Original Article

Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over

Sofía Martínez[1]*, Carlos Gómez[1], Lucia Navarro[2]

[1]Department of Prescription Verification and AI, Faculty of Pharmacy, Polytechnic University of Madrid, Madrid, Spain. [2]Department of Pharmacy Safety and Human Oversight, Faculty of Pharmacy, University of Barcelona, Barcelona, Spain.

Abstract

Prescription-verification systems increasingly combine rules, clinical knowledge bases, predictive models, and workflow automation. Yet the central safety problem is not simply whether a system can classify a prescription as high or low risk. It is whether the system can recognize when its output is insufficient for continued automation and responsibility must transfer to a pharmacist. Confidence-only automation is inadequate because a high score may coexist with poor calibration, missing clinical context, conflicting evidence, unusual patient–therapy combinations, vulnerability, or high-consequence decisions. This article proposes a non-empirical Prescription-Verification State Model linked to a Pharmacist-Takeover Trigger Framework. The model separates context assembly, automated checking, evidence reconciliation, clarification, pharmacist takeover, adjudication, disposition, and lifecycle monitoring. Takeover is organized around six proposed trigger classes: uncertainty, conflicting evidence, missing context, unusual combinations, vulnerable patients, and high-consequence decisions. A related takeover contract specifies the reason for escalation, unresolved information, urgency, completed checks, required professional action, fallback behavior, and documentation. Evaluation must extend beyond model discrimination to calibration, trigger sensitivity, missed takeover, unnecessary takeover, timing, workload displacement, interface comprehension, subgroup performance, and post-deployment change. Governance must define authority for activation, modification, suspension, requalification, and retirement. The framework is an original conceptual synthesis rather than a validated clinical system, guideline, or regulatory standard. Its trigger logic, state transitions, thresholds, and responsibility boundaries require empirical testing in specific pharmacy settings before decision use.

Keywords: Digital pharmacy, Artificial intelligence, Pharmacy practice, Medication safety, Human–AI collaboration, Clinical decision support

Introduction

Verification Is More than Error Classification

Medication-related processes contribute to preventable patient harm, but the safety relevance of a prescription-verification system cannot be reduced to whether it correctly labels an order as erroneous [1]. Artificial intelligence in pharmacy already encompasses heterogeneous functions, including prediction, prioritization, information retrieval, automation, and decision support, each with different evidential and professional implications [2]. "Verification" is therefore broader than error detection: it includes establishing patient and prescription identity, assembling the clinical context, checking internally and externally derived evidence, recognizing unresolved uncertainty, determining whether a proposed action remains within system coverage, and assigning responsibility for the next decision.

Existing medication AI has concentrated substantially on optimizing alerts and prioritizing prescriptions at elevated medication-error risk rather than defining an explicit responsibility-transfer contract [3, 4]. These functions may help organize attention, but prioritization does not demonstrate that lower-ranked prescriptions are safe to process without professional review. Context also matters. Indication documentation is often incomplete or poorly integrated into prescribing systems, limiting the capacity of computerized support to evaluate dose, interaction, duplication, and therapeutic appropriateness in relation to clinical purpose [5]. Similar limitations arise when renal or hepatic function, allergy details, pregnancy status, laboratory trends, prior response, current medicines, or the intended treatment duration are absent, stale, or internally inconsistent.

Address for correspondence: Sofía Martínez, Department of Prescription Verification and AI, Faculty of Pharmacy, Polytechnic University of Madrid, Madrid, Spain.
sofia.martinez@upm.es

How to cite this article: Martínez S, Gómez C, Navarro L. Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over. Arch Pharm Pract. 2025;16(1):69-77. https://doi.org/10.51847/acWJvHeaeM

This article therefore treats prescription verification as a sequence of conditional decision states rather than a single classifier output. Its original contribution is a proposed Prescription-Verification State Model and Pharmacist-Takeover Contract. The model asks not only "What risk does the system predict?" but also "Is the context adequate, are evidence sources concordant, is the case within validated coverage, what are the consequences of error, and who is authorized to decide next?" The contract defines when automated support must stop progressing, what evidence must be transferred, how urgently pharmacist review is required, and how the final disposition is recorded.

The framework does not assume that every uncertainty requires interruption or that every pharmacist review prevents harm. Excessive takeover could create delay, fatigue, and queue displacement. Conversely, permissive automation could create false reassurance or allow high-consequence errors to pass. The design problem is therefore a boundary-setting problem: preserving useful automation while making uncertainty, conflict, missingness, vulnerability, and consequence visible enough to trigger timely professional control.

Failure of Confidence-Only Automation

Confidence-only automation is defined here as a design pattern in which a model score, probability, or threshold is treated as sufficient authorization to continue or complete verification. That pattern conflates distinct properties. Decision use requires uncertainty characterization, uncertainty communication, and calibration rather than reliance on a raw confidence score [6-8]. A model may be confident because the input resembles its training data while still being wrong because the target was poorly specified, relevant context was omitted, or the clinical environment has changed. Conversely, a low-confidence output may reflect ambiguity that is clinically trivial or readily resolved. Confidence is therefore one signal, not a transferable measure of medication safety.

Explanations do not repair this problem automatically. A plausible feature attribution or rationale can make an output easier to inspect without establishing that the model is valid, causally appropriate, or safe for the intended decision [9]. Human review is also not an infallible backstop. Automation bias and verification complexity can lead users to accept incorrect advice, neglect contradictory information, or perform only superficial checking when the system appears authoritative [10]. The relevant design question is not whether a human remains nominally "in the loop," but whether the workflow creates a meaningful opportunity, sufficient information, and clear authority for independent adjudication.

The proposed framework therefore makes automation conditional on several jointly satisfied requirements: adequate and current context; authorized intended use; acceptable calibration and coverage for the case; no material conflict among rules, models, patient data, and treatment history; no unresolved unusual combination; no vulnerability condition requiring professional interpretation; and no high-consequence decision that local governance has designated for mandatory review. Failure of any requirement does not necessarily imply that the prescription is wrong. It means that the system lacks sufficient grounds to continue without clarification or pharmacist takeover.

Alternative designs may use universal pharmacist review, selective automation, tiered clarification, or non-interruptive decision support. The framework does not presume that one configuration is always superior. It proposes that the authority boundary be explicit, testable, and resistant to the false inference that confidence alone establishes readiness for automated progression.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Table 1. Definitions, components, relationships, and intended uses for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Confidence and calibration	A high score may be mistaken for safety	Predicted probability, calibration evidence, decision threshold	**Proposed:** confidence informs but does not authorize progression; calibration is assessed separately [8]	Reduces score-to-decision conflation	External and local calibration across relevant cases	Miscalibration and threshold instability	No score alone establishes safe non-takeover
Uncertainty communication	Users may not recognize system limits	Predictive uncertainty, data quality, coverage	**Proposed:** uncertainty is displayed with its source and possible consequence [7]	Supports informed review	Comprehension and response testing	False precision or vague warnings	Displayed uncertainty is not proof of validity

Context completeness	Clinically material information may be absent	Indication, organ function, allergies, pregnancy, laboratory and treatment history	**Proposed:** missing material context creates clarification or takeover [5]	Prevents unsupported contextual inference	Completeness, freshness, provenance, and clinical adjudication	Missing, stale, or incorrect data	Completeness rules require local validation
Evidence conflict	Rules, models, records, or history may disagree	Rule outputs, model outputs, patient record, prior response	**Proposed:** unresolved material conflict blocks automated progression	Makes disagreement actionable	Conflict reference standards and pharmacist adjudication	Silent precedence or contradictory advice	Conflict does not identify the correct action by itself
Unusual combination and coverage	A case may lie outside ordinary or validated use	Drug–dose–route–patient pattern and coverage indicators	**Proposed:** weak coverage or unusual combinations increase review requirements [6]	Limits unsupported extrapolation	Rare-case and out-of-distribution evaluation	Novelty missed or over-triggered	"Unusual" is model- and setting-dependent
Vulnerability and consequence	Low-probability errors may remain unacceptable	Patient susceptibility, severity, reversibility, monitoring capacity	**Proposed:** consequence can mandate takeover despite high confidence	Aligns authority with tolerance for error	Clinical-severity and subgroup validation	Stereotyping or incomplete categories	No universal vulnerability list is asserted
Pharmacist takeover	An alert may not clearly transfer responsibility	Trigger reason, urgency, unresolved evidence, completed checks	**Proposed:** a takeover contract transfers decision authority and information	Creates an auditable professional handoff	Workflow, human-factors, delay, and safety evaluation	Nominal review, unclear ownership, queue delay	Takeover activation does not prove improved outcome
Monitoring and requalification	Performance and context may change	Incidents, missed triggers, false triggers, workload, system change	**Proposed:** monitoring can modify, suspend, or requalify trigger logic [10]	Supports lifecycle control	Prospective surveillance and change-control evidence	Undetected deterioration or unsafe adaptation	Initial implementation does not establish continuing validity

Prescription-Verification Workflow and Decision States

A verification workflow requires explicit states because the meaning of a system signal depends on what has already been checked, what remains unresolved, and which action follows. Prospective evidence on medication-related decision-support overrides shows that overrides are heterogeneous: some reflect appropriate clinical judgment, whereas others expose unsafe responses or poorly designed alerts [11]. A binary alert–override model therefore conceals important distinctions among clarification, justified nonacceptance, unresolved conflict, professional adjudication, and silent bypass.

Risk prediction can be useful near the beginning of the workflow by prioritizing prescriptions or patients for attention [12]. In the proposed model, however, prioritization is not equivalent to authorization. A high-risk classification may accelerate pharmacist review, whereas a low-risk classification remains conditional on context completeness, model coverage, evidence agreement, and consequence. The relevant unit of safety is the combined human–technical process, not the classifier alone.

Usability is likewise part of task performance. A medication-related decision-support application can affect how efficiently users find relevant information, interpret warnings, and complete actions, so evaluation must include interaction and workflow effects rather than only algorithmic accuracy [13]. Poorly timed requests, fragmented evidence, hidden provenance, or unclear responsibility can turn an otherwise informative output into delay, workaround, or false reassurance.

The proposed Prescription-Verification State Model contains nine states:

S0—Prescription intake and identity assembly: links the correct patient, prescriber, medicine, strength, dose, route, frequency, duration, and available indication.

S1—Context-completeness assessment: determines whether required clinical information is present, current, attributable, and internally coherent.

S2—Automated mechanism checks: applies authorized rules, interaction resources, dose checks, allergy checks, and predictive models.

S3—Evidence reconciliation: examines agreement, provenance, uncertainty, applicability, and coverage.

S4—Automation-supported proceed: permits continued processing only within the validated intended use and when no takeover trigger applies.

S5—Clarification hold: pauses progression while missing or ambiguous but potentially resolvable information is obtained.

S6—Mandatory pharmacist takeover: transfers decision authority to a pharmacist.

S7—Pharmacist adjudication and disposition: records approval, modification, clarification, rejection, referral, or continued hold.

S8—Monitoring and requalification: evaluates incidents, misses, excessive triggers, delays, workload effects, subgroup performance, drift, and system change.

Clinical decision support is sociotechnical: benefits and risks arise from interactions among data, software, users, tasks, organizations, and implementation conditions [14]. Accordingly, transitions between the proposed states must be governed rather than inferred from a single output. A prescription cannot enter S4 merely because no rule fired; absence of a signal may reflect incomplete data, a knowledge-base gap, or weak model coverage. Nor should every deviation automatically enter S6. S5 can resolve nonurgent missingness without imposing an interruptive takeover, provided delay does not create material risk.

The state model therefore separates "no problem detected," "sufficient grounds to proceed," "additional information required," and "professional authority required." It also separates activation of takeover from completed review: a trigger has not protected the patient when the pharmacist cannot access the evidence, the queue delays action, or responsibility remains ambiguous.

These states are an original synthesis, not an empirically complete taxonomy. Their definitions, permitted transitions, responsibility assignments, and fallback behavior require validation across community, hospital, ambulatory, specialty, and remote-pharmacy settings.

Pharmacist-Takeover Trigger Framework

Drug-allergy checking illustrates why detection cannot be equated with disposition. Allergy decision support varies in knowledge representation, specificity, and clinical relevance, leaving cases in which a warning, absent warning, or ambiguous record still requires professional interpretation [15].

Figure 1 presents the original conceptual synthesis for prescription-verification workflow with pharmacist-takeover triggers, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

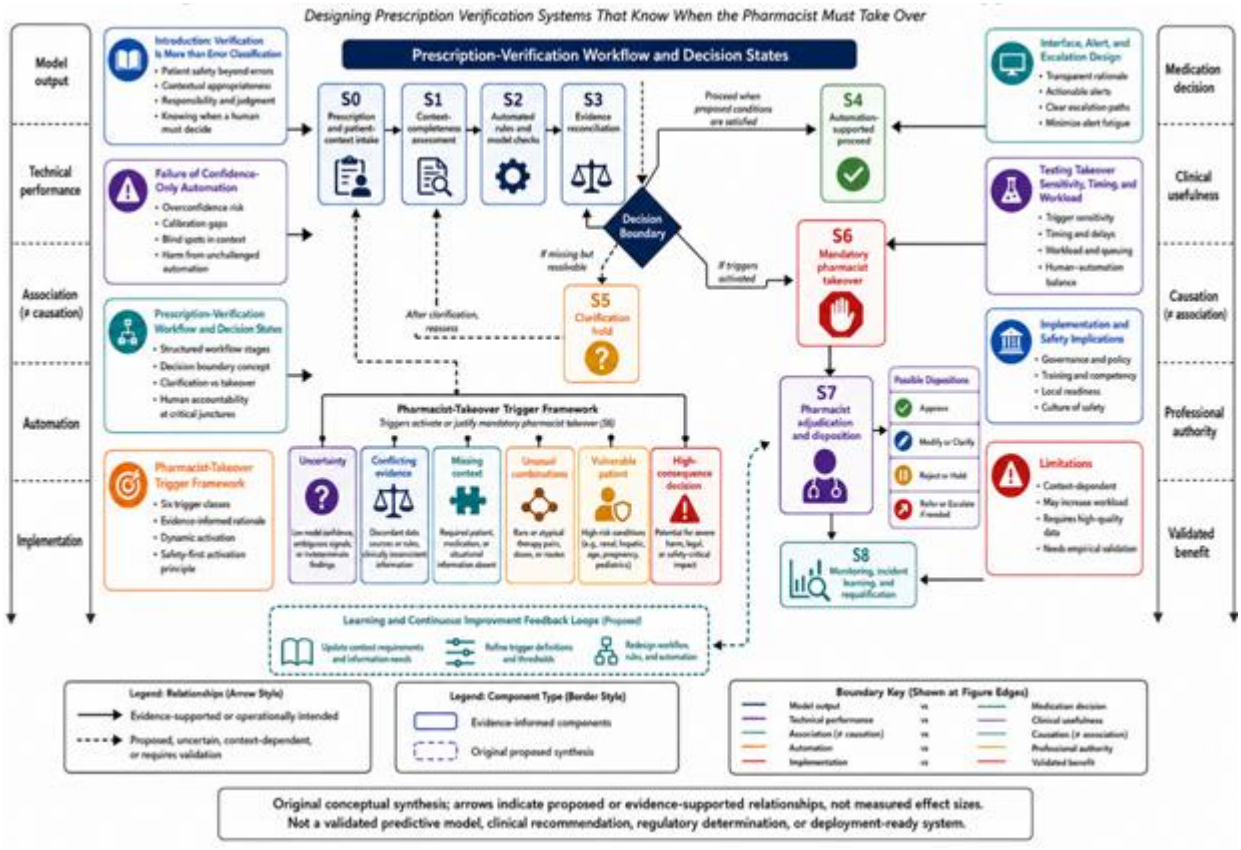

Figure 1. Prescription-Verification Workflow with Pharmacist-Takeover Triggers

The first proposed trigger is uncertainty. It includes poor calibration, unstable output, weak coverage, ambiguous records, or uncertainty that cannot be reduced within the available workflow. It does not mean every low-confidence output is unsafe; activation depends on whether uncertainty is material to the pending decision.

The second trigger is conflicting evidence. Context-aware drug–drug interaction screening demonstrates that relevance can change when patient and clinical information is incorporated [16]. In the proposed framework, disagreement among rules, model outputs, laboratory values, treatment history, allergy data, or documented indication requires reconciliation. A material unresolved conflict activates takeover, whereas a documentation discrepancy that can be resolved safely may enter clarification hold.

The third trigger is missing context. Material missingness is defined relative to the decision rather than a generic completeness score. Missing indication, organ function, pregnancy status, concurrent treatment, allergy detail, laboratory results, or monitoring data should activate clarification or takeover when the absent information could alter dose, suitability, interaction interpretation, monitoring, or urgency.

The fourth trigger is unusual combinations. This category covers cases outside validated or familiar coverage, including uncommon drug–dose–route patterns, rare comorbidity combinations, contradictory treatment histories, or atypical patient–therapy relationships. Unusualness is not evidence that a prescription is wrong. It is a boundary marker against unsupported extrapolation.

The fifth trigger is vulnerable patients, and the sixth is high-consequence decisions. Evidence that alert timing and screening intervals affect burden shows that safety cannot be separated from when and how review is requested [17]. Medication-related harm among older adults further supports treating patient susceptibility and potential consequence as dimensions distinct from predicted probability [18]. The framework therefore permits mandatory takeover when tolerance for error is low, even when estimated risk is low or model confidence is high.

The proposed Pharmacist-Takeover Contract accompanies every transition into S6. It communicates:

- the trigger class and triggering evidence;
- the affected prescription element;
- missing, conflicting, uncertain, or unusual information;
- evidence provenance and timestamp;
- estimated urgency and potential consequence;
- checks already completed and actions not completed;
- the required pharmacist role;
- permissible dispositions;
- fallback behavior if review is delayed or unavailable; and
- the final adjudication and rationale.

The contract is intended to prevent an alert from being mistaken for a completed handoff. It also makes escalation auditable by distinguishing system detection, transfer of authority, pharmacist receipt, completed adjudication, and final disposition.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
S0–S1: intake and context	Establish identity and required context	Prescription, patient, indication, history, and current clinical data	Determines whether subsequent checks are interpretable	Completeness, freshness, provenance, correct patient linkage	Wrong-patient association, stale data, silent missingness	Pharmacy and information governance define required fields	Populated fields do not by themselves establish adequate context
S2: automated checks	Apply authorized rules and models	Structured context enters checks; outputs enter reconciliation	Produces evidence rather than final authority	Technical validity within the intended use	False negative, false positive, or weak coverage	Model and knowledge-base owners maintain versions and limitations	Technical performance does not establish clinical usefulness
S3: evidence reconciliation	Identify agreement, conflict, and uncertainty	Rules, models, records, history, and provenance are compared	Routes the case to proceed, clarification, or takeover	Conflict detection, source traceability, and reproducibility	Silent precedence or unresolved contradiction	Pharmacist-led governance defines material conflict	Association does not determine the correct medication action
S4: automation-supported proceed	Continue only when proposed conditions are satisfied	Reconciled evidence and absence of an active trigger	Conditional transition rather than autonomous approval	Appropriate non-takeover and absence of missed mandatory cases	False reassurance or unauthorized progression	Organization defines the permissible scope of automation	"Proceed" is not equivalent to independent clinical authorization

S5: clarification hold	Obtain resolvable information	Missing item, request, response, and updated context	Returns to S1 or escalates to S6	Resolution quality, delay, and handoff completeness	Indefinite hold, informal workaround, or lost request	Assigned staff obtain and verify information	Clarification is unsuitable when delay creates material risk
S6–S7: takeover and adjudication	Transfer authority and document disposition	Trigger contract moves to the pharmacist; rationale returns to the record	Produces approval, modification, clarification, rejection, referral, or hold	Trigger sensitivity, response time, comprehension, and traceability	Nominal review, queue delay, unclear ownership, or automation bias [10]	Pharmacist adjudicates; organization ensures staffing and fallback	Takeover activation alone does not establish improved safety
S8: monitoring and requalification	Detect deterioration, inequity, burden, and new failure	Incidents, misses, false triggers, workload, and system changes	Feeds redesign, suspension, or requalification	Timely detection and corrective control	Normalized workaround or unmanaged change	Multidisciplinary governance owns surveillance and change control	Continued operation does not prove continuing validity
Trigger framework	Convert uncertainty, conflict, missingness, unusualness, vulnerability, and consequence into action	Trigger evidence is combined with urgency and consequence	Activates S5 or S6 under proposed local rules	Trigger-class sensitivity, unnecessary takeover, and subgroup behavior	Over-triggering, under-triggering, or omitted trigger categories	Local clinical governance validates thresholds and exceptions	Trigger classes, thresholds, and exceptions remain proposed

The trigger framework is not a clinical rule set. It supplies categories and responsibility-transfer fields that may be operationalized differently across settings. Hard triggers, conditional triggers, escalation timing, permitted exceptions, and fallback arrangements require local empirical validation, human-factors testing, workflow assessment, and governance authorization.

Interface, Alert, and Escalation Design

A takeover interface must communicate transferred authority, not merely another warning. Interruptive medication alerts have variable effects, so interruption should be reserved for conditions in which bypass or delay could materially affect the prescription decision [19]. Lower-consequence uncertainty may enter clarification or queued review.

Role-tailored interaction design may reduce medication-alert fatigue [20]. Pharmacist-facing escalation should show the trigger, affected prescription element, unresolved evidence, provenance, urgency, completed checks, and permitted dispositions rather than only a risk score.

Workload, complexity, and repeated alerts influence alert fatigue [21]. Queue position, response time, staffing, and repeated-trigger burden are safety variables. Mandatory triggers should remain visible until resolved, while redundant signals may be consolidated without hiding distinct high-consequence concerns.

Collaborative alert review can support accountable removal or redesign of low-value decision support [22]. Interface changes require testing for comprehension, verification, response quality, workload displacement, and workarounds.

Testing Takeover Sensitivity, Timing, and Workload

Evaluation should document live human–AI interaction, AI-specific trial conduct, and prespecified handling of inputs, outputs, errors, and professional involvement [23-25]. The unit is the prescription–system–pharmacist episode. Outcomes should distinguish correct, missed, unnecessary, delayed, and incomplete takeover.

Prediction-model reporting should define intended use, targets, data sources, missing-data handling, performance, and limitations [26]. These properties are insufficient alone. Testing must also examine trigger-class sensitivity, conflict detection, pharmacist comprehension, fallback, subgroup behavior, and whether escalation occurs before an irreversible action.

Sensitivity must be evaluated under nominal and stressed workload. Long queues, reduced staffing, or competing tasks may delay escalation; excessive takeover may displace attention from more consequential prescriptions. No universal performance or workload threshold is proposed.

Implementation and Safety Implications

Clinical impact depends on generalizability, workflow integration, data quality, and real-use conditions [27]. **Figure 2** presents the original conceptual synthesis for takeover testing matrix across sensitivity, timing, and workload, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

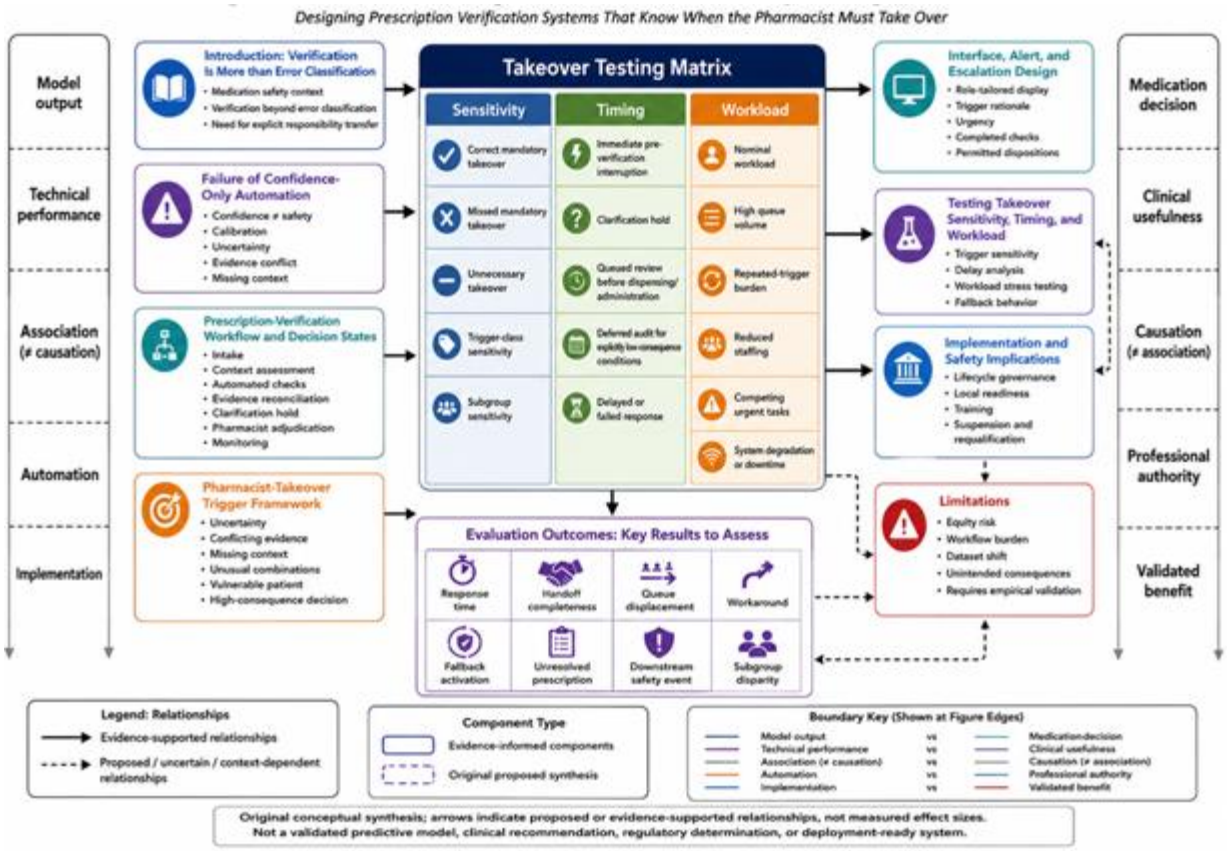

Figure 2. Takeover Testing Matrix across Sensitivity, Timing, and Workload

Responsible health-care machine learning requires lifecycle monitoring and mechanisms for intervention when assumptions fail [28]. Governance must control intended use, updates, incidents, suspension, requalification, and retirement.

Adoption and response timing are workflow-mediated [29]. Evaluation should distinguish alert exposure, pharmacist receipt, completed adjudication, and final disposition. High use is not equivalent to safe use.

Implementation should be assessed across the intervention, setting, individuals, and process [30]. Local readiness requires staffing, role clarity, training, downtime procedures, auditable provenance, and authority to suspend unsafe automation.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article Designing Prescription Verification Systems That Know When the Pharmacist Must Take Over.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary
Intended use	Multidisciplinary safety body	Approve scope and exclusions	Intended-use dossier	Out-of-scope use	Recurrent scope breach	Restrict or suspend	Approval does not establish benefit
Trigger and interface control	Pharmacist-led informatics team	Change logic and presentation	Versioned rationale	Misses, false takeover, or confusion [22]	Material recurring failure	Redesign and retest	Configuration remains locally validated
Lifecycle safety	System and pharmacy owners	Review, rollback, or suspend	Incident and update records	Drift or data failure [28]	Assumptions no longer hold	Correct and requalify	Operation does not prove continuing validity

Workflow capacity	Pharmacy leadership	Assign queues and fallback	Staffing and downtime plan	Delayed takeover [29]	Inadequate response capacity	Reallocate or activate fallback	Technology cannot replace unavailable expertise
Implementation	Implementation leads	Assess fit and adaptation	Training and readiness record	Persistent process failure [30]	Setting–system mismatch	Adapt or withdraw	Adoption is not proof of safety
Equity	Safety and patient representatives	Require subgroup review	Subgroup audit	Differential misses or delay	Unjustified disparity	Investigate, redesign, or suspend	One fairness metric cannot settle acceptability

Limitations

Proxy variables and target definitions can produce inequitable decisions despite strong apparent performance [31]. Missing data, uneven documentation, differential pharmacist access, or incomplete vulnerability categories may reproduce such disparities. Fairness metrics cannot resolve ethical and distributive questions alone [32].

Machine learning can also generate unintended clinical and organizational effects [33]. Takeover logic may increase delay, superficial review, workarounds, or attention displacement. Dataset shift may invalidate assumptions after implementation [34]. The proposed states, triggers, thresholds, and governance responsibilities therefore require external, prospective, subgroup-sensitive, and lifecycle evaluation.

Conclusion

Prescription verification is not completed when a system classifies risk or produces a confident recommendation. It requires context assembly, evidence reconciliation, responsibility allocation, professional adjudication, and documented disposition.

The proposed Prescription-Verification State Model and Pharmacist-Takeover Contract identify uncertainty, conflicting evidence, missing context, unusual combinations, vulnerable patients, and high-consequence decisions as takeover classes. They distinguish alert generation from transfer of authority and takeover activation from completed review.

The conceptual framework establishes no universal threshold, mandate, improved outcome, regulatory acceptance, or deployment readiness. Its value lies in defining testable responsibility boundaries and failure conditions. Operational use requires local validation, pharmacist and patient involvement, implementation assessment, accountable governance, and continuing authority to modify or suspend the system.

ACKNOWLEDGMENTS: None
CONFLICT OF INTEREST: None
FINANCIAL SUPPORT: None
ETHICS STATEMENT: None

References

1. Panagioti M, Khan K, Keers RN, Abuzour A, Phipps D, Kontopantelis E, et al. Prevalence, severity, and nature of preventable patient harm across medical care settings: Systematic review and meta-analysis. BMJ. 2019;366. doi:10.1136/bmj.l4185
2. Chalasani SH, Syed J, Ramesh M, Patil V, Pramod Kumar TM. Artificial intelligence in the field of pharmacy practice: A literature review. Explor Res Clin Soc Pharm. 2023;12:100346. doi:10.1016/j.rcsop.2023.100346
3. Graafsma J, Murphy RM, van de Garde EMW, Karapinar-Çarkit F, Derijks HJ, Hoge RHL, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411-22. doi:10.1093/jamia/ocae076
4. 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):1688-94. doi:10.1093/jamia/ocaa154
5. Feather C, Appelbaum N, Darzi A, Dean Franklin B. Indication documentation and indication-based prescribing within electronic prescribing systems: A systematic review and narrative synthesis. BMJ Qual Saf. 2023;32(6):357-68. doi:10.1136/bmjqs-2022-015452
6. Begoli E, Bhattacharya T, Kusnezov D. The need for uncertainty quantification in machine-assisted medical decision making. Nat Mach Intell. 2019;1(1):20-3. doi:10.1038/s42256-018-0004-1
7. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
8. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW; Topic Group 'Evaluating Diagnostic Tests and Prediction Models' of the STRATOS Initiative. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
9. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9
10. Lyell D, Coiera E. Automation bias and verification complexity: A systematic review. J Am Med Inform Assoc. 2017;24(2):423-31. doi:10.1093/jamia/ocw105
11. Wong A, Amato MG, Seger DL, Rehr C, Wright A, Slight SP, et al. Prospective evaluation of medication-related clinical decision support over-rides in the intensive care unit. BMJ Qual Saf. 2018;27(9):718-24. doi:10.1136/bmjqs-2017-007531
12. Abdo A, Gallay L, Vallecillo T, Clarenne J, Quillet P, Vuiblet V, et al. A machine learning-based clinical predictive tool to identify patients at high risk of medication errors. Sci Rep. 2024;14(1):32022. doi:10.1038/s41598-024-83631-w
13. Nanji KC, Garabedian PM, Langlieb ME, Rui A, Tabayoyong LL, Sampson M, et al. Usability of a perioperative medication-related clinical decision support software application: A randomized controlled trial. J Am Med Inform Assoc. 2022;29(8):1416-24. doi:10.1093/jamia/ocac035
14. Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: Benefits, risks, and strategies for success. NPJ Digit Med. 2020;3:17. doi:10.1038/s41746-020-0221-y
15. Légat L, Van Laere S, Nyssen M, Steurbaut S, Dupont AG, Cornu P. Clinical decision support systems for drug allergy checking: Systematic review. J Med Internet Res. 2018;20(9). doi:10.2196/jmir.8206

16. Muylle KM, Gentens K, Dupont AG, Cornu P. Evaluation of an optimized context-aware clinical decision support system for drug-drug interaction screening. Int J Med Inform. 2021;148:104393. doi:10.1016/j.ijmedinf.2021.104393
17. Van De Sijpe G, Walgraeve K, Van Laer E, Quintens C, Machiels C, Foulon V, et al. The impact of customized screening intervals on the burden of drug-drug interaction alerts: An interrupted time series analysis. J Med Syst. 2024;48(1):93. doi:10.1007/s10916-024-02113-8
18. Parekh N, Ali K, Davies JG, Stevenson JM, Banya W, Nyangoma S, et al. Medication-related harm in older adults following hospital discharge: Development and validation of a prediction tool. BMJ Qual Saf. 2020;29(2):142-53. doi:10.1136/bmjqs-2019-009587
19. Page N, Baysari MT, Westbrook JI. A systematic review of the effectiveness of interruptive medication prescribing alerts in hospital computerized order entry systems to change prescriber behavior and improve patient safety. Int J Med Inform. 2017;105:22-30. doi:10.1016/j.ijmedinf.2017.05.011
20. Hussain MI, Reynolds TL, Zheng K. Medication safety alert fatigue may be reduced via interaction design and clinical role tailoring: A systematic review. J Am Med Inform Assoc. 2019;26(10):1141-9. doi:10.1093/jamia/ocz095
21. Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R; HITEC Investigators. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. BMC Med Inform Decis Mak. 2017;17(1):36. doi:10.1186/s12911-017-0430-8
22. McCoy AB, Russo EM, Johnson KB, Addison B, Patel N, Wanderer JP, et al. Clinician collaboration to improve clinical decision support: The Clickbusters initiative. J Am Med Inform Assoc. 2022;29(6):1050-9. doi:10.1093/jamia/ocac027
23. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al.; DECIDE-AI Expert Group. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
24. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
25. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ; SPIRIT-AI and CONSORT-AI Working Group. Guidelines for clinical trial protocols for interventions involving artificial intelligence: The SPIRIT-AI extension. Nat Med. 2020;26(9):1351-63. doi:10.1038/s41591-020-1037-7
26. Collins GS, Dhiman P, Andaur Navarro CL, Ma J, Hooft L, Reitsma JB, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
27. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2
28. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6
29. Henry KE, Adams R, Parent C, Soleimani H, Sridharan A, Johnson L, et al. Factors driving provider adoption of the TREWS machine learning-based early warning system and its effects on sepsis treatment timing. Nat Med. 2022;28(7):1447-54. doi:10.1038/s41591-022-01895-z
30. Damschroder LJ, Reardon CM, Widerquist MAO, Lowery J. The updated Consolidated Framework for Implementation Research based on user feedback. Implement Sci. 2022;17(1):75. doi:10.1186/s13012-022-01245-0
31. 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):447-53. doi:10.1126/science.aax2342
32. McCradden MD, Joshi S, Mazwi M, Anderson JA. Ethical limitations of algorithmic fairness solutions in health care machine learning. Lancet Digit Health. 2020;2(5). doi:10.1016/S2589-7500(20)30065-0
33. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA. 2017;318(6):517-8. doi:10.1001/jama.2017.7797
34. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041