

Artificial Intelligence in Clinical Pharmacy: Current Applications, Clinical Evidence, Challenges, and Future Perspectives A Critical Narrative Review

Sanket B. Dude¹, Vijay H. Shelke², Rushikesh S. Gawade³, Monika P. Barhate⁴

^{1,2,3}Students, Sahakar Maharshi Kisanrao Varal Patil College of Pharmacy, Ahilyanagar, Maharashtra, India

⁴Guide, Sahakar Maharshi Kisanrao Varal Patil College of Pharmacy, Ahilyanagar, Maharashtra, India

Abstract: Artificial intelligence (AI) is increasingly being explored as a means of supporting medication-related decisions across hospital and community care. In clinical pharmacy, relevant applications include prescription screening, medication-error prioritisation, drug–drug interaction assessment, adverse drug reaction (ADR) detection, pharmacovigilance, individualised dose prediction, medication reconciliation, adherence support, and clinical documentation. This critical narrative review summarises the main AI approaches and considers not only technical performance but also their relevance to pharmacist-led care. Literature was identified through a focused search of PubMed/PubMed Central and publisher or regulatory sources using combinations of terms related to artificial intelligence, machine learning, clinical pharmacy, medication safety, drug interactions, pharmacovigilance, precision dosing, and generative AI. Key reviews, original studies, and authoritative guidance available up to October 2026 were prioritised. The literature indicates that AI can help process large, complex datasets and identify medication-related patterns that may be difficult to detect manually. However, the evidence is heterogeneous: many systems are evaluated retrospectively, use institution-specific datasets, or report model metrics rather than improvements in patient outcomes. Generative AI may assist with information retrieval and documentation, but errors, fabricated statements, incomplete context, and variable performance make unsupervised clinical use unsafe. Practical implementation therefore requires external and prospective validation, transparent outputs, privacy safeguards, workflow integration, monitoring, and clear professional accountability. AI should currently be understood as a decision-support capability that can augment—not replace—the clinical pharmacist. Future research should assess patient-centred outcomes, implementation costs, equity, and performance in diverse healthcare settings, including resource-constrained environments.

Keywords: Artificial intelligence; clinical pharmacy; machine learning; medication safety; pharmacovigilance; clinical decision support; drug–drug interactions; precision dosing; generative AI

Abbreviation / term	Meaning
ADE	Adverse drug event
ADR	Adverse drug reaction
AI	Artificial intelligence
CDSS	Clinical decision support system
DDI	Drug–drug interaction



EHR	Electronic health record
FDA	United States Food and Drug Administration
INR	International normalized ratio
LLM	Large language model
ML	Machine learning
NLP	Natural language processing
WHO	World Health Organization

I. INTRODUCTION

Clinical pharmacy focuses on the safe, effective, and appropriate use of medicines through medication review, monitoring, counselling, prevention of medication-related problems, and collaboration with the wider healthcare team. These tasks have become more complex as patients frequently receive multiple medicines, have several long-term conditions, and move between care settings. Pharmacists must interpret prescription details alongside allergies, laboratory results, diagnoses, organ function, prior treatment response, and patient preferences. Manual review remains essential, but the volume and complexity of available data can make it difficult to prioritise every potential risk consistently.

Artificial intelligence is an umbrella term for computational methods that perform tasks commonly associated with human reasoning, pattern recognition, language processing, or prediction. In clinical pharmacy, AI systems may analyse structured data such as laboratory results and medication lists, or unstructured information such as clinical notes, discharge summaries, and adverse-event narratives. Their potential value lies in identifying patterns, flagging cases for review, and presenting relevant information in a more timely manner. The American Society of Health-System Pharmacists and pharmacy-practice literature have recognised AI as a developing area for pharmacy operations and patient-care services, while also highlighting the need for professional oversight [1,2].

The central question is not whether AI can generate a prediction, but whether using that prediction improves a meaningful clinical process without creating unacceptable risks. High accuracy in a retrospective dataset does not necessarily demonstrate clinical benefit. A model may perform well in one hospital but fail elsewhere because of differences in patient populations, prescribing habits, formularies, documentation, or electronic health record systems. Likewise, a system that produces too many low-value alerts can increase workload and encourage users to ignore warnings. These issues are particularly important when outputs may influence high-risk medicines or urgent treatment decisions [3,4].

This review therefore considers AI through a clinical-pharmacy lens. It describes major technologies and applications, separates technical or preclinical evidence from implementation and patient-outcome evidence, and discusses practical requirements for safe adoption. It does not assume that every application is mature or ready for routine use.

II. REVIEW METHODOLOGY

This article was developed as a critical narrative review, not as a registered systematic review. A focused literature search was conducted in PubMed/PubMed Central and through targeted searches of publisher, journal, and authoritative institutional websites. Search concepts combined terms for artificial intelligence, machine learning, deep learning, natural language processing, large language models, clinical pharmacy, pharmacy practice, medication safety, prescribing error, drug–drug interaction, adverse drug reaction, pharmacovigilance, precision dosing, and clinical decision support. Current literature available through October 2026 was prioritised, while older foundational studies and professional guidance were retained where relevant.



Priority was given to peer-reviewed original research, systematic or scoping reviews, pharmacy-practice publications, and guidance from authoritative bodies such as the World Health Organization (WHO) and the United States Food and Drug Administration (FDA). Studies were considered relevant when they addressed medication use, medication safety, pharmacist-facing workflows, clinical decision support, drug safety surveillance, or individualised pharmacotherapy. General AI studies without a clear connection to medicine use were used sparingly.

Because this was a focused narrative search, the review should not be interpreted as exhaustive and no formal risk-of-bias assessment was undertaken across all included publications. The synthesis is qualitative and distinguishes model development, internal validation, external validation, prospective implementation, and demonstrated patient benefit. The absence of a reproducible, database-by-database search log is a limitation; accordingly, this review does not claim systematic-review completeness.

III. AI TECHNOLOGIES RELEVANT TO CLINICAL PHARMACY

3.1 Machine learning.

Machine learning (ML) algorithms learn statistical relationships from examples rather than relying solely on manually written rules. Supervised learning can classify prescriptions as higher or lower risk, estimate the likelihood of an adverse event, or predict a dose-related outcome from labelled data. Unsupervised methods can group similar patients or identify unusual patterns. Model performance depends on the representativeness of training data, the quality of outcome labels, and the clinical consequences of false-positive and false-negative predictions. The main technology classes and their intended roles are summarised in Table 1.

Table 1. AI technologies and their potential roles in clinical pharmacy

Technology	Potential clinical-pharmacy role	Key caution
Rule-based CDSS	Guideline, allergy, dose-limit and contraindication checks	May generate excessive or non-contextual alerts
Machine learning	Risk ranking, medication-error prediction, patient stratification	Bias, poor calibration and weak external validity
Deep learning	Analysis of complex longitudinal records and clinical text	High data needs and limited interpretability
NLP	Extract medicines, adverse events and discrepancies from text	Negation, abbreviations and context may be misread
Generative AI / LLMs	Drafting, summarisation, counselling text and information navigation	Plausible but false output; source verification is essential

3.2 Deep learning.

Deep learning uses multi-layer neural networks to model complex patterns in large datasets. It may be applied to longitudinal electronic health records, medication sequences, laboratory trends, or text. More complex models can capture relationships that simpler methods miss, but their outputs may be harder to explain and they may require substantial data and computational resources.

3.3 Natural language processing.

Natural language processing (NLP) extracts or interprets information from text, including clinical notes, discharge summaries, spontaneous safety reports, and patient communications. It can support identification of drug names, adverse events, indications, and temporal relationships. Errors may arise from abbreviations, negation, misspellings, multilingual documentation, and ambiguity. Human review remains important when extracted information could affect care [5,6].



3.4 Knowledge-based and hybrid systems.

Traditional clinical decision support often uses explicit rules, drug databases, interaction tables, and guideline logic. These systems are not necessarily AI. Hybrid systems combine established rules with ML-based risk ranking or text processing. This can be useful because some medication risks are governed by well-established contraindications, while other cases require prioritisation based on multiple patient-specific variables.

3.5 Generative AI and large language models.

Generative AI systems produce text or other content in response to prompts. Potential pharmacy uses include drafting patient-friendly explanations, summarising records, preparing documentation, and helping clinicians navigate drug information. However, language fluency is not equivalent to factual accuracy. A model can produce plausible but incorrect content, omit important contraindications, or fail to recognise missing clinical details. Outputs should be checked against authoritative sources and patient-specific records, and sensitive data should not be entered into unapproved tools [7,8].

IV. MAJOR APPLICATIONS IN CLINICAL PHARMACY

4.1 Prescription screening and medication-error prioritisation

Electronic prescribing systems can generate alerts for allergies, duplicate therapy, dose limits, contraindications, and interactions. Rule-based systems are useful for explicit, well-defined risks, but can generate many alerts that are irrelevant to the individual patient. ML approaches may add value by ranking orders according to the likelihood that pharmacist intervention will be needed. Corny and colleagues described a machine-learning-based decision support system designed to identify prescriptions at high risk of medication error [9]. Other work has explored prediction of atypical medication orders and pharmacist intervention needs [10,11].

These approaches may help pharmacists focus on higher-risk prescriptions, especially in settings with high order volumes. Yet a prediction that an order is “atypical” is not proof that it is inappropriate: unusual orders can be clinically justified. Models should therefore support prioritisation, not automatically reject prescriptions. Evaluation should include sensitivity for serious errors, false-alert burden, time to review, pharmacist acceptance, and patient-safety outcomes.

4.2 Drug–drug interactions

Drug–drug interaction (DDI) assessment is challenging because risk depends on more than the presence of two medicines. Dose, route, timing, duration, renal and hepatic function, therapeutic indication, monitoring availability, and patient susceptibility may all matter. AI and data-mining methods have been investigated for identifying potential interactions and detecting signals from real-world data [12]. Network-based and ML approaches can identify candidate associations, while NLP may extract interaction-related information from clinical notes or published literature.

A key limitation is that a statistical association does not automatically establish a clinically meaningful interaction or causal relationship. DDI models may also struggle with rare combinations, incomplete medication lists, and changes in treatment over time. The pharmacist must assess the severity, mechanism, evidence quality, and practical management options. Newer evaluations of large language models have also found that performance can vary across interaction identification, clinical relevance assessment, and recommendation quality [13]. Generative systems should not replace validated interaction databases or professional judgement.

4.3 Adverse drug reaction detection and pharmacovigilance

Pharmacovigilance involves detecting, assessing, understanding, and preventing adverse effects or other medicine-related problems. AI may assist by processing spontaneous reports, electronic health records, scientific literature, social-media content, and clinical narratives. NLP can extract drug-event pairs and relevant context; ML can prioritise



reports or identify patterns that warrant further investigation. Reviews have mapped the use of AI in pharmacovigilance, including adverse-event extraction and safety-signal detection [5,6,14]. Text-mining and NLP studies illustrate how medication-event narratives and other clinical text can be analysed, although extraction performance and clinical meaning require verification [24-27].

The benefit is potentially faster processing of large volumes of text and more efficient triage. Nevertheless, models can confuse symptoms with adverse reactions, fail to account for underlying disease, or overstate signals from biased reporting. A signal is a reason for assessment, not confirmation that a medicine caused an event. Clinical validation, pharmacological plausibility, temporal assessment, and established signal- evaluation procedures remain necessary. AI should complement, rather than bypass, formal pharmacovigilance systems.

4.4 Individualised dosing and precision pharmacotherapy

Dose requirements vary with age, body size, renal and hepatic function, genetics, comorbidities, concomitant medicines, adherence, and treatment response. ML models have been investigated for predicting dose requirements and treatment outcomes, including warfarin dosing. Warfarin is a useful example because of its narrow therapeutic range and the need for repeated international normalised ratio (INR) monitoring. Research has explored ML approaches to predict stable dose requirements and support anticoagulation management [15,16]. Pharmacogenetic guidelines also illustrate how genetic information can be combined with clinical factors when individualising warfarin therapy [17].

Prediction accuracy alone is insufficient for adopting an AI dosing tool. An incorrect dose can cause serious harm, and an apparently accurate model may perform differently across populations. Clinical use requires external validation, comparison with accepted dosing strategies, prospective assessment, and monitoring of bleeding, thrombosis, time in therapeutic range, and other clinically meaningful outcomes. Pharmacists should verify the data used by the model and ensure that recommendations align with current guidelines and the patient's circumstances.

4.5 Medication reconciliation and transitions of care

Medication reconciliation aims to establish an accurate list of medicines when a patient is admitted, transferred, or discharged. Differences between lists may reflect omissions, duplications, changes in dose, or intentional discontinuation. NLP and ML may help extract medicines from documents and identify discrepancies across records. Such tools could reduce manual data-entry burden and prioritise discrepancies for review. However, an apparent discrepancy is not always an error; the reason for each change may require clinical context or direct confirmation from the patient, caregiver, prescriber, or dispensing record. AI outputs should therefore be treated as candidate discrepancies requiring verification. Early evaluations of AI-enabled reconciliation tools have also highlighted the importance of usability and clinician acceptance [28].

4.6 Medication adherence and patient support

AI-enabled applications may use refill histories, patient-reported information, digital reminders, or communication patterns to identify patients who could benefit from adherence support. Predictive models may help target counselling or follow-up, but they should not assume that a missed refill means deliberate non-adherence. Cost, side effects, access barriers, health literacy, complex regimens, and cultural or language factors may explain gaps in medicine use. Pharmacists can use a risk flag as a prompt for supportive, non-judgemental assessment rather than as a label applied to a patient.

4.7 Clinical documentation and drug-information services

NLP and generative AI may help draft medication histories, summarise long records, prepare counselling material, and convert technical information into plain language. These uses may improve efficiency when outputs are reviewed. However, generative models may omit key facts, invent references, confuse drug strengths, or produce advice that is



not appropriate for the specific patient. A safe workflow should require source verification, explicit review of dose and units, protection of patient information, and clear labelling of AI-generated drafts. Patient-facing advice should be checked for readability, language suitability, and consistency with approved product information and local guidance. Broader reviews of generative AI in healthcare similarly caution that utility must be balanced against factual errors, omissions, and the need for human verification [23].

4.8 Hospital pharmacy operations and inventory

A cross-domain summary of potential value and evidence requirements is provided in Table 2.

Table 2. Application areas, potential value, and evidence questions

Application	Potential value	Evidence needed before routine use
Prescription screening	Prioritise high-risk orders for pharmacist review	Serious-error sensitivity, false-alert burden, prospective workflow outcomes
Drug–drug interactions	Identify candidate interactions and patient-specific risk	Clinical relevance, causal evidence, recommendation quality
ADR / pharmacovigilance	Triage reports and extract drug-event information	Signal validity, time saved, confirmed safety outcomes
Individualised dosing	Estimate dose or treatment response using patient features	External validation, dose safety, patient outcomes and subgroup performance
Medication reconciliation	Find discrepancies across medication lists and records	Accuracy of discrepancy detection and confirmed clinical significance
Adherence support	Identify patients who may need additional support	Fairness, patient acceptability, adherence and health outcomes
Documentation / generative AI	Summarise records and draft counselling or notes	Factual error rate, omissions, readability, privacy and review time
Inventory / operations	Forecast demand and reduce shortages or waste	Stock-out rate, expiry, cost and continuity of supply

AI has also been explored for medication identification, forecasting medicine demand, anticipating shortages, and improving workflow prioritisation. These functions may support continuity of supply and operational efficiency, indirectly benefiting patient care. Nevertheless, inventory optimisation is not the same as clinical decision-making, and performance should be evaluated against stock-out rates, wastage, expiry, service continuity, and cost. Data from one hospital may not generalise to another because of different formularies, procurement cycles, seasonal demand, and supplier reliability [1,2,18].

Table 4. Selected evidence examples and interpretation

Application source	Study type or setting	Reported contribution	Key limitation / interpretation
Prescription-error prioritisation [9]	Hospital EHR data; model development and independent validation	Hybrid machine-learning and rule-based system ranked prescriptions for pharmacist review; reported better discrimination than comparators in that setting.	Single institutional context; prediction performance does not itself prove reduced patient harm.
Drug–drug interaction assessment [13]	Evaluation using 15 constructed patient cases and three LLMs	Performance differed by model and task; identification did not guarantee accurate clinical relevance or optimal management advice.	Small case set; results should not be generalised to unsupervised clinical use.



Medication reconciliation [28]	Small observational pilot involving 10 patients	AI-supported tool was usable in the pilot and helped patients identify issues in medication lists.	Very small sample and usability focus; not definitive evidence of improved clinical outcomes.
AI-enabled early warning system [29]	Prospective, multi-site cohort study in sepsis care	Provider interaction with a deployed alert system was associated with better outcomes in the studied setting.	Indirect evidence for clinical pharmacy; should not be presented as proof for medication-specific AI tools.
Clinical-pharmacy review landscape [19]	Systematic review	Summarises AI applications in clinical pharmacy and reinforces need for prospective, multicentre evaluation.	A review is secondary evidence; its methods and included studies should be assessed alongside original research.

Evidence note: These examples illustrate different evidence levels and are not a formal risk-of-bias assessment or exhaustive evidence synthesis.

V. CLINICAL EVIDENCE AND READINESS FOR PRACTICE

The available literature contains several different levels of evidence that should not be treated as interchangeable. A model-development study demonstrates that an algorithm can be trained on a dataset. Internal validation estimates performance on held-out data from a similar source. External validation tests whether performance persists in a different institution, time period, or population. Prospective implementation studies assess the tool under real workflow conditions, while controlled studies can provide stronger evidence about whether its use improves care or patient outcomes. Table 3 presents a practical maturity framework for interpreting these levels.

Table 3. Evidence maturity framework for evaluating AI tools

Stage	What it establishes	What it does not establish
1. Model development	An algorithm can be trained on available data	Reliable performance in new patients or settings
2. Internal validation	Performance on held-out data from a similar source	Transportability to another institution
3. External validation	Performance in a different population/site/time period	Improved care in routine practice
4. Prospective implementation	Feasibility and performance in a live workflow	Definitive causal benefit unless comparator design supports it
5. Patient-outcome evaluation	Impact on safety, effectiveness, experience or cost	Automatic generalisation to all settings

Reviews of AI in pharmacy practice have described applications in drug safety, pharmacy operations, and precision medicine, but have also highlighted transparency, generalisability, and external validation as continuing challenges [1]. A systematic review focused on clinical pharmacy similarly identified medication-error detection, adverse-event assessment, clinical decision support, and dose-related applications, while noting the need for stronger prospective and multicentre evidence [19]. A 2026 scoping review mapped AI across clinical pharmacy services and reported that much of the literature remained retrospective or computational, with limited evidence of sustained integration into routine pharmacist workflows [20]. These recent reviews are important context, but their existence also means a new broad review must offer a clear angle rather than simply repeat a general catalogue of applications.

For clinical pharmacy, useful endpoints extend beyond area under the receiver operating characteristic curve or classification accuracy. Relevant measures include serious medication errors detected, false alerts per patient or



prescription, time to pharmacist intervention, accepted recommendations, preventable adverse drug events, readmissions, length of stay, patient understanding, treatment control, workload, cost-effectiveness, and equity. A model may have excellent discrimination yet be poorly calibrated, generate too many alerts, or fail to improve decisions. The model should be evaluated in the specific clinical role for which it is intended.

Implementation also changes the human-technology interaction. If pharmacists do not understand why a system flagged a prescription, they may ignore it or over-trust it. Training, clear escalation pathways, feedback mechanisms, and routine monitoring are therefore part of the intervention, not optional additions. The safest current model is a pharmacist-led workflow in which AI identifies candidates or synthesises information, while a trained professional verifies the output and remains responsible for the clinical decision. Evidence from prospective implementation of machine-learning systems in other clinical settings further illustrates why workflow and patient outcomes should be evaluated alongside model metrics [29,30].

VI. CHALLENGES, SAFETY, AND ETHICAL CONSIDERATIONS

6.1 Data quality and representativeness.

Electronic records contain missing values, inconsistent coding, duplicated entries, and documentation practices that vary across institutions. Training data may under-represent certain age groups, ethnicities, socioeconomic groups, languages, or people with complex conditions. Such gaps can produce unequal performance. Models should be tested across relevant subgroups and settings before use.

6.2 Explainability and calibration.

A risk score is most useful when clinicians understand what it means and how much confidence to place in it. Complex models can be difficult to interpret. Even a well-ranked risk score may be poorly calibrated, meaning that predicted probabilities do not correspond to observed event rates. Clear limitations, confidence information, and clinically meaningful explanations can support safer use.

6.3 Automation bias and alert fatigue.

Clinicians may accept computer-generated recommendations without adequate independent review, or ignore alerts after repeated low-value warnings. Systems should be designed to prioritise clinically important alerts, show relevant evidence, allow appropriate override, and record the reason for consequential decisions. AI must not create the illusion that every alert is correct or every unflagged prescription is safe.

6.4 Privacy, security, and confidentiality.

Clinical pharmacy data may include sensitive personal information. Data access, storage, model training, vendor contracts, retention, and cross-border processing require governance. Staff should use only institutionally approved tools for identifiable patient data and follow applicable law and policy. De-identification reduces some risks but does not eliminate re-identification or security concerns.

6.5 Generative AI errors.

Large language models can produce plausible but false statements, outdated recommendations, inaccurate citations, or incomplete safety advice. They may also respond confidently when the question lacks essential information. High-risk tasks such as dose selection, contraindication assessment, and emergency advice require authoritative-source verification and professional oversight. A language model should not be treated as a drug database merely because it can discuss medicines fluently.



6.6 Accountability and regulation

Responsibility for a clinical decision cannot be assumed to transfer to the algorithm or software vendor. Institutions need clear policies defining the tool's intended use, professional review requirements, incident reporting, audit trails, updates, and withdrawal criteria. WHO guidance emphasises human autonomy, safety, transparency, accountability, inclusiveness, and sustainability in AI for health [7,8]. The FDA likewise highlights lifecycle management and transparency considerations for AI-enabled medical devices [21,22]. Regulatory requirements differ by jurisdiction and intended use; authors and implementers should avoid assuming that one country's regulatory framework applies universally.

6.7 Cost, infrastructure, and digital equity

AI deployment may require reliable electronic records, interoperability, technical support, cybersecurity, staff training, and ongoing model maintenance. Hospitals with fewer resources may be unable to implement systems that depend on extensive infrastructure. The value proposition should include total cost, maintenance, staff time, accessibility, and the consequences of system failure—not only initial model accuracy.

VII. INDIA-SPECIFIC CONSIDERATIONS

India has a diverse health system with differences in hospital infrastructure, access to clinical pharmacists, language, medicine availability, and electronic health record maturity. AI tools could potentially support medication review, pharmacovigilance, patient counselling, and identification of high-risk prescriptions, especially where clinical teams face high workloads. However, a model developed using data from a different country may not transfer reliably because of differences in prescribing practices, disease prevalence, formulations, brand names, documentation, and patient characteristics.

Indian implementation should prioritise local validation, support for relevant languages, secure data governance, and integration with existing workflows rather than adding a separate system that increases burden. Tools should be tested in both tertiary hospitals and resource-constrained settings, with attention to connectivity, staffing, cost, and usability. Pharmacists and other healthcare professionals should be involved in design and evaluation so that model outputs fit actual medication-use processes. Where clinical evidence is limited, AI should be introduced as a monitored support tool rather than as an autonomous decision-maker. The use of AI should also align with applicable Indian law, institutional policies, and current national guidance; legal and regulatory requirements should be checked at the time of implementation.

VIII. PRACTICAL IMPLEMENTATION FRAMEWORK FOR CLINICAL PHARMACY TEAMS

A responsible implementation plan should begin with a clearly defined medication-related problem rather than with the availability of a particular algorithm. Examples include excessive low-value alerts, delayed identification of high-risk prescriptions, incomplete medication reconciliation, or inefficient review of large volumes of safety reports. The pharmacy team should define the intended users, the clinical decision to be supported, the point in the workflow at which the output will appear, and the action expected after a positive signal. A system without a defined action pathway may generate information without improving care.

Before deployment, the local team should assess data completeness, coding practices, representativeness, integration requirements, and the consequences of false-positive and false-negative predictions. Validation should include a comparison with current practice and should be performed on data that were not used for model development. Where feasible, a silent prospective phase can measure how the tool performs without influencing care. This can reveal workflow issues, drift in input data, and unexpected alert burden before clinicians rely on the output.

Implementation should include pharmacist training, a named clinical owner, escalation rules, documentation standards, and a process for reporting unsafe recommendations. The interface should communicate the reason for a flag, the level of uncertainty, and the limits of the model. Clinicians must retain the ability to override



recommendations, but overrides should be reviewed to identify both inappropriate alerts and missed risks. Monitoring should continue after launch because prescribing patterns, formularies, patient populations, and clinical guidance change over time.

Evaluation should use a balanced set of measures. Technical measures may include sensitivity, specificity, calibration, and performance across patient subgroups. Workflow measures may include pharmacist review time, alert acceptance, override frequency, and time to intervention. Clinical measures may include preventable medication errors, adverse drug events, treatment delays, and patient understanding. Economic measures should consider implementation, maintenance, training, and downstream costs. Patient-centred measures, including trust, accessibility, and equity, should be considered alongside efficiency.

A staged implementation pathway is preferable: (1) define the problem and intended use; (2) map the workflow and risks; (3) validate locally; (4) conduct a limited pilot; (5) evaluate process and patient outcomes; and (6) continue monitoring, audit, and recalibration. This framework helps prevent the common mistake of treating a high-performing model as automatically ready for clinical use. The overall workflow is illustrated in Figure 1, while Figure 2 summarises the evidence-to-implementation pathway.

Figure 1. AI-supported clinical pharmacy workflow



Figure 1. AI-supported clinical pharmacy workflow.

Figure 2. Evidence pathway before routine implementation



Figure 2. Evidence-to-implementation pathway.

IX. RESEARCH GAPS AND FUTURE PERSPECTIVES

First, more prospective and multicentre studies are needed. Many publications describe retrospective prediction or technical validation, but routine practice requires evidence that tools work across sites, populations, and changing



clinical conditions. Second, research should use clinically meaningful endpoints. Future evaluations should measure medication harm, intervention quality, treatment control, patient experience, workload, and costs, not only model discrimination.

Third, explainability and uncertainty reporting need improvement. A system should communicate why a case was flagged, what information is missing, and when it is outside its validated use. Fourth, external validation and ongoing monitoring should be built into the deployment plan. Changes in prescribing behaviour, patient mix, clinical guidelines, and electronic records can degrade performance over time. Fifth, evaluations should examine bias and equity, including differences by age, sex, language, comorbidity, and access to care.

Sixth, generative AI should be evaluated for specific, bounded pharmacy tasks rather than judged by general conversational quality. Useful research questions include whether a model can accurately summarise medication histories, identify clinically important omissions, draft understandable counselling text, or retrieve evidence with traceable citations. These tasks require comparison against expert pharmacist review and authoritative references, with errors classified by severity. Seventh, implementation research should assess how AI affects pharmacist workload and professional judgement. An effective tool should improve care or release time for meaningful patient contact, not simply move work from one screen to another.

Finally, research from low- and middle-income countries is needed to assess feasibility, affordability, language performance, and generalisability. Collaborative datasets, privacy-preserving approaches, transparent reporting, and shared evaluation standards may accelerate progress. Future systems are likely to be most useful when they combine reliable clinical rules, patient-specific data, evidence retrieval, and human review within a clearly defined workflow.

X. DISCUSSION

The most defensible near-term role for AI in clinical pharmacy is to augment attention and information processing. Medication review is a high-information task in which a pharmacist must identify a small number of clinically important issues among many prescriptions and patient details. AI may help triage records, extract relevant text, identify unusual patterns, and organise evidence for review. These functions are potentially valuable even when the final decision remains with a human professional.

However, there is a risk of confusing technical novelty with clinical value. A model that predicts an outcome accurately may not indicate an effective intervention, and a system that improves speed may still increase harm if it produces misleading recommendations. The appropriate benchmark is not simply whether AI performs a task, but whether AI-supported care is safer, more effective, more equitable, or more efficient than current practice under realistic conditions. Comparisons should include the standard workflow and should consider the added burden of reviewing alerts and correcting errors.


The recent publication of broad reviews specifically addressing AI in clinical pharmacy also changes the publication landscape [19,20]. A manuscript that repeats the same list of applications is unlikely to offer sufficient novelty. The contribution of this review is therefore its emphasis on evidence maturity, clinical workflow, implementation risk, and decision points for pharmacist oversight. For journal submission, this emphasis should be strengthened further by a reproducible updated search, a transparent evidence table with verified study details, and a clearly stated selection approach. If the target journal expects a systematic or scoping review, the methods and reporting would need to be upgraded accordingly.

The pharmacist's role is not made redundant by AI. Clinical interpretation, patient communication, shared decision-making, ethical judgement, and accountability remain central. AI outputs can be wrong, incomplete, biased, or inappropriate for a particular patient. A human-in-the-loop approach is therefore a practical safety requirement, not merely a transitional preference.



Figure 3. Human-AI safety loop.

Figure 3. Human-AI safety loop



XI. CONCLUSION

Artificial intelligence offers a range of potential applications in clinical pharmacy, including prescription screening, drug-interaction assessment, adverse-event detection, pharmacovigilance, individualised dosing, medication reconciliation, adherence support, documentation, and pharmacy operations. Current evidence suggests that AI can assist with pattern recognition and prioritisation, but the strength of evidence differs across applications. Retrospective model performance should not be interpreted as proof of improved patient outcomes, and generative AI should not be used as an unsupervised source of clinical drug advice.

Responsible implementation requires high-quality and representative data, external and prospective validation, clinically meaningful endpoints, privacy protection, transparent outputs, workflow integration, and ongoing monitoring. The most appropriate current model is collaborative: AI supports information processing and risk identification, while pharmacists verify evidence, interpret patient context, and make or oversee clinical decisions. Future research should focus on measurable patient benefit, implementation costs, equity, and validation in diverse healthcare settings, including India. With these safeguards, AI may strengthen clinical pharmacy services; without them, it may introduce new risks into already complex medication-use systems.

Acknowledgements

The authors acknowledge the guidance of Monika P. Barhate during the preparation of this review.

Declarations

Funding: The authors must confirm and complete this statement before submission.

Conflicts of interest: All authors must confirm and complete this statement before submission.

Author contributions: To be completed and approved by all authors before submission.

Data availability: No new datasets were generated or analysed for this narrative review.

Ethics approval: Not applicable; this review synthesises previously published literature.

Use of AI tools: The authors should disclose any AI-assisted writing or editing in accordance with the target journal's current policy and verify all scientific statements and references.

REFERENCES

1. Schutz N, Olsen CA, McLaughlin AJ, Yi WM, Nelson SD, Kalichira AL, et al. ASHP statement on the use of artificial intelligence in pharmacy. *Am J Health Syst Pharm.* 2020;77(23):2015-2018.
2. Flynn A. Using artificial intelligence in health-system pharmacy practice: finding new patterns that matter. *Am J Health Syst Pharm.* 2019;76(9):622-627.



3. Wong A, Wentz E, Palisano N, Dirani M, Elsamadisi P, Qashou F, et al. Role of artificial intelligence in pharmacy practice: a narrative review. *J Am Coll Clin Pharm*. 2023;6(11):1237-1250. doi:10.1002/jac5.1856.
4. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med*. 2019;17:195.
5. Li Y, Tao W, Li Z, Sun Z, Li F, Fenton S, et al. Artificial intelligence-powered pharmacovigilance: a review of machine and deep learning in clinical text-based adverse drug event detection for benchmark datasets. *J Biomed Inform*. 2024;152:104621. doi:10.1016/j.jbi.2024.104621.
6. Bate A, Evans SJW. Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol Drug Saf*. 2009;18(6):427-436.
7. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021. ISBN:978-92-4-002920-0.
8. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2024. ISBN:978-92-4-008475-9.
9. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O, et al. A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. *J Am Med Inform Assoc*. 2020;27(11):1688-1694.
10. Balestra M, Chen J, Iturrate E, Aphinyanaphongs Y, Nov O. Predicting inpatient pharmacy order interventions using provider action data. *JAMIA Open*. 2021;4(3):ooab083.
11. Hogue SC, Chen F, Brassard G, Lebel D, Bussi eres JF, Durand A, Thibault M. Pharmacists' perceptions of a machine learning model for the identification of atypical medication orders. *J Am Med Inform Assoc*. 2021;28(8):1712-1718. doi:10.1093/jamia/ocab071.
12. Hauben M. Artificial intelligence and data mining for the pharmacovigilance of drug-drug interactions. *Clin Ther*. 2023;45(2):117-133. doi:10.1016/j.clinthera.2023.01.002.
13. Chase A, Most A, Sikora A, Smith SE, Devlin JW, Xu S, Liu T, Murray B. Evaluation of large language models' ability to identify clinically relevant drug-drug interactions and generate high-quality clinical pharmacotherapy recommendations. *Am J Health Syst Pharm*. 2026;83(13):e494-e500. doi:10.1093/ajhp/zxaf168.
14. Kompa B, Hakim JB, Palepu A, Kompa KG, Smith M, Bain PA, 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.
15. Asimwe IG, Blockman M, Cohen K, Cupido C, Hutchinson C, Jacobson B, et al. Stable warfarin dose prediction in sub-Saharan African patients: a machine-learning approach and external validation of a clinical dose-initiation algorithm. *CPT Pharmacometrics Syst Pharmacol*. 2022;11(1):20-29.
16. Petch J, Nelson W, Wu M, Ghassemi M, Benz A, Fatemi M, et al. Optimizing warfarin dosing for patients with atrial fibrillation using machine learning. *Sci Rep*. 2024;14(1):4516. doi:10.1038/s41598-024-55110-9.
17. Johnson JA, Caudle KE, Gong L, Whirl-Carrillo M, Stein CM, Scott SA, et al. Clinical Pharmacogenetics Implementation Consortium guideline for pharmacogenetics-guided warfarin dosing: 2017 update. *Clin Pharmacol Ther*. 2017;102(3):397-404.
18. Liu I, Colmenares E, Tak C, Vest MH, Clark H, Oertel M, Pappas A. Development and validation of a predictive model to predict and manage drug shortages. *Am J Health Syst Pharm*. 2021;78(14):1309-1316. doi:10.1093/ajhp/zxab152.
19. Alqahtani SS, Menachery SJ, Alshahrani A, Albalkhi B, Alshayban D, Iqbal MZ. Artificial intelligence in clinical pharmacy—a systematic review of current scenario and future perspectives. *Digit Health*. 2025;11:20552076251388145. doi:10.1177/20552076251388145.
20. Al-Taie A, Abdullah AS, Arueyingho O. Integration of artificial intelligence applications in clinical pharmacy services: a scoping review. *Future Healthc J*. 2026;13(3):100547. doi:10.1016/j.fhj.2026.100547.
21. US Food and Drug Administration. CDRH issues guiding principles for transparency of machine learning-enabled medical devices. FDA; 2024. Available from: <https://www.fda.gov/medical-devices/medical-devices-news-and>



events/cdrh-issues-guiding-principles-transparency-machine-learning-enabled-medical-devices (accessed 12 October 2026).

22. US Food and Drug Administration. Artificial intelligence-enabled medical devices. FDA Digital Health Center of Excellence. Available from: <https://www.fda.gov/medical-devices/digital-health-center-excellence/artificial-intelligence-enabled-medical-devices> (accessed 12 October 2026).

23. Sallam M. ChatGPT utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. *Healthcare (Basel)*. 2023;11(6):887.

24. Rezaei Z, Ebrahimpour-Komleh H, Eslami B, Chavoshinejad R, Totonchi M. Adverse drug reaction detection in social media by deep learning methods. *Cell J*. 2020;22(3):319-324.

25. Carrell DS, Cronkite D, Palmer RE, Saunders K, Gross DE, Masters ET, et al. Using natural language processing to identify problem usage of prescription opioids. *Int J Med Inform*. 2015;84(12):1057-1064.

26. Dong X, Deng J, Hou W, Rashidian S, Rosenthal RN, Saltz M, et al. Predicting opioid overdose risk of patients with opioid prescriptions using electronic health records based on temporal deep learning. *J Biomed Inform*. 2021;116:103725.

27. Lester CA, Kessler JM, Modisett T, Chui MA. A text mining analysis of medication quality related event reports from community pharmacies. *Res Social Adm Pharm*. 2019;15(7):845-851. doi:10.1016/j.sapharm.2018.09.013.

28. Long J, Yuan MJ, Poonawala R. An observational study to evaluate the usability and intent to adopt an artificial intelligence-powered medication reconciliation tool. *Interact J Med Res*. 2016;5(2):e14.

29. Adams R, Henry KE, Sridharan A, Soleimani H, Zhan A, Rawat N, et al. Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. *Nat Med*. 2022;28(7):1455-1460.

30. Stephenson J. WHO offers guidance on use of artificial intelligence in medicine. *JAMA Health Forum*. 2021;2(7):e212467. doi:10.1001/jamahealthforum.2021.2467.

