



**INTERNATIONAL JOURNAL OF
PHARMACEUTICAL SCIENCES**
[ISSN: 0975-4725; CODEN(USA): IJPS00]
Journal Homepage: <https://www.ijpsjournal.com>



Review Article

Artificial Intelligence in Pharmaceuticals: Ethical Risks, Regulatory Challenges, and a Framework for Responsible Governance

A. Sangdokar*

Independent Researcher.

ARTICLE INFO

Published: 28 Oct 2026

Keywords:

artificial intelligence, pharmaceuticals, AI ethics, pharmaceutical regulation, responsible governance, patient safety

DOI:

10.5281/zenodo.23722087

ABSTRACT

Artificial intelligence (AI) is used in drug discovery, development, manufacturing, medication management, and patient care. These uses differ in supporting evidence and potential harm, so oversight must match their risks. This critical narrative review evaluates selected research, official US and Indian regulatory documents, and international and Indian ethical guidance. Concerns include biased or unrepresentative data, privacy, cybersecurity, unreliable outputs, limited transparency, divided responsibility, staff readiness, and unequal access. Much of the evidence comes from reviews, surveys, interviews, and limited performance tests. US Food and Drug Administration (US FDA) documents emphasize intended use, risk-based evidence of reliability, and lifecycle oversight. Discussion papers, draft guidance, and guiding principles differ from binding requirements. World Health Organization (WHO) and Indian Council of Medical Research (ICMR) guidance provide ethical foundations, while Indian medical-device rules and Central Drugs Standard Control Organization (CDSCO) software guidance connect intended use with regulation, local suitability, and monitoring. The review proposes seven governance stages: define intended use; assess risk and responsibilities; manage data and vendors; establish reliability; approve controlled introduction; monitor and manage changes; and address harm or end use. Each stage identifies responsibilities, evidence, and decisions. Assessment also considers human–AI performance, effects across groups, service resilience, and opportunities to question decisions. Responsible pharmaceutical AI requires continuing organizational oversight matched to effects on product quality, regulatory evidence, and patient care.

INTRODUCTION

Artificial intelligence is becoming part of how medicines are discovered, evaluated, manufactured, supplied, and used. Prediction models can help screen

*Corresponding Author: A. Sangdokar

Address: Independent Researcher.

Email ✉: ameysangodkar@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.



molecules and design formulations. Systems that use data can support manufacturing control and clinical development. Generative tools can produce drug-information summaries or explanations for patients. Reviews describe many pharmaceutical uses, but a promising use does not prove clinical effectiveness, value for money, or readiness for routine use. How well a tool performs technically, how useful it is in practice, and how much it benefits patients are separate questions that require different evidence.^{1,2,3}

This distinction is especially important in pharmaceuticals because people may give increasing weight to computer-generated results as those results move through an organization. A model's suggestion may influence which drug candidate receives resources for experiments. A predicted safety outcome may influence development decisions. A convincing medication recommendation may influence a pharmacist or patient. Errors can therefore affect both individual care and the evidence used to guide future care. Ethical assessment should examine what a model is meant to do and how its failure could cause harm, rather than treating all pharmaceutical AI as one category.^{4,5}

Governance covers more than the model itself. Those responsible for data, sponsors, manufacturers, software vendors, healthcare organizations, professionals, and regulators control different parts of a system. Responsibility can become unclear when a vendor supplies a tool whose workings are difficult to understand, an organization changes the way work is done, and a clinician is expected to detect the resulting errors. Habi and colleagues argue that responsibility and safety checks must reflect how control and understanding are divided between these groups. Their analysis concerns moral responsibility; it does not determine who is legally liable.⁴

This review examines three connected questions: which ethical risks matter most throughout the pharmaceutical lifecycle; why existing regulatory approaches struggle with AI; and how organizations can turn broad principles into recorded decisions that can be checked and reconsidered after a system enters use. It focuses on US FDA discussion papers and guidelines, the joint US FDA–European Medicines Agency (EMA) principles and World Health Organization

(WHO) guidance on large multimodal models (LMMs). Indian sources include the Indian Council of Medical Research (ICMR) AI ethics guidelines, the Medical Devices Rules, 2017, and Central Drugs Standard Control Organization (CDSCO) medical-device software guidance. These sources cover related ethical and product-regulatory questions, but they do not all apply to the same pharmaceutical uses. The proposed governance framework brings these ideas together for practical use.

MATERIALS AND METHODS

This article critically examines a collection of publications and regulatory materials from regulatory authorities in USA, Europe and India. It also draws on online searches for supporting original studies, reporting guidelines, and official governance documents through 17 October 2026. Official US FDA webpages were checked to confirm the status of key documents. This is a critical narrative review, not a systematic review. The collection was not assembled through a database search that others could repeat using selection criteria defined in advance. The review therefore does not claim to cover all relevant literature or provide a formal count of all records screened. It covers pharmaceutical research, development, manufacturing, and pharmacy practice. It does not compare all AI, privacy, professional-practice, or product-liability laws in every jurisdiction.

Source selection gave priority to original studies directly relevant to pharmacy, official regulatory documents, and reviews that explained AI uses across the lifecycle or problems with oversight. WHO's 2024 guidance was included for its international ethical recommendations on generative AI. A book covering AI ethics across disciplines was used selectively to explain bias concepts, not as evidence that pharmaceutical AI is effective. The Indian sources were ICMR's 2023 AI ethics guidelines, a compilation of the Medical Devices Rules, and CDSCO's 2026 software guidance. They were examined for what they cover, their legal or advisory status, and their recommendations for pharmaceutical governance. This selected set does not amount to a complete review of Indian law or all amendments to the Rules. The January



2026 US FDA–EMA principles were also examined.^{6,7} Findings were grouped by AI use, ethical risk, regulatory relevance, and practical ways to reduce risk. Source assessment considered study design, the sample or testing setting, clarity of reporting, whether important claims could be traced to evidence, and whether a source reported original research or summarized other work. Studies from wider healthcare and general AI research helped explain risks that may also occur in pharmaceuticals; they were not treated as direct evidence of pharmaceutical outcomes.

Four distinctions guide how the evidence is interpreted. Surveys describe participants' views; they do not

measure how often AI actually causes injury. Interviews identify experiences and areas of disagreement; they do not establish how common these are in the wider population or define official regulatory policy. Benchmark tests assess particular systems and tasks at a particular time; they do not prove that systems can safely work on their own in routine care. Finally, regulatory documents have different roles and legal status: an US FDA discussion paper is not guidance, and draft guidance is not a binding rule. ICMR ethical recommendations and CDSCO explanatory guidance must also be distinguished from Indian statutory rules. Table 1 summarizes the original studies that provide evidence for the discussion.^{8,9,10,11,12}

Table 1: Selected Original Studies and Limits on Their Interpretation

Study	Design and scope	Contribution to this review	Main limitation
Hasan et al., 2024 ⁸	Online survey; 501 pharmacy professionals; seven Middle East and North Africa (MENA) countries	Concerns about privacy, cybersecurity, employment, and regulation	Convenience sample and self-reported views; concerns do not measure actual harm rates
Hasan et al., 2024 ⁹	Survey; 875 pharmacy students and faculty; six countries	Gaps in AI knowledge and formal education	One-time educational survey; cannot show workplace competence
Li et al., 2025 ¹⁰	Eight generative systems; 48 pharmacy questions; six pharmacist raters	Errors in specific tasks and the need to test the intended use	Tests on one day; limited questions; no assessment of patient outcomes
Gregory and Austin, 2025 ¹²	Interviews with 12 pharmacy regulators in Canada and the US	Principles, professional oversight, ability to withstand disruption, and limits of regulatory authority	Explores participants' views; not enacted policy or national agreement
Austin and Gregory, 2026 ¹¹	Interviews with 24 Ontario community pharmacists	Informal use and demand for practical guidance	Local, self-reported experience; may not apply widely

RESULTS

3.1. Applications and the consequences of failure

Pharmaceutical AI should first be grouped by what it does and how much an important decision depends on it. In drug discovery, a model may rank compounds for later laboratory testing.^{13,14,15} Its immediate influence is limited when independent experiments still determine the outcome. However, biased selection can still waste resources or repeatedly overlook certain targets. In manufacturing, a model that influences process control or release testing can directly affect product quality. In pharmacy, a tool may help a professional, communicate with a patient, or trigger an action with little chance to correct it. These settings need different types

and levels of safety checks and supporting evidence.^{1,5,16}

Reviews of generative AI in drug development discuss possible savings in time and resources. However, faster computation does not automatically mean a shorter overall development process or greater clinical success. Experimental confirmation, unsuccessful candidates, work to integrate the system, data preparation, and continued monitoring still matter. Similarly, robotics and automated dispensing do not always involve AI. A process that follows fixed instructions and a prediction system that learns from data can fail in different ways. Treating them as the same can hide both why they are useful and which safeguards they need.^{2,17}



The randomized phase 2a trial of rentosertib shows that a discovery program involving AI can reach clinical testing. On its own, it does not prove that this approach is better than conventional drug development.¹⁸

Table 2 gives examples of how harm could occur. These examples support analysis; they are not formal

US FDA risk classifications. Risk depends on the actual use, the availability of independent evidence, whether decisions can be reversed, the population affected, and the consequences of error. Even an apparently administrative system may affect clinical care if its output determines who can obtain a medicine or which safety reports are reviewed.

Table 2: Examples of Risks and Priorities for Checks and Supporting Evidence

Application	How failure could cause harm	Priority checks and evidence
Discovery and formulation	Biased or unreliable predictions direct experiments toward unsuitable priorities	Data sources and history, limits of use, reproducibility, independent experiments
Clinical development	Unsuitable predictions or selection methods distort trial evidence	Defined use, evaluation outside development settings, subgroup analysis, traceable decisions
Manufacturing	Model errors or unnoticed changes weaken process control	Process-specific validation, independent controls, change management within the quality system
Safety surveillance	Missed signals or too many false alerts hinder review	Sensitivity and workload assessment, referral to a human reviewer, traceable case handling
Medication management	Incorrect doses, contraindications, or advice unsuitable for the situation reach patient care	Local clinical evaluation, verified sources, meaningful professional review
Operational allocation	Forecasting errors cause shortages or unequal access	Monitor service levels and effects on different groups; plan for disruptions

3.2. Ethical risks across the pharmaceutical lifecycle

3.2.1. Bias, population representation, and unequal benefit

AI systems reflect limitations in the data and objectives used to build them. In pharmacogenomics and individualized therapy, missing or poorly represented populations can make predictions less suitable for those groups. Differences in coding, language, coexisting illnesses, access to testing, and prescribing patterns can also affect whether a model works in another clinical setting. The ethical problem goes beyond balancing demographic groups in training data. It also concerns whether the model addresses the right question and whether people have fair access to its possible benefits.^{19,20}

Medical-imaging research shows that models can infer race even when explicit demographic labels are not provided as inputs. Removing these labels alone therefore does not establish fairness.²¹

Oversight should assess fairness in relation to the particular decision. For a medication-risk model, questions include which groups have more missed high-risk

cases, whether results are more uncertain for poorly represented patients, and what happens when the model cannot give a reliable recommendation. For a tool that allocates resources, unequal access or delays may matter as much as prediction accuracy. These are proposed priorities for assessment. This article does not establish one measure of fairness that resolves all these competing concerns.

Overall performance can hide clinically important failures in small groups.²² However, estimates for groups with very few observations may also be unreliable. Reports should therefore identify gaps in representation, give estimates of uncertainty where appropriate, and state when evidence is insufficient for a proposed use. Limiting use until more evidence is available may be more justified than treating an uncertain comparison between groups as proof of fairness. This approach agrees with the US FDA–EMA principles' focus on the intended setting, relevant data, and performance assessment based on risk.²³

Bias assessment should also consider groups with overlapping characteristics, such as older patients with disabilities or patients whose preferred language is



poorly represented. Such analysis requires enough data and a justified purpose. Rane and colleagues distinguish ways to address bias during data preparation, model training, and the handling of model outputs. These categories help describe possible controls, but they do not prove that a particular control improves pharmaceutical outcomes.⁷

Measures of fairness answer different questions. Demographic parity asks how often a model gives a positive prediction in different groups. Equalized odds compare true-positive and false-positive rates between groups. Calibration asks whether predicted risks match observed outcomes. These goals can conflict under some data and modelling conditions. They should not be treated as identical requirements or as goals that always conflict. For a medication-risk tool, the organization should explain why its chosen measure addresses the harm it wants to prevent. It should report clinically relevant error rates and uncertainty, and check whether measures to reduce bias weaken other safety requirements. The proposed framework uses these concepts to guide assessment; it does not set one universal cutoff for fairness.⁷

ICMR also emphasizes training data that represent the intended population, independent checks of algorithms, and protection of poorly represented groups. Its validation guidance recommends evaluation using independently collected data and local assessment before a system is used in healthcare. These recommendations support testing whether a system suits the intended population rather than relying on its overall performance elsewhere.²⁴ (sections 1.9–1.10 and 2.2–2.3)

3.2.2. Privacy, data control, and cybersecurity

Pharmaceutical AI may use sensitive clinical, genomic, commercial, or research data. Risks can arise when data is collected, transferred, used to develop models, recorded in logs, accessed for vendor support, or reused for another purpose. Hasan and colleagues found that 58.9% of 501 surveyed pharmacy professionals were concerned about patient-data privacy. The same proportion were concerned about cybersecurity. These figures show that the issues mattered to the people surveyed. They do not measure how often breaches

occur or prove that AI is less secure than existing alternatives.⁸

Permission to access data should be distinguished from permission to reuse them for another purpose.²⁵ Research showing that training data can be extracted from language models demonstrates another risk of information disclosure. It does not demonstrate a pharmaceutical data breach.²⁶ The proposed assessment should record how data will move, who can access them, how they will be kept, what vendors may do with them, and who responds to incidents. Explanations for patients should be understandable and match the role AI plays. Consent and notification requirements depend on the activity and the applicable law. Ethical transparency requires more than applying the same consent checkbox to every situation.

ICMR connects data use with informed consent, confidentiality, whether individuals can be identified, and ethics review. It discusses waivers of consent for research using existing anonymized data. This does not provide a general exemption from ethics review or applicable law. Governance records should distinguish permission for the original data collection from permission for the proposed AI use, and explain the basis for any waiver.²⁴ (sections 2.1 and 4)

Security and safety are connected. A compromised data source, changed model version, or unavailable service can alter a recommendation or interrupt care even when information has not been disclosed.²⁷ In Gregory and Austin's interviews, regulators linked privacy, the integrity of information, backup capacity, and the ability of systems to work together. These findings support including the ability to withstand disruption and dependence on vendors in decisions about adoption. They do not establish that any particular security design is effective.¹²

3.2.3. Unreliable outputs and misleading confidence

Generative systems pose a particular risk when fluent writing is mistaken for verified pharmaceutical knowledge. An answer can sound reasonable and consistent while missing a contraindication, applying a guideline incorrectly, or overlooking local practice. Li and colleagues tested eight systems using 48 clinical-pharmacy questions. They reported weaknesses involving



contraindications, adaptation to local settings, and complex reasoning. The study shows possible ways systems can fail. However, its limited questions and testing on a single day cannot establish a lasting ranking of commercial systems or estimate actual patient injuries.¹⁰

The key question is whether the system reliably supports the decision it is meant to inform. A drug-information assistant should be checked for factual accuracy, clinically important omissions, accurate use of sources, and up-to-date information. It should also be assessed for whether it appropriately declines to answer or refers the question for further review when information is insufficient. For individual recommendations, the patient's circumstances and the way the tool interacts with professional judgment also matter. A convincing explanation does not by itself prove that an answer is correct or that the explanation accurately describes the model's reasoning.

Retrieving information from carefully selected sources and requiring professional review are useful design options, but their effectiveness must be tested. Reviewers need access to the evidence, enough time to assess it, and clear authority to reject the output. Where errors could have serious consequences, the organization should explain why the remaining risk is acceptable and how errors missed during review will be detected. This links AI hallucinations, or plausible but incorrect outputs, to the wider need for continuing safety checks.^{4,5}

3.2.4. Responsibility and the limits of human oversight

Responsibility is often addressed by keeping a human "in the loop." However, a formal approval step may provide little protection if the person cannot properly assess the recommendation. A professional may have authority on paper while facing time pressure, poor explanations, incomplete patient data, or organizational pressure to follow the system. Habli and colleagues explain why responsibility depends on control and understanding, not only on who takes the final action.⁴

Experimental research on diagnostic advice found that incorrect advice could negatively influence physicians' judgments. This supports testing how people and AI

perform together rather than assuming that human review ensures safety.²⁸

The interview studies reveal a practical difference in expectations. Ontario pharmacists wanted clearer direction and support. The regulators interviewed in Canada and the United States preferred guidance based on principles and emphasized the limits of professional regulation. These findings should not be applied to all pharmacists or regulators. Taken together, however, they suggest a possible gap between the reassurance practitioners want and what professional regulators believe they can provide.^{11,12}

The proposed response is to share responsibility while assigning each duty clearly. Developers should record known limitations and changes. Organizations using AI should assess local suitability and provide safe working procedures. Professionals should use their judgment within roles for which they receive support. Those responsible for governance should keep the authority to suspend use. Each important safeguard needs a named person responsible for it. Shared responsibility does not work when everyone can blame another party for a failure that could have been foreseen.

ICMR's guidance on introducing AI into use calls for human oversight and the ability to override the system manually where AI could cause severe or bodily injury. It also calls for an alternative way to provide care if the system fails. Its ethics-review section recommends access to expertise in AI, data science, and relevant law. These recommendations support informed oversight and workable backup arrangements. They do not, by themselves, legally assign liability.²⁴ (sections 2.3 and 3)

3.2.5. Staff capability, access, and sustainability

Using AI requires more than knowing its terminology. In a survey across several countries involving 875 pharmacy students and faculty members, Hasan and colleagues reported an average knowledge score of 42.3 out of 100. Only 15.3% reported coursework in AI or machine learning. These findings show an educational gap in that sample; they do not measure the readiness of all pharmacists. Training should cover assessing evidence, recognizing missing context,



protecting data, reporting incidents, and retaining skills needed when the system is unavailable.⁹

Concerns about employment also need careful interpretation. In the ethics survey, 62.9% were concerned about job displacement. This is not a prediction of actual job losses. However, changing staff roles can affect who performs difficult checking tasks, who is responsible, and whether efficiency gains improve care or simply increase workloads. Assessment should consider staff workload, freedom to exercise professional judgment, and the quality of communication with patients, as well as productivity.⁸

Access and sustainability matter even when they do not directly affect model performance. Subscription costs, infrastructure needs, and dependence on vendors may exclude smaller organizations or leave services vulnerable to disruption. Reviews of AI adoption and robotics discuss resource and implementation barriers. However, the supplied evidence does not support a reliable estimate of overall net costs or environmental benefits across the sector. Organizations should assess total operating costs and resource use rather than assume that faster computing or fewer laboratory activities make a system sustainable.^{17,29}

3.2.6. Research integrity and pharmaceutical evidence

AI-generated literature summaries, manuscripts, and regulatory documents can spread incorrect references or unsupported claims if sources are not checked carefully.³⁰ A published letter by Glynn raises concerns about suspected use of generative AI without disclosure. It expresses concerns about transparency and integrity; it does not prove misconduct in every case discussed. The wider lesson for governance is that

important claims should have clearly identified sources and a clear record of verification.³¹

Pharmaceutical research must distinguish measured results, model-generated estimates, and the author's interpretation. AI assistance should be disclosed as required by the publication or submission process, and identifiable people should take responsibility for the final document. A convincing literature summary cannot replace checking the original evidence, especially when it supports claims about effectiveness, safety, or regulatory duties.

3.3. Regulatory challenges and emerging responses

3.3.1. What the documents cover and their regulatory status

AI may be used to develop a medicine, support manufacturing, help professional practice, or form part of a separately regulated product. These roles should be distinguished. The applicable regulatory process depends on the intended function and setting, not simply on whether the technology uses machine learning. US FDA's cross-center paper describes how its medical-product centers work together and identifies priorities for regulatory approaches, standards, and research. It sets out a strategy; it does not provide a general route to approval for all pharmaceutical AI.³²

The selected documents also show why it is inaccurate to say that pharmaceutical AI has no regulatory oversight. US FDA's manufacturing discussion considers AI within existing quality and oversight systems. The challenge is to decide what evidence and controls are needed to meet established responsibilities for each AI use. Existing oversight does not mean that all new ways in which AI can fail have already been addressed.¹⁶

Table 3 distinguishes selected US, joint US FDA–EMA, and Indian documents in the collection.

Table 3: Selected Regulatory and Ethical Documents and HOW to Interpret Them

Document	Status and scope	Contribution to oversight	Interpretation to avoid
Artificial Intelligence in Drug Manufacturing, 2023 ¹⁶	Discussion paper; not draft or final guidance	Questions about oversight, data, validation, and changes during use	Treating discussion questions as new binding requirements
Considerations for AI to Support Regulatory Decision-Making, 2025 ⁵	Draft guidance; information supporting regulatory decisions on drugs and biologics	Assessing reliability for a particular use according to risk	Calling it final guidance or approval of all AI tools



Document	Status and scope	Contribution to oversight	Interpretation to avoid
AI & Medical Products, revised 2025 ³²	Strategy across US FDA centers	Coordination, regulatory approaches, standards, and research priorities	Assuming all planned activities have been carried out
US FDA–EMA Good AI Practice principles, 2026 ²³	Joint guiding principles for drug development	Design around people's needs, data management, performance, and lifecycle oversight	Treating the principles as a separate law
External Engagements with US FDA ³³	Webpage explaining opportunities to engage with US FDA	Ways to discuss particular development questions	Treating participation as product approval or proof of compliance
ICMR AI ethics guidelines, 2023 ²⁴	Ethical guidance for biomedical research and healthcare	Consent, fairness, validation, ethics review, and human oversight	Treating ethical guidance as an AI law or product authorization
Medical Devices Rules, 2017 ³⁴	Legally binding rules for products they cover	Risk classification, licensing, and quality systems	Applying medical-device rules to every pharmaceutical AI use
CDSCO medical-device software guidance, 2026 ³⁵	Explains existing device rules	Software evidence, suitability for India, change control, and monitoring	Treating it as a new, separate AI law

3.3.2. Reliability must be assessed for the specific use

The January 2025 US FDA draft covers AI used to produce information or data that support regulatory decisions about drug safety, effectiveness, or quality. It proposes seven steps: define the question; define how and where the model will be used; assess the model's risk; plan how to assess its credibility, or reliability for that use; carry out the plan; record results and departures from the plan; and decide whether the evidence is sufficient for the intended use. Risk depends on both the model's influence on a decision and the consequences of a wrong decision. This assessment depends on the situation; it is not one numerical formula for all risks. ⁵

A model that passes one validation exercise is not therefore reliable for every purpose. The same output may support a decision in one process but determine it in another. The strength of a claim depends on whether the data are relevant, how testing is designed, whether testing is independent, how uncertain the results are, and how the model's output relates to other evidence. The draft does not cover every pharmaceutical AI activity. It excludes certain discovery uses and operational uses that do not affect patient safety, product quality, or the reliability of relevant studies. Exclusion from the draft does not mean that a use has no ethical risks. ⁵

The draft gives a useful structure for deciding what evidence and checks are needed, but it is not a complete ethical framework. Privacy, cybersecurity, access, effects on staff, and ways to address harm need further attention. The framework below therefore builds on the draft's approach to reliability while also considering effects on organizations and different groups of people.

3.3.3. Manufacturing, vendors, and change control

The 2023 manufacturing discussion paper raises questions about cloud computing and oversight of external providers, the amount of data and how long they are kept, how existing oversight applies, validation, and changes throughout a system's life. These questions are connected. A manufacturer may depend on a vendor whose model changes, even though the evidence supporting process control relates to an older version. It is therefore essential to keep a clear, auditable link between the system in use and the evidence supporting it. ¹⁶

Changes involve more than retraining a model. Changes to input definitions, instruments, how data are collected and processed, information sources, user interfaces, or working procedures can alter system behaviour and its effects. For generative systems, prompts and available knowledge sources are also part of the setup being evaluated. The proposed response is to record that setup, assess the effects of changes, and decide how much reassessment is needed before the system informs important decisions. Required

regulatory reports or submissions depend on the product and the change. This review does not suggest that every update requires a new US FDA application.⁵

3.3.4. Shared principles do not mean identical requirements

The joint US FDA–EMA principles provide common ideas: design around people's needs, assess risk, use standards, define the context of use, involve experts from several fields, manage data responsibly, develop models soundly, assess performance, manage the full lifecycle, and provide essential information. Their value is in helping developers and regulators ask similar questions. They do not make legal duties or evidence requirements identical across jurisdictions.²³

EMA's 2024 reflection paper also addresses AI throughout the life of a medicinal product, linking oversight to the intended use and risk. It is a scientific reflection paper, not a separate AI law.³⁶

US FDA identifies opportunities for discussion through programs and meetings suited to particular development questions. Early discussion may help clarify the intended use or the plan for gathering evidence. However, joining a program does not mean that an AI system is approved or that compliance is assured. Professional regulation, product oversight, data protection, and organizational governance may each cover different parts of the same AI use. The continuing challenge is to coordinate them without leaving gaps in responsibility.^{12,33}

3.3.5. WHO ethical principles and international governance

WHO's 2024 guidance, *Ethics and governance of artificial intelligence for health: Guidance on large multimodal models*, adds to the US FDA–EMA principles' focus on drug development. It covers generative models in health and medicine and repeats six ethical principles from WHO's earlier guidance: protect people's autonomy, or ability to make their own decisions; promote well-being, safety, and the public interest; provide transparency and understandable explanations; ensure responsibility and accountability; support inclusion and fairness; and remain responsive and sustainable. These principles connect technical checks with patients' rights and how benefits and burdens are

shared. They are ethical recommendations, not a replacement for the law in a particular jurisdiction.^{6,37} (WHO 2024, pp. 4–5)

WHO recommends assessing impacts throughout the lifecycle, covering ethics, human rights, safety, and data protection, with independent examination and public reporting. For large-scale use of LMMs, it recommends that governments introduce independent audits and impact assessments after release, including results reported separately for relevant user groups. It also distinguishes those who develop, provide, and put systems into use. An organization can create risk by using a system inappropriately even when it did not build the model. These recommendations support separate responsibilities for evaluation, clear information about limitations, and explicit authority to suspend use.⁶ (pp. 51, 54–57)

The guidance calls for more than technical checks by experts. It discusses participation by patients and health professionals, training to reduce excessive reliance on automated advice, understandable risk information, affordability, and availability in relevant languages. It also considers the concentration of technological power and the use of carbon and water resources. This broadens oversight beyond the accuracy of individual outputs. In pharmaceuticals, these issues are especially relevant to assistants used by patients and information tools used by professionals. Applying them elsewhere, such as manufacturing control, requires adjustment to the decision and the people affected.⁶ (pp. 24–29, 55–57)

WHO's 2023 *Regulatory considerations on artificial intelligence for health* offers additional points for regulatory discussion. It explicitly states that it is not guidance, a regulatory framework, or policy.³⁸

3.3.6. Indian ethical guidance and medical-device rules

The documents relating to Indian regulatory authorities show why ethical oversight and product regulation should be distinguished. ICMR's 2023 guidelines cover AI in biomedical research and healthcare through principles centred on patients, guidance on development and validation, ethics review, and informed consent. They provide an ethical basis for responsible



research and healthcare use. They do not themselves authorize a product or establish a comprehensive law for pharmaceutical AI. Historical references to data-protection proposals in the document should not be treated as current law.²⁴

For products that they cover, the Medical Devices Rules, 2017 provide the legal framework. Rule 4 sets out four risk classes, A through D. The First Schedule links classification to intended purpose. It states that software driving a device or influencing its use falls in the same class as that device. The Fifth Schedule sets out quality-management requirements. This does not mean that every pharmaceutical prediction model or administrative tool is a medical device. Whether a product falls within the Rules must be established before device-specific duties are applied.³⁴ (rule 4; First and Fifth Schedules)

CDSCO's Guidance Document on Medical Device Software, identified as CDSCO/MD/GD/MDSW/01/2026, explains existing oversight under these Rules. Its notice refers readers to the legal provisions, and its scope section clearly states that it does not introduce a new regulatory control. It links software oversight to the intended medical purpose and covers classification, quality systems, submissions, and activities after market introduction. This matters when comparing it with US FDA drug-development guidance. CDSCO's document concerns software that qualifies as a medical device. US FDA's draft concerns AI-generated information used to support regulatory decisions on drugs and biological products.^{5,35}

The CDSCO guidance sets out specific expectations relevant to responsible AI. For models trained or validated outside India, their suitability for Indian clinical settings must be justified. It describes disclosure of what training, validation, and testing datasets contain, including demographic, geographic, and clinical diversity. It also describes evidence addressing bias, whether results apply beyond the original setting, and reliable performance across relevant Indian population groups. Its discussion of risk recognizes indirect harm from misleading or delayed software information. It describes an Algorithm Change Protocol that may be developed where applicable. Monitoring after market introduction includes performance, bias, changes in

data or performance over time, cybersecurity, complaints, and corrective action. These provide useful foundations for practical oversight, but they apply within the document's medical-device scope.³⁵ (pp. 42–44, 49, 56–57)

DISCUSSION

4.1. A proposed framework for responsible governance

4.1.1. Framework design and oversight matched to risk

The proposed framework combines the US FDA draft's approach to reliability for a specific use, the US FDA–EMA principles, WHO and ICMR ethical guidance, relevant CDSCO software controls, reasoning about organizational safety, and evidence from people involved in pharmacy. It brings these ideas together specifically for this review and has not been tested prospectively in practice. Its contribution is to connect each lifecycle stage with a responsible person, recorded evidence, and a point at which a decision must be made. For example, transparency becomes practical when an organization states what information it will provide, to whom, when, and how it will judge whether that information is sufficient.^{4,5,6,12,23,24,35}

This lifecycle approach also agrees with the FUTURE-AI healthcare consensus and the voluntary NIST AI Risk Management Framework and its generative-AI profile, which apply across sectors. None of these documents test or endorse the seven-stage framework proposed here.^{39,40,41}

The amount of evidence and checking required should match the risk. A drafting tool whose errors would have limited consequences may need basic restrictions on data use, checks of its outputs, and a clearly responsible person. A system that strongly influences an outcome used in regulatory assessment, product release, or medication decisions needs stronger independent evaluation, a recorded decision accepting the remaining risk, and active monitoring. The framework does not give each technology a permanent risk class. Oversight depends on how its output affects decisions, the consequences of those decisions, and how well independent safeguards work.



Table 4: Proposed Framework for Oversight Throughout the Lifecycle

Sr. No.	Stage	Responsible lead*	Main evidence	Decision to make
1	Define intended use	Scientific or clinical lead	Statement of use, users, population, excluded uses, applicable regulation	Is AI suitable for the problem?
2	Assess risk and assign responsibility	Governance lead with subject experts	How harm could occur, affected groups, assigned duties, oversight process	Are consequences and responsibilities clear?
3	Manage data and vendors	Data lead with privacy/security leads	Data sources and history, dataset composition, permissions, vendor terms, version record	Are data and external dependencies acceptable?
4	Establish reliability for the intended use	Evaluation lead with independent review	Independent test data, local suitability, subgroup results, testing of people and AI together	Is there enough evidence for this use?
5	Approve controlled introduction	Clinical or quality-system authority	Training, operating procedure, override and backup arrangements, decision on remaining risk	Can use proceed safely within defined limits?
6	Monitor and manage changes	Operational lead with quality support	Performance and bias trends, incidents, changes over time, recorded change assessment	Continue, limit, reassess, or suspend use?
7	Address harm or stop using the system	Responsible executive with clinical/quality leads	Impact review, corrective action, communication, continuity and record-retention plan	Has harm been addressed, and is closure justified?

*These roles describe proposed responsibilities, not job titles required by law. Organizations may combine roles but should preserve enough independence for important decisions. The records and depth of review should match the risk.

Table 5 shows the ethical and regulatory sources behind the framework. The proposed controls are this review's practical interpretation of the cited principles.

The organizations cited have not endorsed the seven-stage framework. WHO's recommendations concern health uses of LMMs, while the joint principles concern AI throughout drug development. ICMR provides biomedical and healthcare ethics guidance. CDSCO provides controls for software that falls within medical-device regulation. These differences in coverage should remain clear when choosing controls.^{6,23,24,35}

Table 5: Linking Ethical Principles to Proposed Oversight Measures

Ethical aim	Source of the principle	Proposed measure	Framework stage
Autonomy and informed use	WHO autonomy; US FDA–EMA 1 and 10; ICMR consent and autonomy	Explain AI's role clearly, provide meaningful review, allow important outputs to be questioned	1, 5, 7
Safety and public benefit	WHO safety; US FDA–EMA 2, 4 and 8; ICMR validation	Define use, match independent evaluation to risk, test how people and AI perform together	1, 2, 4
Transparency and traceability	WHO transparency; US FDA–EMA 6, 7 and 10; CDSCO dataset disclosure	Record data sources and history, analysis decisions, limitations, and important updates	3, 4, 6
Accountability and remedy	WHO accountability; US FDA–EMA 5 and 9; ICMR oversight and remedy	Name responsible people, use independent checks when needed, refer concerns and correct problems	2, 6, 7
Fairness and participation	WHO equity; US FDA–EMA 1 and 8; ICMR fairness; CDSCO Indian applicability	Justify subgroup testing, use accessible languages, involve patients and professionals	2, 4, 5, 6



Ethical aim	Source of the principle	Proposed measure	Framework stage
Sustainability and responsiveness	WHO sustainability; US FDA–EMA 9; CDSCO post-market monitoring	Assess resource needs, schedule monitoring, reassess periodically, plan for ending use	2, 6, 7

4.1.2. From intended use to reliable evidence

The first stage should produce a short statement that can be tested. It should explain who uses the system, what it produces, which decision it informs, which population or process it covers, and what it must not do. It should also record the applicable oversight process, including whether the system is a regulated product. If this is unclear, suitable regulatory experts should resolve the question before the system enters use. Claims of benefit should identify what the system is compared with and which meaningful outcome will be measured. Saying that it “improves efficiency” is not enough. The claim should explain whose time is saved, whether checking creates more work elsewhere, and whether quality is maintained.

The second stage identifies foreseeable harms and assigns responsibility for preventing, detecting, and responding to them. Assessment should consider how serious an error could be, how much influence the system has, whether decisions can be reversed, which groups are affected, and how effective human review is in practice. Patients, pharmacists, manufacturing specialists, or other affected people should be involved when their experience helps explain the risk. For relevant health applications, assessment should explicitly cover autonomy, human rights, data protection, access, and effects on different groups, following WHO's recommended approach to impact assessment.⁶ It should also cover transfers of work or responsibility to vendors and external services, rather than stop at the organization's boundary.

The third stage checks whether the data and external services or suppliers are suitable. At a minimum, records should describe the sources and history of data, what was included or excluded, missing data, coverage of the relevant population or process, permitted uses, and limitations. Vendor assessment should address access to performance information, notice of important changes, cooperation during incidents, continuity of service, and arrangements for exporting data or moving

to another system. These are proposed safeguards when obtaining a tool or service. Which safeguards are needed and how they are applied should follow the risk assessment, rather than one identical checklist for every tool.

The fourth stage tests the complete system for its intended use. The testing plan should be set before testing begins. It should use relevant independent data where possible, include meaningful comparisons, and examine errors with important consequences rather than rely only on an average score. External evaluations of pneumonia detection and sepsis prediction show why performance must be checked beyond the settings in which a model was developed.^{42,43} Assessment should also test how users work with the system. Can professionals identify incorrect outputs? Is uncertainty explained clearly? Does referral for further review work under realistic conditions? Records should distinguish developer testing, independent review, and evaluation of local working procedures. For large-scale use of LMMs in healthcare, independent audits and public summaries of impact assessments can put WHO's recommendations for external scrutiny into practice, while protecting sensitive information. These are proposed governance arrangements, not claims that every jurisdiction currently requires them. In India, assessment should explain why the evaluation population, language, clinical working procedures, and setting are relevant. For software that qualifies as a medical device, this also reflects CDSCO's emphasis on suitability for Indian settings and reliable performance across population groups.^{24,35} Gaps in evidence should lead to more testing, narrower claims, limited use, or a decision not to introduce the system.⁶

Reporting should suit the study design. SPIRIT-AI covers trial protocols, CONSORT-AI covers trial reports, and DECIDE-AI covers early clinical evaluation of decision-support systems.^{44,45,46} TRIPOD+AI covers clinical prediction models, TRIPOD-LLM covers studies using large language models, and STARD-AI



covers diagnostic-accuracy studies where relevant.^{47,48,49} For research on prediction models, PROBAST+AI can help assess quality, risk of bias, and whether the findings apply to the intended use. Complete reporting alone does not establish safety.⁵⁰

4.1.3. Introducing systems, monitoring them, and addressing harm

Approval to use a system should identify the person responsible for accepting the risks that remain after safeguards are applied. Those risks should be recorded. For uses with important consequences, a limited introduction may be appropriate before wider use, provided there are clear safeguards and it does not replace necessary testing beforehand. Staff should know the approved uses, prohibited uses, checking requirements, and how to report concerns. Backup arrangements must work in practice. An instruction to return to manual work is not enough if staff numbers, access to information, or skills make this impossible.

Monitoring should lead to action. Relevant indicators may include serious output errors, missed high-risk cases, performance in different groups, decisions to override the system and the reasons, workload caused by alerts, downtime, unauthorized data use, or changes in the pattern of input data. Not every indicator suits every use. Action thresholds and review intervals should be justified for the particular application. The literature does not support one set of numerical cutoffs for all systems. A serious safety event, unacceptable performance in a population group, an important change that has not been reviewed, or an inability to reconstruct an output used for an important decision should trigger urgent review. Use should be suspended when justified.

A change record should identify the model, data-processing steps, knowledge sources, and working procedures affected. It should describe the expected effect, the evidence needed before release, and how to return to the previous version if necessary. CDSCO describes an Algorithm Change Protocol for medical-device software where applicable. This review adapts that planning idea for governance. It does not assume that such a protocol gives advance regulatory permission for every update.³⁵

Addressing harm is part of being accountable. An organization should be able to identify affected decisions or records, correct wrong outputs where possible, support clinical or quality follow-up, and communicate through appropriate channels. Users and affected patients need a practical way to question important decisions. When a system is retired, plans should cover continuity of service, keeping necessary evidence, removing access, and deciding what happens to retained data. These are proposed governance responsibilities. How they must be carried out legally depends on the setting.

4.1.4. Examples in pharmacy and manufacturing

Consider a generative assistant that drafts medication-counselling text. Its use could be limited to preparing drafts from approved sources for a pharmacist to review, without independently prescribing for individual patients. Testing would examine clinically important omissions, whether the language is accessible, whether sources are accurate, and whether pharmacists detect errors deliberately inserted into realistic test situations. Monitoring would track harmful inaccuracies and the work required for review. Changes to the model or knowledge sources would trigger assessment of their effects. These safeguards respond to the failures found in Li and colleagues' limited evaluation and the guidance needs described by Austin and Gregory. They do not prove that this hypothetical system is safe.^{10,11}

For this counselling example in India, assessment should also examine relevant languages and local medication practices. Whether the tool is covered by medical-device oversight must be assessed from its intended function and claims. Pharmacist review alone does not replace that assessment.^{24,35}

For a manufacturing model that supports process control, the evidence and safety checks would focus on whether the data and testing represent the process, whether measurements are reliable, the operating limits, independent controls, and how the model fits within the pharmaceutical quality system. A vendor change that alters model behaviour would need assessment against the validated process and the applicable change-control procedure. Both examples follow the same governance stages, but need different evidence,



responsible people, and criteria for approval. This is the practical value of a framework based on the intended use.^{5,16}

4.2. Research priorities and limitations

The most urgent gap is evidence showing whether governance measures improve real-world outcomes. Studies that follow systems over time should assess medication errors, clinically important omissions, quality failures, unequal effects, workload, and the cost of maintaining checks and supporting evidence. Comparisons should distinguish how a model performs alone from how people perform when working with it. Training studies should test whether users can detect and manage failures, rather than measure confidence alone. Evaluations should clearly report system versions, intended use, the setting and nature of the data, and changes, so that findings can be understood and studies repeated.

The proposed framework also needs testing. Pilot studies at several sites could assess whether it is practical, the burden on staff, consistency of risk decisions, how completely incidents are detected, and how quickly corrective action is taken. Patients and practitioners should help assess whether communication and ways to challenge decisions are usable. Implementation studies should also examine whether smaller organizations can maintain the required oversight. Even an effective framework could increase inequality if its costs prevent some organizations from adopting AI responsibly.

CONCLUSION

AI offers opportunities across pharmaceuticals. Whether its use is ethically acceptable depends on the decision it influences, the supporting evidence, and the ability to detect and address failure. The reviewed literature supports concerns about limited data, unreliable outputs, unclear responsibility, staff readiness, and unequal access. It provides direct evidence of long-term benefits for patients or organizations. US FDA materials increasingly emphasize reliability for a specific use and oversight throughout the lifecycle. Their different legal or advisory status and areas of coverage must remain clear. WHO and ICMR guidance strengthen the ethical foundation. Indian medical-

device rules and CDSCO software guidance show how intended use, suitability for local settings, and continued oversight connect to regulation of particular products. Responsible governance should combine these regulatory approaches with meaningful human review, clear responsibility for data and vendors, assessment of effects on different groups, the ability to withstand disruptions, and action to address harm. The proposed seven-stage framework provides a practical structure for this work and a basis that future implementation studies can test.

Acknowledgment

There was no support or assistance received from any institution or other individuals in preparing this review article.

REFERENCES

1. Das S, Dey R, Nayak AK. Artificial Intelligence in Pharmacy. *Indian Journal of Pharmaceutical Education and Research*. 2021;55 (2):304–318. <https://doi.org/10.5530/ijper.55.2.68>
2. Riemer A, Freund V. Generative artificial intelligence in pharmaceutical drug development: A systematic review of time and cost efficiency across discovery, preclinical, and clinical phases. *Intelligent Pharmacy*. 2026;4:145–158. <https://doi.org/10.1016/j.ipha.2025.12.006>
3. Alam A, Shah SS, Rabbani SA, El-Tanani M. The Role of Artificial Intelligence in Pharmacy Practice and Patient Care: Innovations and Implications. *BioMedInformatics*. 2025;5:65. <https://doi.org/10.3390/biomedinformatics5040065>
4. Habli I, Lawton T, Porter Z. Artificial intelligence in health care: accountability and safety. *Bulletin of the World Health Organization*. 2020;98:251–256. <https://doi.org/10.2471/BLT.19.237487>
5. US 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 and other interested parties. January 2025. Draft status verified 15 October 2026. <https://www.USFDA.gov/media/184830/download>



6. World Health Organization. Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. Geneva: World Health Organization; 2024. ISBN 978-92-4-008475-9. Relevant sections: Box 1, pp. 4–5; societal concerns, pp. 24–29; governance and deployment recommendations, pp. 51–57.
7. Rane J, Chaudhari RA, Rane NL. Ethical Considerations and Bias Detection in Artificial Intelligence/Machine Learning Applications. Deep Science Publishing; 2025. Chapter 1: Algorithmic Fairness and Statistical Bias Mitigation in Machine Learning Systems: A Framework for Bias Detection and Algorithm Transparency. Used as a secondary conceptual source. <https://doi.org/10.70593/978-93-7185-870-0>
8. Hasan HE, Jaber D, Khabour OF, Alzoubi KH. Ethical considerations and concerns in the implementation of AI in pharmacy practice: a cross-sectional study. *BMC Medical Ethics*. 2024;25:55. <https://doi.org/10.1186/s12910-024-01062-8>
9. Hasan HE, Jaber D, Al Tabbah S, Lawand N, Habib HA, Farahat NM. Knowledge, attitude and practice among pharmacy students and faculty members towards artificial intelligence in pharmacy practice: A multinational cross-sectional study. *PLoS ONE*. 2024;19 (3):e0296884. <https://doi.org/10.1371/journal.pone.0296884>
10. Li L, Du P, Huang X, Zhao H, Ni M, Yan M, et al. Comparative Analysis of Generative Artificial Intelligence Systems in Solving Clinical Pharmacy Problems: Mixed Methods Study. *JMIR Medical Informatics*. 2025;13:e76128. <https://doi.org/10.2196/76128>
11. Austin Z, Gregory P. Guidance for Use of Artificial Intelligence in Community Pharmacy Practice: Perspectives and Needs of Pharmacists in Ontario, Canada. *Pharmacy*. 2026;14:57. <https://doi.org/10.3390/pharmacy14020057>
12. Gregory PAM, Austin Z. Responsible Adoption of Artificial Intelligence (AI) in Pharmacy Practice: Perspectives of Regulators in Canada and the United States. *Pharmacy*. 2025;13:152. <https://doi.org/10.3390/pharmacy13060152>
13. Stokes JM, Yang K, Swanson K, et al. A Deep Learning Approach to Antibiotic Discovery. *Cell*. 2020;180 (4):688–702.e13. <https://doi.org/10.1016/j.cell.2020.01.021>
14. Wong F, Zheng EJ, Valeri JA, et al. Discovery of a structural class of antibiotics with explainable deep learning. *Nature*. 2024;626:177–185. <https://doi.org/10.1038/s41586-023-06887-8>
15. Zhavoronkov A, Ivanenkov YA, Aliper A, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*. 2019;37:1038–1040. <https://doi.org/10.1038/s41587-019-0224-x>
16. US Food and Drug Administration, Center for Drug Evaluation and Research. Artificial Intelligence in Drug Manufacturing. Discussion paper. 2023. Not draft or final guidance. <https://www.USFDA.gov/media/165743/download>
17. Stasevych M, Zvarych V. Innovative Robotic Technologies and Artificial Intelligence in Pharmacy and Medicine: Paving the Way for the Future of Health Care—A Review. *Big Data and Cognitive Computing*. 2023;7:147. <https://doi.org/10.3390/bdcc7030147>
18. Xu Z, Ren F, Wang P, et al. A generative AI-discovered TNF inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial. *Nature Medicine*. 2025;31:2602–2610. <https://doi.org/10.1038/s41591-025-03743-2>
19. Taherdoost H, Ghofrani A. AI's role in revolutionizing personalized medicine by reshaping pharmacogenomics and drug therapy. *Intelligent Pharmacy*. 2024;2:643–650. <https://doi.org/10.1016/j.ipha.2024.08.005>
20. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366 (6464):447–453. <https://doi.org/10.1126/science.aax2342>
21. Gichoya JW, Banerjee I, Bhimireddy AR, et al. AI recognition of patient race in medical imaging: a modelling study. *The Lancet Digital Health*. 2022;4:e406–e414. [https://doi.org/10.1016/S2589-7500\(22\)00063-2](https://doi.org/10.1016/S2589-7500(22)00063-2)



22. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine*. 2021;27:2176–2182. <https://doi.org/10.1038/s41591-021-01595-0>
23. US Food and Drug Administration and European Medicines Agency. Guiding Principles of Good AI Practice in Drug Development. January 2026. Official US FDA document; accompanying US FDA webpage checked 15 October 2026. https://www.US_FDA.gov/media/189581/download
24. Indian Council of Medical Research. Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare. New Delhi: ICMR; 2023. ISBN 978-93-5811-343-3. Sections 1.9–1.10, 2–4. https://www.icmr.gov.in/icmrobject/custom_data/pdf/Ethical-guidelines/Ethical_Guidelines_AI_Healthcare_2023.pdf
25. Price WN II, Cohen IG. Privacy in the age of medical big data. *Nature Medicine*. 2019;25:37–43. <https://doi.org/10.1038/s41591-018-0272-7>
26. Carlini N, Tramèr F, Wallace E, et al. Extracting Training Data from Large Language Models. In: Proceedings of the 30th USENIX Security Symposium. USENIX Association; 2021. p. 2633–2650. <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>
27. Finlayson SG, Bowers JD, Ito J, et al. Adversarial attacks on medical machine learning. *Science*. 2019;363 (6433):1287–1289. <https://doi.org/10.1126/science.aaw4399>
28. Gaube S, Suresh H, Raue M, et al. Do as AI say: susceptibility in deployment of clinical decision-aids. *npj Digital Medicine*. 2021;4:31. <https://doi.org/10.1038/s41746-021-00385-9>
29. Fahim MIA, Tonny TS, Al Noman A. Realizing the potential of AI in pharmacy practice: Barriers and pathways to adoption. *Intelligent Pharmacy*. 2024;2:308–311. <https://doi.org/10.1016/j.ipha.2024.02.003>
30. Walters WH, Wilder EI. Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific Reports*. 2023;13:14045. <https://doi.org/10.1038/s41598-023-41032-5>
31. Glynn A. Suspected undeclared use of generative artificial intelligence. *Intelligent Pharmacy*. 2024;2:596–597. <https://doi.org/10.1016/j.ipha.2024.03.003>
32. US Food and Drug Administration. Artificial Intelligence & Medical Products: How CBER, CDER, CDRH, and OCP are Working Together. March 2024; revised February 2025. https://www.US_FDA.gov/media/177030/download?attachment=
33. US Food and Drug Administration. External Engagements with US FDA for Artificial Intelligence in Drug Development. Webpage. Accessed 15 October 2026. https://www.US_FDA.gov/about-US_FDA/artificial-intelligence-drug-development/external-engagements-US_FDA-artificial-intelligence-drug-development
34. Government of India, Ministry of Health and Family Welfare. Medical Devices Rules, 2017. G.S.R. 78 (E), 31 January 2017. CDSCO compilation supplied for review; rule 4 and First and Fifth Schedules. Subsequent applicable amendments must be consulted. https://cdsco.gov.in/opencms/resources/UploadCDSCOWeb/2022/m_device/Medical%20Devices%20Rules%2C%202017.pdf
35. Central Drugs Standard Control Organization, Medical Devices and IVD Division. Guidance Document on Medical Device Software. CDSCO/MD/GD/MDSW/01/2026. 2026. Relevant sections: 1–2, 7, 9, 12.4–12.5. https://www.cdsco.gov.in/opencms/export/sites/CDSO_WEB/Pdf-documents/Guidance-document-on-Medical-Device-Software-under-MDR-2017.pdf
36. European Medicines Agency. Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle. EMA/CHMP/CVMP/83833/2023. 2024. <https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial->



[intelligence-ai-medicinal-product-lifecycle_en.pdf](#)

37. World Health Organization. Ethics and governance of artificial intelligence for health. Geneva: WHO; 2021. ISBN 978-92-4-002920-0. <https://www.who.int/publications/i/item/9789240029200>
38. World Health Organization. Regulatory considerations on artificial intelligence for health. Geneva: WHO; 2023. ISBN 978-92-4-007887-1. <https://www.who.int/westernpacific/publications/i/item/9789240078871>
39. Lekadir K, Frangi AF, Porras AR, et al. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. *BMJ*. 2025;388:e081554. <https://doi.org/10.1136/bmj-2024-081554>
40. Tabassi E. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg, MD: National Institute of Standards and Technology; 2023. NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
41. Autio C, Schwartz R, Dunietz J, et al. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. Gaithersburg, MD: National Institute of Standards and Technology; 2024. NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>
42. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Medicine*. 2018;15 (11):e1002683. <https://doi.org/10.1371/journal.pmed.1002683>
43. Wong A, Otles E, Donnelly JP, et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Internal Medicine*. 2021;181 (8):1065–1070. <https://doi.org/10.1001/jamainternmed.2021.2626>
44. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. *Nature Medicine*. 2020;26:1351–1363. <https://doi.org/10.1038/s41591-020-1037-7>
45. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nature Medicine*. 2020;26:1364–1374. <https://doi.org/10.1038/s41591-020-1034-x>
46. Vasey B, Nagendran M, Campbell B, et al. Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *BMJ*. 2022;377:e070904. <https://doi.org/10.1136/bmj-2022-070904>
47. Collins GS, Moons KGM, Dhiman P, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ*. 2024;385:e078378. <https://doi.org/10.1136/bmj-2023-078378>
48. Gallifant J, et al. The TRIPOD-LLM reporting guideline for studies using large language models. *Nature Medicine*. 2025;31:60–69. <https://doi.org/10.1038/s41591-024-03425-5>
49. Sounderajah V, et al. The STARD-AI reporting guideline for diagnostic accuracy studies using artificial intelligence. *Nature Medicine*. 2025;31:3283–3289. <https://doi.org/10.1038/s41591-025-03953-8>
50. Moons KGM, Damen JAA, Kaul T, et al. PROBAST+AI: an updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. *BMJ*. 2025;388:e082505. <https://doi.org/10.1136/bmj-2024-082505>

HOW TO CITE: A. Sangdokar, Artificial Intelligence in Pharmaceuticals: Ethical Risks, Regulatory Challenges, and a Framework for Responsible Governance, *Int. J. of Pharm. Sci.*, 2026, Vol 4, Issue 10, 4375-4391. <https://doi.org/10.5281/zenodo.23722087>

