMP) data model, using retrieval-augmented generation (RAG) architectures that ground LLM outputs in verified source documents to mitigate hallucination risk [19]. A related application, termed QA-RAG, integrates generative AI with retrieval mechanisms specifically for pharmaceutical regulatory compliance question-answering, allowing regulatory professionals to query complex, evolving guideline text with improved traceability to source documents [20]. These RAG-based architectures are notable because they directly address one of the central risks of LLM deployment in regulated environments — the tendency of generative models to produce plausible but unverified content — by anchoring outputs to authoritative regulatory text.

Knowledge-graph-based approaches represent a further evolution, enabling structured representation of relationships between regulatory requirements, product attributes, and prior submission history, thereby supporting both AI system auditing and downstream transparency for regulators reviewing AI-assisted submissions [21]. Collectively, these developments suggest a trajectory toward "structured content authoring," in which regulatory documents are generated from validated, machine-readable data elements rather than assembled as free text, with AI serving as both the generation engine and the consistency-checking layer.

5. AI in eCTD Submissions and Regulatory Publishing

The electronic Common Technical Document (eCTD) has served as the near-universal backbone of regulatory submission architecture for major health authorities — including the FDA, Health Canada, the EMA, and the PMDA — for over two decades, and is now transitioning toward the more granular, structured eCTD 4.0 standard. This transition, oriented around discrete, reusable data components rather than static PDF documents, is widely regarded as a key enabler for deeper AI integration into submission workflows, since structured data is substantially more amenable to machine processing than unstructured document sets.

Within this evolving architecture, NLP tools are used to structure submission content, validate internal consistency across modules, and detect discrepancies such as mismatched cross-references or inconsistent adverse event terminology prior to formal submission [22]. Automated pre-submission validation tools check content against health-authority-specific technical and structural requirements — for example, FDA, EMA, or PMDA validation criteria — before final compilation, reducing the incidence of technical rejections that historically delayed review timelines [22,23]. Predictive analytics models are additionally being explored to forecast regulatory approval probability and inform submission strategy, complementing rather than replacing the judgement of experienced regulatory strategists [22].

Machine learning applications have also been proposed for automating the generation of clinical data standards artefacts — such as Study Data Tabulation Model (SDTM) mappings — by extracting metadata from raw datasets and predicting case report form (CRF) annotations with reference to prior submission learnings, thereby reducing the substantial manual burden historically associated with clinical data standardisation for regulatory submission [23]. For organisations managing eCTD lifecycle events across dozens of jurisdictions simultaneously, augmented collaboration tools incorporating NLP-based content tagging and automated policy-checking are increasingly embedded within regulatory information management (RIM) platforms, supporting more consistent global lifecycle management even as regional variation formatting requirements persist.

6. AI for Regulatory Intelligence and Compliance Monitoring

Regulatory intelligence — the continuous monitoring, interpretation, and operationalisation of evolving regulations, guidance documents, and health authority communications across jurisdictions — is an inherently information-intensive discipline well suited to AI-based automation. NLP-driven monitoring systems can scan agency websites, official gazettes, and guidance repositories in near real time, flagging relevant regulatory changes for review by regulatory affairs personnel and substantially reducing the lag between publication and organisational awareness [1,2].

Compliance monitoring applications extend this capability to internal documentation, using automated policy-checking to identify discrepancies between current organisational practice and updated regulatory requirements, and, where discrepancies are found, to draft first-pass updates to standard operating procedures or submission content for expert review [1]. This human-in-the-loop model — AI performing initial detection and drafting, with regulatory professionals retaining final judgement and sign-off — recurs as a consistent theme across the regulatory intelligence literature and is broadly consistent with the risk-based, human-oversight expectations articulated in both FDA and EMA AI guidance [7,8].

A persistent challenge in this domain is the absence of standardised, machine-readable regulatory text across jurisdictions; regulatory guidance is published in heterogeneous formats, legal registers, and languages, complicating the training and validation of NLP models intended for cross-jurisdictional intelligence gathering. This structural heterogeneity is a recurring justification, in both academic and regulatory literature, for greater international harmonisation of not only regulatory content itself but also the digital formats in which it is published [1,9].

7. AI in Pharmacovigilance and Signal Detection

Pharmacovigilance — encompassing the detection, assessment, and prevention of adverse drug events — has emerged as one of the most mature domains of AI application within regulatory affairs, driven by escalating case report volumes that have outpaced the capacity of purely manual review. An early and

influential demonstration was Schmider and colleagues' pilot evaluation of AI and robotic process automation for adverse event case processing, which established the technical feasibility of AI-assisted extraction from adverse event source documents and confirmed its capacity to differentiate vendor solution quality during procurement [24]. Subsequent systematic review by Salas and colleagues, synthesising literature published between 2015 and 2021, characterised the expanding scope of AI applications across signal detection, case triage, duplicate detection, and Individual Case Safety Report (ICSR) narrative processing [25].

A landmark special issue in Drug Safety (2022) further consolidated the field's maturity, with Ball and Dal Pan critically assessing whether AI is genuinely "ready for prime time" in pharmacovigilance and cautioning that transparency, validation rigor, and regulatory acceptance remain incomplete despite substantial technical progress [26]. Complementary contributions in the same issue characterised AI-based scoping methodologies for pharmacovigilance applications [27], industry perspectives on practical AI/ML implementation challenges [28], and the broader safety implications of AI and ML adoption for medicines regulation [29]. More recent work employing GPT-4o-class large language models has demonstrated feasibility in automating extraction of case elements — including risk factors and dechallenge/rechallenge responses — from ICSR narratives, while emphasising that human expert verification of model outputs remains essential given variable model performance across extraction categories.

Beyond case processing, machine translation is increasingly used to enable near-real-time evaluation of foreign-language adverse event reports, a capability of particular importance for multinational sponsors monitoring safety signals across linguistically diverse markets. Looking forward, experimental architectures combine adverse event reports with knowledge graphs encoding drug mechanisms and patient characteristics, with the aim of identifying high-risk patient subsets that might otherwise be obscured within aggregate signal detection — an approach that, if validated, could meaningfully advance the precision of post-marketing safety surveillance. Nonetheless, the consistent conclusion across this literature is that AI currently functions most defensibly as a triage and efficiency tool operating under continuous human oversight, rather than as an autonomous causality-assessment system [24–29].

8. AI in Clinical Trial Regulatory Processes

Clinical trials remain the most resource-intensive component of drug development, and AI is increasingly deployed across protocol design, patient recruitment, and regulatory-facing trial documentation. Askin and colleagues' widely cited analysis identifies patient recruitment efficiency and protocol design optimisation as the areas of clearest near-term AI value, noting that machine-learning models trained on historical trial design and patient characteristic data can help predict both trial success probability and regulatory approval likelihood, informing more robust Phase II/III design decisions [30].

AI-based clinical trial matching systems, which mine electronic health records and structured clinical data to identify eligible candidates against complex protocol inclusion and exclusion criteria, have demonstrated real-world feasibility; Alexander and colleagues' evaluation of such a system in Australian lung cancer patients represents an early, empirically grounded example of this application [31]. However, this same literature consistently flags an important caution: because AI-based recruitment tools learn from historical enrolment data reflecting well-documented diversity gaps in clinical research, they risk systematically reproducing — or even amplifying — those gaps unless deliberately audited and corrected for representativeness, a concern with direct relevance to regulatory expectations around trial generalisability.

On the regulatory-process side specifically, the FDA has acknowledged active exploration of AI for predictive modelling and in silico simulation to inform clinical trial design, alongside its use in constructing electronic health record-based patient cohorts and phenotypes to support both trial planning and post-approval evidence generation [7]. The ICH's guideline on Model-Informed Drug Development (MIDD) explicitly incorporates AI/ML methodologies within its scope, formally situating machine-learning-based modelling and simulation as a category of regulatory evidence subject to established methodological expectations for model credibility and evidentiary reporting [32]. Explainability remains a cross-cutting concern here as elsewhere: Amann and colleagues' multidisciplinary analysis underscores that explainability requirements for AI in healthcare — and by extension in AI-supported trial design — must be calibrated to the clinical and regulatory stakes of the specific application rather than applied uniformly [33].

9. AI in CMC Documentation and Quality Systems

Chemistry, Manufacturing, and Controls (CMC) documentation sits at the intersection of regulatory affairs and pharmaceutical quality systems, and has become a focal point for AI integration under the Quality by Design (QbD) paradigm established by ICH Q8–Q11. Zhu's AI-powered information management framework exemplifies this convergence, combining machine learning and deep learning models to predict Critical Quality Attributes (CQAs) from Critical Process Parameters (CPPs) with NLP-based management of the accompanying unstructured regulatory documentation, while incorporating explainable AI (XAI) techniques — including SHAP and LIME — specifically to preserve interpretability in alignment with ICH Q8–Q11 expectations [34]. This explicit engineering-in of explainability reflects a broader recognition that CMC applications of AI, given their direct bearing on product quality and patient safety, warrant particularly rigorous interpretability standards.

Complementary work by Duarte and colleagues reframes QbD itself through a digital lens, arguing that AI, digital twins, and advanced data analytics are evolving traditional QbD into a "quality by digital design" framework capable of real-time monitoring, predictive process control, and continuous improvement across the manufacturing lifecycle [35]. Within this framework, Process Analytical Technology (PAT) — long a cornerstone of regulatory-accepted real-time release testing — is being extended through AI-enabled predictive maintenance and anomaly detection, with regulators including the FDA supporting continuous manufacturing approaches formalised in ICH Q13 as a natural complement to AI-enabled process control [36,37].

The EMA's reflection paper explicitly addresses AI/ML use within CMC and manufacturing contexts, recommending that applicants document data acquisition and processing activities with sufficient traceability to satisfy GxP data integrity expectations, and that Good Laboratory Practice (GLP) principles for computerised systems be extended, where applicable, to AI/ML tools used in non-clinical and manufacturing-adjacent contexts [8]. A recurring implementation challenge identified across this literature is the tension between the adaptive, continuously learning nature of many ML models and the change-control expectations of validated GxP computerised systems, which traditionally assume a fixed, previously validated state — a tension that has directly motivated the Predetermined Change Control Plan (PCCP) mechanisms discussed in Section 10.

10. Global Regulatory Perspectives

Regulatory approaches to AI in pharmaceutical development and oversight have evolved largely indepen-

dently across jurisdictions, producing a landscape that is directionally convergent — most agencies endorse risk-based, lifecycle-oriented, human-centric principles — but procedurally fragmented in scope, terminology, and enforceability. Fu and colleagues' comparative global analysis characterises this landscape as "nascent," with agencies at markedly different stages of formalising AI-specific policy despite broadly shared underlying objectives of patient safety, data integrity, and innovation facilitation [9].

10.1 United States (FDA)

The FDA has adopted the most extensively documented regulatory posture to date, beginning with its 2023 discussion paper inviting public comment on AI/ML use in drug and biological product development [6], followed by cross-centre alignment across CDER, CBER, CDRH, and the Office of Combination Products [38]. Its January 2025 draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," introduces a seven-step, risk-based credibility assessment framework centred on a model's defined "Context of Use" (COU) — requiring sponsors to define the regulatory question, specify the COU, assess model risk, develop and execute a credibility assessment plan, and document results accordingly [7,39,40]. On the medical device side, the FDA has additionally finalised guidance on Predetermined Change Control Plans (PCCPs), permitting manufacturers to pre-specify anticipated AI model modifications and validation protocols at initial authorisation, thereby reducing the need for repeated marketing submissions as adaptive algorithms evolve [41].

10.2 European Union (EMA)

The EMA's finalised reflection paper (September 2024) addresses AI/ML across the full medicinal product lifecycle — development, authorisation, and post-authorisation — for both human and veterinary medicines [8]. Distinctively, the EMA explicitly permits "black-box" models where developers can substantiate that transparent alternatives are unsatisfactory, provided applicants draw on explainable AI methods where feasible and adopt a human-centric approach grounded in the EU's broader Trustworthy AI principles and the Assessment List for Trustworthy Artificial Intelligence (ALTAI) [8]. This comparatively pragmatic stance toward model transparency — contrasted with more prescriptive explainability expectations emerging elsewhere — reflects the EMA's stated intent to avoid foreclosing legitimately high-performing but less interpretable architectures, while still requiring rigorous risk-proportionate validation.

10.3 United Kingdom (MHRA)

The MHRA's 2024 strategic approach to AI addresses risk and opportunity across three lenses — AI as a regulated product, AI as a tool used by the regulator itself, and AI's broader impact on the pace of medicine development and personalisation — and commits to ensuring regulatory pathways remain "sufficiently agile and robust" to accommodate the technology's trajectory. The MHRA's AI Airlock regulatory sandbox, pairing manufacturers and approved bodies with the National Health Service to pilot novel AI-enabled products pre-market, exemplifies a more experimental, iterative regulatory posture relative to the more codified US and EU frameworks.

10.4 Canada (Health Canada)

Health Canada has pursued parallel tracks for AI in medical devices and in evidentiary submissions supporting reimbursement decisions. Its draft pre-market guidance for machine-learning-enabled medical devices complements its status as co-signatory, alongside the FDA and MHRA, of the ten Good Machine Learning Practice (GMLP) guiding principles first issued in 2021 — covering representative training data,

appropriate human oversight, and lifecycle performance monitoring. Separately, Canada's Drug Agency has issued a position statement on AI use in the generation and reporting of health technology assessment evidence, explicitly drawing on the UK National Institute for Health and Care Excellence's (NICE) 2024 position statement and emphasising that AI methods should be used "judiciously," only where they demonstrably add value over conventional approaches.

10.5 Japan (PMDA)

Japan's PMDA has focused its formal AI guidance on the medical device domain, issuing points to consider on the handling of changes to AI-based Software as a Medical Device (AI-SaMD) in March 2023, alongside its established Post-Approval Change Management Protocol (PACMP) framework, which — while not AI-specific in origin — is increasingly referenced as a template for managing anticipated AI model updates within a pre-agreed regulatory pathway. Relative to the FDA and EMA, formal PMDA guidance specific to AI in drug and biologic (as opposed to device) regulatory decision-making remains comparatively less developed as of this review, representing a notable gap in the current global landscape.

10.6 China (NMPA)

China's regulatory approach, as characterised by Fu and colleagues, emphasises data sovereignty and localised validation, with the NMPA requiring extensive validation using Chinese population-specific data to address concerns regarding genetic and environmental differences that may affect AI model performance [9]. This reflects a broader national policy architecture — including the 2021 State Council implementation opinions on strengthening drug regulatory capacity through informatisation — that positions AI governance within a wider framework of data security and algorithmic sovereignty legislation, resulting in a regulatory posture that is procedurally more conservative than the FDA's but is reported to be modernising through pilot programmes and regulatory sandboxes.

10.7 Comparative Synthesis

Table 1 synthesises these jurisdictional approaches across six comparative dimensions: primary guidance instrument, lifecycle scope, explainability stance, validation and data expectations, adaptive/continuous-learning accommodation, and current stage of formalisation. The comparison underscores that while risk-based, lifecycle-oriented principles are broadly shared, meaningful divergence persists in explainability tolerance, data localisation requirements, and the procedural mechanisms available for managing adaptive AI models post-authorisation — divergence that carries direct practical implications for sponsors pursuing simultaneous multi-jurisdictional development and submission strategies.

Jurisdiction / Agency	Primary AI Guidance Instrument	Lifecycle Scope	Explainability Stance	Adaptive / Continuous Learning Accommodation
USA (FDA)	Draft Guidance on AI for Regulatory Decision-Making (Jan 2025); GMLP; PCCP guidance	Nonclinical, clinical, manufacturing, post-marketing (excludes discovery)	Risk-based credibility framework tied to Context of Use; explainability expected to scale with risk	Predetermined Change Control Plans permit pre-authorised model updates

Jurisdiction / Agency	Primary AI Guidance Instrument	Lifecycle Scope	Explainability Stance	Adaptive Continuous Learning / Accommodation
EU (EMA)	Reflection Paper on AI in the Medicinal Product Lifecycle (Sep 2024)	Full lifecycle: discovery, non-clinical, clinical, authorisation, post-authorisation	Pragmatic; black-box models acceptable if justified; ALTAI-based trustworthy AI approach	Risk-based deployment and performance-monitoring expectations; multi-annual workplan (2025–2028)
UK (MHRA)	MHRA Strategic Approach to AI (2024); AI Airlock sandbox	Regulator-as-user, product regulation, and R&D pace considered jointly	Agility-oriented; no fixed explainability threshold published	Sandbox-based iterative evaluation of adaptive AI-enabled products
Canada (Health Canada / CDA-AMC)	Draft pre-market ML-device guidance; GMLP (joint with FDA/MHRA); CDA-AMC AI position statement	Medical devices and HTA evidence generation	Judicious-use principle; AI used only where it demonstrably adds value	GMLP lifecycle performance-monitoring principles
Japan (PMDA)	AI-SaMD Change-Handling Guidance (2023); PACMP framework	Primarily medical devices; drug/biologic AI guidance less developed	Not extensively codified for drugs/biologics as of this review	PACMP-based pre-agreed change pathways
China (NMPA)	NMPA AI-related regulatory science initiatives; State Council informatisation policy	Drug regulation broadly; strong manufacturing/data-governance emphasis	Requires local population validation; source-code/registration scrutiny	Emphasis on sovereignty over data/algorithms; sandbox pilots emerging

Table 1. Comparative summary of global regulatory approaches to artificial intelligence in pharmaceutical regulation, based on published agency guidance and peer-reviewed comparative analyses [1,6–9,38–41].

11. Regulatory Challenges and Ethical Considerations

Despite considerable enthusiasm for AI's efficiency potential, its integration into regulatory affairs surfaces a distinct set of challenges that extend well beyond conventional software validation concerns. The "black-box" problem — the opacity of complex model architectures, particularly deep neural networks, to human interpretation — represents perhaps the most persistent obstacle to regulatory trust, since decisions materially affecting patient safety or product approval require a level of accountability that opaque models cannot straightforwardly provide [8]. Explainable AI (XAI) techniques such as SHAP and LIME are increasingly proposed as mitigations, translating opaque model outputs into human-interpretable feature attributions, but their adoption remains inconsistent across regulatory applications and their sufficiency for high-stakes regulatory decisions remains actively debated [33,34].

Algorithmic bias constitutes a related and equally consequential concern. Because AI models learn from historical data, they risk encoding and perpetuating pre-existing inequities — for example, the well-documented overrepresentation of European and Caucasian populations in clinical research datasets — with direct implications for the generalisability and fairness of AI-supported clinical trial recruitment, safety signal detection, and, ultimately, regulatory decisions premised on those outputs. The EMA's reflection paper explicitly requires that data acquisition and augmentation processes be documented with sufficient granularity to permit bias identification and mitigation, reflecting the level of regulatory seriousness now attached to this issue [8].

A further structural challenge concerns the mismatch between adaptive, continuously learning AI models and the traditionally static, validated-state assumptions underpinning GxP computerised systems and conventional marketing authorisation frameworks. Predetermined Change Control Plans and Post-Approval Change Management Protocols represent regulatory attempts to reconcile this tension by pre-authorising defined categories of model modification, but questions remain regarding the appropriate scope, frequency, and evidentiary threshold for such pre-authorised changes, particularly for models exhibiting emergent or unanticipated behavioural drift over time [8,41].

Workforce readiness constitutes an underappreciated but consequential barrier. Boopathi and Raja explicitly identify workforce training requirements as a significant hindrance to AI/ML adoption in regulatory affairs, noting that regulatory professionals require new competencies spanning data science literacy, model validation principles, and critical evaluation of AI-generated outputs — competencies not traditionally embedded within regulatory affairs curricula [1]. This gap has direct implications for pharmacy education, including postgraduate programmes in regulatory affairs, which will need to incorporate AI literacy as a core competency rather than an elective specialisation if the profession is to meaningfully participate in, rather than merely react to, AI-driven regulatory transformation.

Finally, the WHO's ethical framework foregrounds six overarching principles — protection of human autonomy, promotion of human well-being and safety, transparency and explainability, responsibility and accountability, inclusiveness and equity, and promotion of AI that is responsive and sustainable — that collectively provide a normative foundation against which sector-specific regulatory guidance can be evaluated [4]. Applied to regulatory affairs specifically, these principles imply that AI deployment decisions should not be evaluated on efficiency grounds alone, but against their capacity to preserve meaningful human oversight, equitable outcomes, and accountable decision trails throughout the product lifecycle.

12. Data Privacy and Cybersecurity Concerns

AI systems deployed within regulatory affairs frequently require access to sensitive datasets — including patient-level clinical trial data, individual case safety reports, proprietary manufacturing process parameters, and unpublished submission content — raising data privacy and cybersecurity concerns that are distinct from, though related to, the model-performance challenges discussed above. Compliance with established privacy frameworks, including the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in the European Union, is a baseline expectation, but AI-specific risks extend beyond conventional data protection compliance.

A particular concern is the inadvertent exposure of confidential regulatory or clinical data through unauthorised use of general-purpose, cloud-hosted generative AI tools by regulatory or clinical personnel, a risk that industry surveys suggest remains inadequately addressed by automated technical controls across much of the pharmaceutical sector. This risk is compounded where AI vendor systems operate as third-party processors of regulated data, necessitating robust contractual, technical, and audit safeguards to maintain data integrity and confidentiality consistent with GxP expectations. The EMA's reflection paper explicitly extends Good Laboratory Practice principles for computerised systems and data integrity to AI/ML tools, signalling that data governance obligations traditionally associated with validated laboratory and manufacturing systems now extend, by regulatory expectation, to AI systems processing GxP-relevant data [8].

Cybersecurity considerations are further complicated by the fact that many AI models — particularly those with adaptive or continuous-learning architectures — present an evolving attack surface, since model behaviour, and therefore vulnerability profile, may shift over time as the model is updated or retrained. Adversarial vulnerabilities, whereby carefully crafted inputs can induce incorrect model outputs, represent an additional and still-maturing area of concern for AI systems supporting safety-critical regulatory functions such as pharmacovigilance signal triage or CMC quality prediction. Regulatory frameworks have generally been slower to formalise AI-specific cybersecurity expectations than data privacy expectations, representing an identifiable gap addressed further in Section 14.

13. Discussion

Synthesising the evidence reviewed above, several cross-cutting patterns emerge. First, AI adoption within regulatory affairs is proceeding unevenly across functional domains: pharmacovigilance and CMC/QbD applications exhibit comparatively mature evidence bases and explicit regulatory acknowledgment, whereas dossier authoring and regulatory intelligence applications — though rapidly advancing, particularly with generative AI — remain less thoroughly validated against regulatory-grade evidentiary standards. This unevenness suggests that blanket assessments of AI "readiness" for regulatory affairs are of limited analytical value; readiness must instead be assessed domain-by-domain and use-case-by-use-case, consistent with the risk-based, Context-of-Use framing adopted by the FDA [7].

Second, a persistent tension runs through the reviewed literature between the pursuit of model performance and the preservation of interpretability, transparency, and human oversight. Agencies have not converged on a uniform resolution to this tension: the FDA's risk-tiered credibility framework and the EMA's conditional acceptance of black-box models under justified circumstances both represent principled but distinct compromises, reflecting different underlying risk tolerances and legal-administrative traditions [7,8]. For sponsors operating across jurisdictions, this divergence introduces genuine strategic complexity,

since a validation and documentation package sufficient for one agency may not straightforwardly satisfy another's expectations.

Third, the reviewed literature reveals a notable asymmetry between the pace of AI technological advancement — particularly the rapid emergence of generative AI and LLM capabilities since 2022–2023 — and the pace of regulatory guidance development, which necessarily proceeds through more deliberative, consultative processes. This asymmetry is not unique to regulatory affairs, but its consequences are particularly acute in a domain where guidance vacuum can translate directly into either overly cautious underutilisation of beneficial AI tools or, conversely, premature deployment of inadequately validated systems in consequential regulatory workflows.

Fourth, workforce and organisational readiness emerge as under-examined but material determinants of successful AI integration. Much of the reviewed technical and regulatory literature focuses on model capability and regulatory permissibility, with comparatively less systematic attention to the practical competencies, governance structures, and change-management processes required for regulatory affairs organisations to responsibly operationalise these tools. This represents an important complement to the more technically oriented literature reviewed in Sections 4–9.

14. Future Perspectives

Several trajectories appear likely to shape the future evolution of AI in regulatory affairs. First, continued convergence toward structured, machine-readable regulatory data — exemplified by the eCTD 4.0 transition and IDMP-aligned data models — will likely serve as a foundational enabler for more sophisticated AI applications, since AI systems perform markedly better against structured, validated data than against heterogeneous free text. Regulatory affairs professionals should anticipate that data structuring and standardisation, though less visibly transformative than generative AI applications themselves, will constitute an essential and substantial share of near-term implementation effort.

Second, retrieval-augmented generation and knowledge-graph-grounded architectures are likely to become the dominant pattern for generative AI deployment in regulated documentation contexts, given their capacity to constrain LLM outputs to verifiable source material and thereby directly address hallucination risk — a capability increasingly viewed as a prerequisite, rather than an optional enhancement, for regulatory-grade generative AI applications [19,20,21].

Third, international regulatory harmonisation efforts — potentially channelled through ICH working groups, building on precedent established by guidelines such as ICH M15 on model-informed drug development — represent a plausible and valuable mechanism for narrowing the jurisdictional divergence documented in Table 1, particularly regarding explainability expectations and change-control mechanisms for adaptive models. Given the substantial compliance burden that jurisdictional fragmentation imposes on multinational sponsors, industry and academic stakeholders alike have strong incentives to support such harmonisation initiatives.

Fourth, explainable AI is likely to transition from a largely optional technical enhancement to an increasingly explicit regulatory expectation, particularly for higher-risk applications such as pharmacovigilance causality assessment and CMC critical quality attribute prediction, mirroring the trajectory already evident in the EMA's and FDA's respective guidance documents [7,8]. Relatedly, formal frameworks for auditing AI system behaviour over time — addressing model drift, performance degradation, and emergent bias — are likely to mature substantially, potentially drawing on knowledge-graph-based auditing approaches already proposed in the academic literature [21].

Finally, from an educational and professional-development perspective, regulatory affairs curricula — including postgraduate programmes such as M.Pharm specialisations in Regulatory Affairs — will need to systematically incorporate AI literacy, covering not only the operational use of AI-enabled regulatory tools but also critical evaluation skills: the capacity to interrogate model provenance, assess validation adequacy, and recognise the limitations of AI-generated regulatory content. This competency shift is likely to become a meaningful differentiator for regulatory professionals entering the field over the coming decade.

14.1 Synthesis of Benefits, Challenges, and Regulatory Considerations

Table 2 consolidates the benefits, challenges, and corresponding regulatory considerations of AI implementation in regulatory affairs identified throughout this review, organised by functional domain, to support structured risk assessment by regulatory affairs professionals and organisational decision-makers.

Functional Domain	Key Benefits	Key Challenges	Regulatory Considerations
Dossier preparation & documentation	Faster CTD authoring; improved cross-module consistency; automated labelling/ADME classification [17,18]	Hallucination risk in generative drafting; variable terminology standardisation across sources	RAG-grounded generation preferred; human medical-writer review mandatory before submission
eCTD submissions & publishing	Automated pre-submission validation; reduced technical rejections; SDTM/CRF automation [22,23]	Regional variance; formatting legacy document/data structure incompatibility	Alignment with eCTD 4.0 structured-content transition; agency-specific validation rule compliance
Regulatory intelligence & compliance	Near real-time regulatory change monitoring; automated gap analysis vs. internal SOPs [1,2]	Heterogeneous, non-machine-readable source formats across jurisdictions	Human-in-the-loop review of AI-flagged changes prior to policy or SOP amendment
Pharmacovigilance & signal detection	Faster ICSR triage; improved duplicate detection; scalable case processing [24–29]	Validation rigor and transparency gaps; variable extraction accuracy across case elements	GVP-aligned logging of AI actions; mandatory expert review of

Functional Domain	Key Benefits	Key Challenges	Regulatory Considerations
			AI-flagged signals
Clinical trial regulatory processes	Improved recruitment efficiency; MIDD-aligned predictive trial design [30–32]	Risk of amplifying historical enrolment bias; limited external validation of matching algorithms	Bias auditing of training data; ICH MIDD credibility documentation for regulatory submissions
CMC documentation & quality systems	Predictive CQA modelling; real-time release testing support; XAI-enabled interpretability [34–37]	Adaptive-model vs. fixed-validated-state conflict; data integrity across structured/unstructured sources	ICH Q8–Q13 alignment; PCCP/PACMP-based change control for adaptive models
Data privacy & cybersecurity	Enables large-scale analysis of clinical/manufacturing data for compliance and safety insight	Unauthorised exposure via general-purpose AI tools; adversarial vulnerability; third-party processor risk	HIPAA/GDPR compliance; GLP-aligned data integrity for AI/ML tools; vendor audit and contractual safeguards [8]

Table 2. Benefits, challenges, and regulatory considerations of artificial intelligence implementation across functional domains of regulatory affairs.

15. Knowledge Gaps and Future Research Directions

This review identifies several specific knowledge gaps warranting further research. First, there is a paucity of prospective, regulatory-grade validation studies for generative AI applications in dossier authoring; most currently available evidence derives from retrospective feasibility studies or proof-of-concept demonstrations rather than validation against pre-specified regulatory acceptance criteria. Second, comparative empirical research quantifying the real-world impact of divergent jurisdictional explainability standards on multinational submission timelines and outcomes remains largely absent, despite its evident practical importance to industry. Third, the long-term implications of adaptive, continuously learning AI models for post-authorisation change control remain conceptually addressed (through PCCP and PACMP mechanisms) but empirically under-studied, particularly regarding appropriate monitoring intervals and drift-detection thresholds. Fourth, research specifically addressing AI literacy curricula and competency frameworks for regulatory affairs professionals — as distinct from data scientists or biostatisticians —

remains sparse, representing an important gap for pharmacy and regulatory affairs education research. Finally, empirical evaluation of cybersecurity vulnerabilities specific to AI systems operating within GxP-regulated environments, as opposed to AI cybersecurity generally, constitutes an emerging but still nascent area meriting dedicated regulatory science attention.

16. Conclusion

Artificial intelligence is progressively transforming pharmaceutical regulatory affairs, offering substantive efficiency and quality gains across dossier preparation, eCTD submission management, regulatory intelligence, pharmacovigilance, clinical trial operations, and CMC documentation. Major regulatory agencies — including the FDA, EMA, MHRA, Health Canada, PMDA, and NMPA — have responded with an expanding, though still fragmented, body of guidance grounded in broadly shared risk-based and human-centric principles, even as meaningful divergence persists in explainability expectations, data localisation requirements, and mechanisms for managing adaptive AI models. Realising AI's full transformative potential in this domain will require sustained collaboration between regulators, industry, and academia; continued investment in explainable, auditable AI architectures; deliberate attention to algorithmic bias and data governance; and a parallel commitment to building AI literacy within the regulatory affairs workforce. As AI capabilities continue to advance rapidly, regulatory frameworks — and the professionals who operate within them — must evolve in step, preserving the core regulatory mandate of ensuring safe, effective, and high-quality medicines reach patients, while responsibly harnessing AI's capacity to make that mandate more efficient, more consistent, and more responsive to an increasingly complex global drug development landscape.

References

1. Boopathi LK, Raja N. Artificial intelligence and machine learning: enhancing the future of regulatory affairs. *Int J Drug Regul Aff.* 2024;12(4):53-60. doi:10.22270/IJDRA.V12I4.718
2. Patil RS, Kulkarni SB, Gaikwad VL. Artificial intelligence in pharmaceutical regulatory affairs. *Drug Discov Today.* 2023;28(9):103700. doi:10.1016/j.drudis.2023.103700
3. Somara S, Joshi AM, Mitra K, Desai S, Lundberg MS, Bhasin S, Hunsberger J. Artificial intelligence in biotechnology and pharmaceuticals: evolution, applications, and regulatory frontiers. *Curr Stem Cell Rep.* 2025;11(1).
4. World Health Organization. *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance.* Geneva: World Health Organization; 2021.
5. Oualikene-Gonin W, Jaulent MC, Thierry JP, Oliveira-Martins S, Belgodère L, Maison P, Ankri J. Artificial intelligence integration in the drug lifecycle and in regulatory science: policy implications, challenges and opportunities. *Front Pharmacol.* 2024;15:1437167. doi:10.3389/fphar.2024.1437167
6. U.S. Food and Drug Administration. *Using Artificial Intelligence and Machine Learning in the Development of Drug and Biological Products: Discussion Paper and Request for Feedback.* Silver Spring (MD): FDA; 2023 May (revised 2025 Feb).
7. U.S. Food and Drug Administration. *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products: Draft Guidance for Industry.* Silver Spring (MD): FDA; 2025 Jan. Available from: <https://www.fda.gov/media/184830/download>
8. European Medicines Agency. *Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle.* EMA/CHMP/CVMP/83833/2023. Amsterdam: EMA; 2024 Sep.

Available from: https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

9. Fu L, Jia G, Liu Z, Pang X, Cui Y. The applications and advances of artificial intelligence in drug regulation: a global perspective. *Acta Pharm Sin B*. 2025;15(1):1-14. doi:10.1016/j.apsb.2024.11.006
10. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. *Nat Rev Drug Discov*. 2019;18(6):463-477. doi:10.1038/s41573-019-0024-5
11. Fleming N. How artificial intelligence is changing drug discovery. *Nature*. 2018;557(7706):S55-S57. doi:10.1038/d41586-018-05267-x
12. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, et al. A deep learning approach to antibiotic discovery. *Cell*. 2020;180(4):688-702.e13. doi:10.1016/j.cell.2020.01.021
13. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
14. Kim J, Hur M, Min M. From RAG to QA-RAG: integrating generative AI for pharmaceutical regulatory compliance process. *Pharmacol Pharm*. 2023;14(11):Article 1030. doi:10.4236/pp.2023.1411030
15. Waltersdorfer L, Sabou M. Leveraging knowledge graphs for AI system auditing and transparency. *J Web Semant*. 2025;84:100849. doi:10.1016/j.websem.2024.100849
16. Pasas-Farmer S, Jain R. From discovery to delivery: governance of AI in the pharmaceutical industry. *Green Anal Chem*. 2025;13:100268. doi:10.1016/j.greeac.2025.100268
17. Challenges of using artificial intelligence (Natural Language Processing) tools for regulatory documentation. *J Appl Pharm Sci*. 2026 (ahead of print).
18. Shi Y, Wang J, Ren P, ValizadehAslani T, Zhang Y, Hu M, Liang H. Fine-tuning BERT for automatic ADME semantic labeling in FDA drug labeling to enhance product-specific guidance assessment. *J Biomed Inform*. 2023;138:104285. doi:10.1016/j.jbi.2023.104285
19. Kadi H, Abdellatif A, Njamen DIK, Pereme F. Automatic extraction of SmPC document for IDMP data model construction using foundation LLM and RAG: a preliminary experiment for pharmaceutical regulatory affairs. *Front Med (Lausanne)*. 2025;12:1598979. doi:10.3389/fmed.2025.1598979
20. Kim J, Hur M, Min M. From RAG to QA-RAG: integrating generative AI for pharmaceutical regulatory compliance process. *Pharmacol Pharm*. 2023;14(11):Article 1030. doi:10.4236/pp.2023.1411030
21. Waltersdorfer L, Sabou M. Leveraging knowledge graphs for AI system auditing and transparency. *J Web Semant*. 2025;84:100849. doi:10.1016/j.websem.2024.100849
22. IQVIA regulatory publishing and eCTD 4.0 automation analyses, as referenced in: Fu L, Jia G, Liu Z, Pang X, Cui Y. The applications and advances of artificial intelligence in drug regulation: a global perspective. *Acta Pharm Sin B*. 2025;15(1):1-14.
23. Method and system for automating clinical data standards [patent/technical disclosure]. U.S. Patent and Trademark Office; filed for automation of SDTM artefact generation using ML and NLP.
24. Schmider J, Kumar K, LaForest C, Swankoski B, Naim K, Caubel PM. Innovation in pharmacovigilance: use of artificial intelligence in adverse event case processing. *Clin Pharmacol Ther*. 2019;105(4):954-961. doi:10.1002/cpt.1255

25. Salas M, Petracek J, Yalamanchili P, Aimer O, Kasthuril D, Dhingra S, Junaid T, Bostic T. The use of artificial intelligence in pharmacovigilance: a systematic review of the literature. *Pharmaceut Med.* 2022;36(5):295-306. doi:10.1007/s40290-022-00441-z
26. Ball R, Dal Pan G. "Artificial intelligence" for pharmacovigilance: ready for prime time? *Drug Saf.* 2022;45(5):429-438. doi:10.1007/s40264-022-01157-4
27. Kompa B, Hakim JB, Palepu A, et al. Artificial intelligence based on machine learning in pharmacovigilance: a scoping review. *Drug Saf.* 2022;45(5):477-491. doi:10.1007/s40264-022-01176-1
28. Kassekert R, Grabowski N, Lorenz D, Schaffer C, Kempf D, Roy P, et al. Industry perspective on artificial intelligence/machine learning in pharmacovigilance. *Drug Saf.* 2022;45(5):439-448. doi:10.1007/s40264-022-01164-5
29. Bate A, Luo Y. Artificial intelligence and machine learning for safe medicines. *Drug Saf.* 2022;45(5):403-405. doi:10.1007/s40264-022-01177-0
30. Askin S, Burkhalter D, Calado G, El Dakrouni S. Artificial intelligence applied to clinical trials: opportunities and challenges. *Health Technol (Berl).* 2023;13(2):203-213. doi:10.1007/s12553-023-00738-2
31. Alexander M, Solomon B, Ball DL, Sheerin M, Mullan ID, Preininger AM, Jackson GP, Herath DM. Evaluation of an artificial intelligence clinical trial matching system in Australian lung cancer patients. *JAMIA Open.* 2020;3(2):209-215. doi:10.1093/jamiaopen/ooaa002
32. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. ICH M15 Guideline on Model-Informed Drug Development (Draft, Step 2). Geneva: ICH; 2024 Oct.
33. Amann J, Blasimme A, Vayena E, Frey D, Madai VI. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Med Inform Decis Mak.* 2020;20:310. doi:10.1186/s12911-020-01332-6
34. Zhu Z. Intelligent information management enables quality-by-design in pharmaceutical production. *Sci Rep.* 2025;15:44201. doi:10.1038/s41598-025-27879-w
35. Duarte JG, Duarte MG, Piedade AP, Mascarenhas-Melo F. Rethinking pharmaceutical industry with quality by design: application in research, development, manufacturing, and quality assurance. *AAPS J.* 2025;27(4):96. doi:10.1208/s12248-025-01079-w
36. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. ICH Q13: Continuous Manufacturing of Drug Substances and Drug Products. Geneva: ICH; 2022 Nov.
37. U.S. Food and Drug Administration. Artificial Intelligence in Drug Manufacturing. Silver Spring (MD): FDA; 2023 Mar.
38. U.S. Food and Drug Administration. Artificial Intelligence and Medical Products: How CBER, CDER, CDRH, and OCP Are Working Together. Silver Spring (MD): FDA; 2024 Mar (revised 2025 Feb).
39. Federal Register. Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products; Draft Guidance for Industry; Availability. 90 Fed Reg (7 Jan 2025). Docket No. FDA-2024-D-4689. Available from: <https://www.federalregister.gov/documents/2025/01/07/2024-31542>
40. Niazi SK. A critical review of the FDA's draft guidance on artificial intelligence in drug and biological product regulation. *J Chem.* 2026;2026:Article 5202999. doi:10.1155/joch/5202999

41. U.S. Food and Drug Administration. Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions: Guidance for Industry and FDA Staff. Silver Spring (MD): FDA; 2024 Dec.
42. UK Medicines and Healthcare products Regulatory Agency. MHRA's Strategic Approach to Artificial Intelligence. London: MHRA; 2024 May.
43. Health Canada. Pre-Market Guidance for Machine Learning-Enabled Medical Devices (Draft). Ottawa: Health Canada; 2023.
44. U.S. Food and Drug Administration, Health Canada, UK Medicines and Healthcare Products Regulatory Agency. Good Machine Learning Practice for Medical Device Development: Guiding Principles. 2021 Oct.
45. Canada's Drug Agency (CDA-AMC). Position Statement on the Use of Artificial Intelligence in the Generation and Reporting of Evidence. Ottawa: CDA-AMC; 2025.
46. National Institute for Health and Care Excellence. Use of AI in Evidence Generation: NICE Position Statement. London: NICE; 2024. Available from: <https://www.nice.org.uk/about/what-we-do/our-research-work/use-of-ai-in-evidence-generation--nice-position-statement>
47. Pharmaceuticals and Medical Devices Agency (Japan). Points to Consider on the Handling of Changes to Artificial Intelligence-Based Software as a Medical Device (AI-SaMD). Tokyo: PMDA; 2023 Mar 30. Available from: <https://www.pmda.go.jp/files/000266100.pdf>
48. Pharmaceuticals and Medical Devices Agency (Japan). Overview of Post-Approval Change Management Protocol (PACMP) for Medical Devices. Tokyo: PMDA; 2020 Aug 31.
49. General Office of the State Council of the People's Republic of China. Implementation Opinions on Comprehensively Strengthening the Capacity Building of Drug Regulation [(2021) No. 16]. Beijing: State Council; 2021 Apr.
50. World Health Organization. Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models. Geneva: World Health Organization; 2024. Available from: <https://www.who.int/publications/i/item/9789240084759>
51. World Health Organization. Global Strategy on Digital Health 2020-2025. Geneva: World Health Organization; 2021.
52. Government of Canada, Innovation, Science and Economic Development Canada. Artificial Intelligence and Data Act (AIDA) – Companion Document. Ottawa: ISED; 2023.
53. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). Off J Eur Union. 2024 Jun 13.
54. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. ICH E6(R3) Guideline for Good Clinical Practice. Geneva: ICH; 2023.
55. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations: Draft Guidance for Industry and FDA Staff. Silver Spring (MD): FDA; 2025 Jan.