Original Article

Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks

Hans Müller[1]*, Anna Schmidt[1], Thomas Weber[2]

[1]Department of Pharmacy Automation and Task Delegation, Faculty of Pharmaceutical Sciences, University of Bonn, Bonn, Germany.
[2]Department of Clinical Pharmacy and Digital Boundaries, Faculty of Pharmacy, LMU Munich, Munich, Germany.

Abstract

Automation in community pharmacy is often discussed as though medication tasks can be transferred to digital systems whenever adequate technical performance is demonstrated. This framing overlooks differences in clinical consequence, uncertainty, contextual dependence, patient vulnerability, error detectability, and the possibility of correcting an inappropriate action. This article develops an original, non-empirical delegation-boundary model for determining when automation may proceed, when professional review is required, and when delegation should stop. Delegability is defined as a conditional property of a specific medication-task instance rather than an inherent characteristic of an entire service, technology, or occupational role. The proposed model assesses six interacting dimensions: clinical consequence, ambiguity, reversibility, contextual dependence, patient vulnerability, and detectability of error. These dimensions inform three conditional states. Proceed permits bounded automated execution under defined safeguards; Escalate requires review by an authorized professional before action; and Stop retains human control because necessary context, recovery capacity, accountability, or safety conditions are insufficient. The architecture also includes task specification, uncertainty communication, professional authority, monitoring, incident review, and boundary reassessment. Application is intended to distinguish automatable components from judgment-dependent components within dispensing, verification, counselling, and follow-up rather than classify whole services as automatable or non-automatable. Empirical evaluation would require task-specific technical validation, workflow simulation, human-factors assessment, medication-safety analysis, patient and equity evaluation, external validation, and lifecycle monitoring. The model is a proposed conceptual synthesis, not a clinical recommendation, legal standard, validated scoring system, or deployment-ready framework. Its principal contribution is to reposition pharmacy automation as a conditional delegation decision requiring evidence about both system performance and the consequences of transferring professional work.

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

Introduction

The Unresolved Boundary of Pharmacy Automation

Artificial intelligence and related digital systems are being applied across pharmacy activities, including information retrieval, medication review, decision support, dispensing processes, patient communication, and operational management. The available pharmacy literature nevertheless describes substantial variation in the purposes, maturity, evaluation, and practical integration of these applications [1].

"Pharmacy automation" therefore does not identify one intervention or one form of delegation. It may refer to transferring physical execution, data processing, prioritization, interpretation, communication, or decision authority, each of which creates different professional and medication-safety consequences.

Technological transformation may redistribute pharmacy work rather than simply accelerate a stable set of existing tasks [2]. A system may reduce manual transcription while creating new verification work, convert direct checking into exception management, or move professional attention from routine processing to ambiguous cases. Whether such redistribution is beneficial cannot be inferred from the presence of automation alone. It depends on what information the system uses, what action it performs, what responsibility

Address for correspondence: Hans Müller, Department of Pharmacy Automation and Task Delegation, Faculty of Pharmaceutical Sciences, University of Bonn, Bonn, Germany.
hans.mueller@uni-bonn.de

This is an open-access article distributed under the terms of the Creative Commons Attribution-Non Commercial-Share Alike 4.0 License, which allows others to remix, tweak, and build upon the work non commercially, as long as the author is credited and the new creations are licensed under the identical terms.

How to cite this article: Müller H, Schmidt A, Weber T. Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks. Arch Pharm Pract. 2025;16(1):50-9. https://doi.org/10.51847/lgc8J2bCwU

remains with the pharmacist, and whether the surrounding workflow permits meaningful oversight.

Clinical decision support illustrates this distinction. Its value depends not only on computational performance but also on usability, workflow integration, data quality, maintenance, evaluation, and management of unintended consequences [3]. A technically accurate system may still interrupt care at the wrong point, suppress relevant context, create excessive alerts, or induce reliance that weakens independent judgment. Conversely, a modestly complex tool may offer value when it performs a bounded function, communicates uncertainty, and supports rather than obscures professional reasoning.

Machine-learning systems can also produce unintended harms that are not represented by average performance measures [4]. Errors differ in their clinical importance, visibility, reversibility, and distribution across patients. The unresolved question is therefore not whether community pharmacy should use automation, but where delegation should stop for a particular medication task under particular conditions. This article proposes that the answer requires a task-specific boundary model integrating technical evidence, professional judgment, medication-safety consequences, patient context, workflow conditions, and continuing accountability. The model is intended to organize evaluation and research. It does not establish that a task or technology is safe, authorized, or ready for implementation.

Medication Tasks Are Not Equally Delegable

Medication work should not be divided only into broad service labels such as dispensing, verification, counselling, and follow-up. Each service contains component tasks that differ in informational requirements, judgment, patient interaction, consequences of error, and dependence on local circumstances. Delegability is therefore proposed as a property of a specific task instance: a defined action, involving a defined patient and medication context, conducted within a defined technical and organizational environment.

The existence of nominal human oversight does not by itself make delegation safe. Automation-bias research indicates that verification may become unreliable when the automated output is difficult to reconstruct or independently check [5]. A pharmacist who receives a recommendation without access to its relevant assumptions, missing data, uncertainty, or alternative interpretations may be positioned as a reviewer without being given a realistic capacity to review. Human involvement must consequently be distinguished from effective verification.

Pharmacists' electronic work also extends beyond selecting or accepting displayed information. It can involve reconciling records, interpreting incomplete histories, coordinating with other professionals, identifying discrepancies, and connecting digital information with the patient's current circumstances [6]. These activities are not interchangeable. Structured transfer of an accurately identified prescription field may require little contextual interpretation, whereas determining whether a discrepancy is clinically meaningful may require access to treatment history, patient preferences, prescriber intent, and local care arrangements.

Medication-related automation can additionally change the type of error rather than eliminate error. Experimental evidence from electronic prescribing shows that erroneous automated cues may contribute to omission and commission errors [7]. This matters for delegation because the same system may reduce one class of manual mistake while creating reliance on incorrect recommendations or incomplete representations. Task assessment must therefore consider not only how often an error may occur, but also whether the error can be recognized before action and whether the workflow encourages independent checking.

Medication review provides a clear example of task heterogeneity. Decision-support architectures may integrate medication lists, clinical variables, guideline logic, visualization, and recommendations, yet professional interpretation remains necessary when evidence conflicts, information is incomplete, or patient priorities alter the meaning of a recommendation [8]. The article therefore rejects permanent classifications such as "medication review is automatable" or "counselling is non-automatable." A single service may contain components that can proceed automatically, components that require escalation, and components for which delegation should stop.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks.

Table 1. Definitions, components, relationships, and intended uses for the article Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Medication-task instance	Broad service labels conceal important	Intended action, patient, medication,	**Proposed:** delegability attaches to a defined task instance	Creates a precise unit for boundary assessment.	Task analysis and workflow observation	Tasks may be specified too	Classification of a service does not

	differences among actions.	timing, setting, responsible actor	rather than an entire pharmacy service.			broadly or omit hidden work.	determine every task within it.
Automation support	Assistance is often confused with transfer of authority.	System function, output, action permitted, professional role	Decision support, prioritization, data transfer, physical execution, and autonomous action are treated as distinct forms of support [3].	Prevents technical capability from being equated with professional delegation.	Functional testing and intended-use specification	Informal workarounds may expand system authority beyond its design.	A recommendation is not equivalent to authorization to act.
Delegation	Responsibility may become unclear when a system performs part of a medication task.	Scope of execution, decision authority, supervision, accountability	**Proposed:** delegation is the conditional transfer of task execution or decision authority while responsibility remains explicitly allocated.	Clarifies who may act, review, interrupt, and recover.	Role mapping and governance evidence	Nominal oversight may conceal ineffective or impossible review.	Human presence alone does not establish meaningful oversight [5].
Task decomposition	Composite services may contain routine and judgment-dependent elements.	Data collection, preparation, interpretation, communication, exception handling	Tasks are separated into components with different informational and professional requirements [6].	Supports selective rather than whole-service automation.	Cognitive task analysis and direct observation	Hidden coordination or contextual work may be omitted.	Automating one component does not establish delegability of the complete service.
Clinical consequence	Similar error rates may have very different implications.	Medication risk, severity, timing, affected patient, exposure duration	**Proposed:** higher foreseeable consequence narrows the conditions under which automation may proceed.	Aligns delegation with potential harm rather than average accuracy alone.	Task-specific safety analysis	Rare but serious consequences may be obscured by aggregate metrics.	No universal consequence threshold is proposed.
Ambiguity	Systems may produce one output despite incomplete or conflicting evidence.	Missing data, conflicting sources, uncertain intent, multiple reasonable interpretations	**Proposed:** material ambiguity moves a task toward escalation unless uncertainty can be safely contained.	Prevents uncertain outputs from being treated as determinate decisions.	Uncertainty evaluation and scenario testing	Hidden ambiguity may appear as unjustified confidence.	Unresolved material ambiguity is incompatible with unreviewed execution.
Reversibility	A correctable administrative action differs from an irreversible medication consequence.	Time to detection, ability to interrupt, corrective pathway, residual harm	**Proposed:** reversibility can widen the Proceed region only when detection and recovery are realistic.	Connects delegation to practical recovery capacity.	Recovery simulation and incident analysis	"Reversible" actions may still cause delay, confusion, or cumulative harm.	Theoretical correctability does not establish practical reversibility.
Contextual dependence	Medication decisions may depend on information outside the automated representation.	Patient history, local practice, staffing, workflow, prescriber intent, care setting	Context dependence increases when task meaning changes across patients or environments [6].	Identifies when standardized processing is insufficient.	Local validation and context-completeness assessment	Relevant contextual variables may be unrecorded or inaccessible.	A task may cross boundaries when its context changes.
Patient vulnerability	Average performance may conceal greater consequences for particular patients.	Frailty, cognition, language, health literacy, comorbidity, access, ability to report harm	**Proposed:** vulnerability lowers tolerance for ambiguity, low detectability, or weak recovery.	Introduces patient-specific safeguards into delegation decisions.	Patient-centred and subgroup evaluation	Vulnerability may be incompletely observed or reduced to crude proxies.	Vulnerability is not a reason for exclusion; it is a reason for stronger safeguards.
Detectability of error	An error that cannot be independently recognized may pass through nominal review.	Output transparency, independent data, reviewer expertise, workload, timing	Verification quality decreases when automated output is difficult to reconstruct or challenge [5].	Connects oversight to actual checking capacity.	Human-factors testing and error-seeding studies	Explanations may increase confidence without improving detection.	Low detectability can justify escalation or stopping even when error frequency appears low.

Proceed, Escalate, and Stop states	Binary automation decisions ignore intermediate and changing conditions.	Six dimensions, context sufficiency, recovery, professional availability	**Original proposed synthesis:** each task instance is assigned a conditional state rather than a permanent automation classification.	Provides a common language for bounded delegation.	Multistage empirical validation	States may be applied inconsistently without operational definitions.	The states are conceptual and not validated clinical thresholds.

Dimensions Determining Delegability

The proposed model contains six interacting dimensions. They are not intended as an additive score, and no dimension is assigned a universal weight. Their purpose is to expose reasons why adequate technical performance may still be insufficient for delegation.

Clinical consequence concerns the foreseeable importance of an erroneous, delayed, omitted, or inappropriate action. Consequence is shaped by the medication, patient condition, duration of exposure, availability of rescue, and point in the medication-use process. In high-stakes settings, reliance on opaque models requires stronger justification, particularly when an interpretable alternative could support meaningful examination of the decision [9]. Clinical consequence does not mean that every high-risk task must remain entirely manual. It means that stronger evidence, more effective verification, and narrower authority are required before automated execution can be justified.

Ambiguity describes uncertainty about the task, available evidence, patient state, prescriber intent, or appropriate response. It includes missing information, conflicting records, uncertain classification, and the presence of multiple reasonable interpretations. Medical machine-learning systems should communicate uncertainty rather than compress it into an apparently definitive output [10]. Within the proposed model, ambiguity may be managed through calibrated uncertainty presentation, abstention, requests for additional information, or professional escalation. A confidence value alone is insufficient unless users understand what it represents and how it should alter action.

Detectability of error concerns whether an incorrect output can be identified before clinically important harm occurs. Detectability depends on access to independent information, reviewer expertise, time, workflow conditions, and the visibility of the system's assumptions. Incorrect or biased AI advice can worsen human decisions, and explanatory material does not necessarily eliminate that influence [11]. A system may therefore require escalation when its errors are subtle, persuasive, or difficult to distinguish from clinically plausible recommendations. Detectability must be evaluated through observed human–system interaction rather than inferred from the existence of an explanation screen.

Patient vulnerability concerns the likelihood that a patient will experience greater harm, receive a less appropriate recommendation, or have less capacity to identify and recover from an error. Vulnerability may arise through clinical instability, frailty, cognitive or communication barriers, limited health literacy, fragmented records, polypharmacy, unequal access, or underrepresentation in development data. Algorithmic systems may reproduce inequity when a convenient proxy does not adequately represent actual need [12]. Patient vulnerability is therefore not reducible to demographic classification. It must include whether relevant needs are represented, whether communication is accessible, and whether the patient can challenge or correct the automated process.

Reversibility refers to the practical ability to interrupt, correct, and remediate an inappropriate action. Reversibility includes the time before harm, availability of a responsible responder, reliability of monitoring, and residual consequences after correction. Reprinting a label may be more reversible than releasing an inappropriate medicine, but even apparently administrative actions may create delay, duplication, privacy loss, or confusion. Reversibility can support Proceed only when recovery is realistic within the operating environment.

Contextual dependence captures the extent to which task meaning changes with patient circumstances, medication history, local workflow, staffing, information availability, professional relationships, and jurisdiction. A rule that is appropriate for a complete and internally consistent record may be inappropriate when adherence, treatment intent, recent hospitalization, or over-the-counter use is unknown. Contextual dependence also means that an automation boundary is not fixed permanently. The same nominal task may proceed in one encounter, escalate in another, and stop when the necessary context or competent reviewer is unavailable.

The dimensions interact. High consequence combined with low detectability may justify stopping even when ambiguity appears modest. High ambiguity may be manageable when the action is reversible and rapid professional escalation is available. Patient vulnerability may narrow the boundary despite unchanged average model performance. These interactions are proposed conceptual relationships requiring empirical investigation; they do not establish validated weights, mandatory thresholds, or universal task classifications.

Proposed Stop–Proceed–Escalate Boundary Model

The proposed model treats automation as a conditional delegation pathway rather than a binary choice between human and machine work. Evidence from other clinical domains suggests that human–computer performance

depends on how complementary capabilities are combined, not simply on replacing human judgment with algorithmic output [13]. In community pharmacy, the central unit is consequently the medication-task instance: what action is proposed, for whom, on what evidence, under which workflow conditions, and with what authority and recovery pathway.

Figure 1 presents the original conceptual synthesis for stop–proceed–escalate boundary model for community-pharmacy task delegation, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

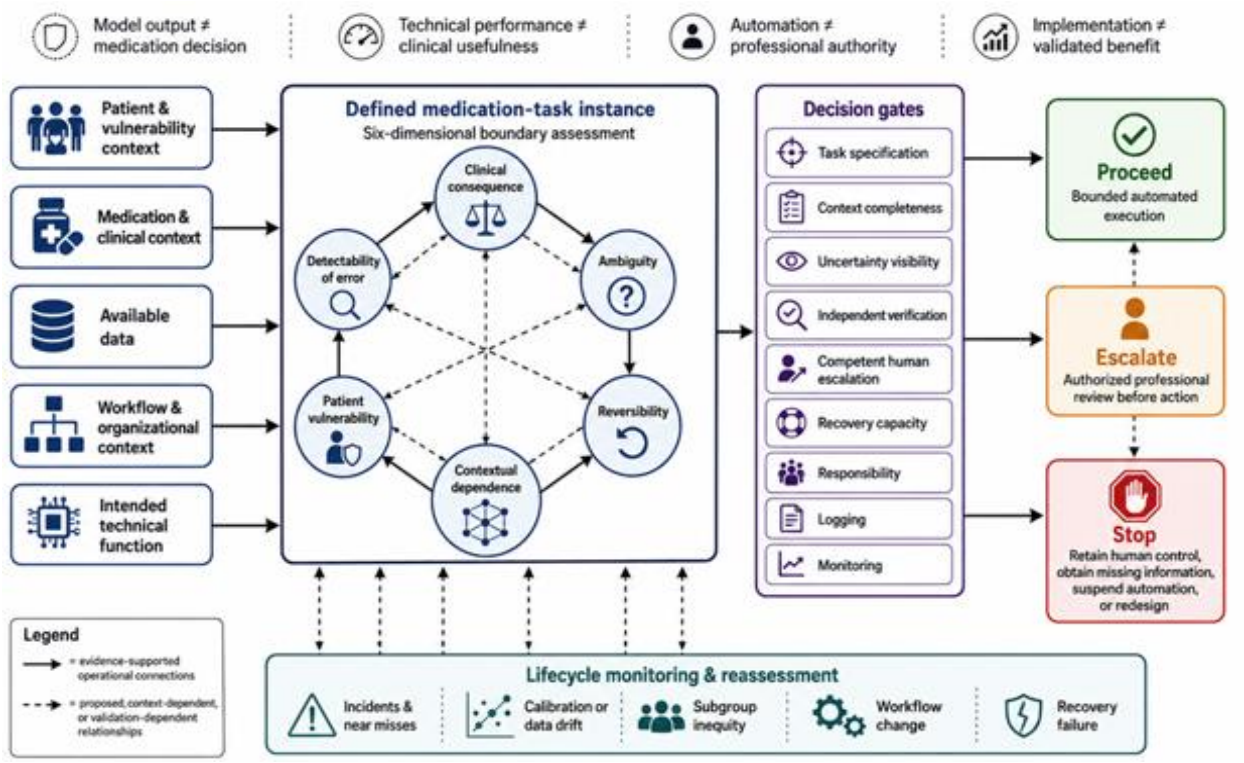

Figure 1. Stop–Proceed–Escalate Boundary Model for Community-Pharmacy Task Delegation.

The figure is an original conceptual synthesis developed for "Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks." A defined medication-task instance is assessed in relation to patient, medication, data, technical, workflow, and organizational context. Six proposed interacting dimensions inform decision gates concerning task specification, context completeness, uncertainty, independent verification, professional escalation, recovery, monitoring, and accountability. Proceed denotes conditional automation within a bounded scope; Escalate denotes required professional review before action; and Stop denotes that delegation is inappropriate under the present conditions. Monitoring, incidents, drift, inequity, and workflow change return the task to boundary reassessment. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

Within the model, Proceed does not mean that a technology is generally safe or that professional responsibility has been transferred. It means that a defined function may execute within a bounded intended use because the relevant context is sufficiently represented, foreseeable consequences are constrained, error detection and recovery are credible, and monitoring remains active. Trust must remain calibrated to demonstrated capability and context; clinician confidence or willingness to use a system does not by itself establish trustworthiness [14].

Escalate applies when automation may support information gathering, prioritization, or presentation but cannot responsibly complete the task. Triggers include material uncertainty, conflicting records, unrepresented patient circumstances, disagreement between the system and patient-reported information, inability to reconstruct the basis of an output, or a vulnerability that changes the consequence of error. Explanation is relevant only when it is usable for the particular professional, purpose, and decision context [15]. A technically detailed explanation that does not help the pharmacist detect an error or decide what action is warranted does not satisfy the escalation function.

Stop applies when delegation is inappropriate under present conditions. Examples include a high-consequence action with poor error detectability, an effectively irreversible action without a recovery pathway, missing information essential to safe interpretation, lack of competent professional escalation, known system malfunction or drift, and undefined responsibility for monitoring or remediation. Stop is not necessarily permanent. It may lead to obtaining further information, redesigning the process, narrowing the intended use, restoring manual handling, or conducting additional validation.

The model is enclosed by continuing governance because a delegation decision cannot be treated as a one-time technical authorization. Responsible healthcare machine learning requires attention to data, stakeholders, evaluation, deployment, and monitoring across the system lifecycle [16]. A task that initially proceeds may later require escalation or stopping when the population, medication environment, workflow, data source, model calibration, or recovery capacity changes.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Where Should Automation Stop in Community Pharmacy? A Boundary Model for Delegating Medication Tasks.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Intended-use specification	Define the exact task, permitted action, user, patient group, and operating conditions.	Governance and workflow requirements enter system design.	Constrains every later gate and state.	Task scope is observable, testable, and distinguishable from adjacent professional decisions.	Scope expands informally or hides composite judgment work.	Organization and system owner define and control use.	No Proceed state when the task or authority transferred is ambiguous.
Data and context acquisition	Assemble medication, patient, workflow, and organizational information required for the task.	Source data enter the boundary assessment.	Supports ambiguity, vulnerability, consequence, and detectability assessment.	Relevant information is sufficiently complete, timely, traceable, and locally applicable.	Missing, stale, inconsistent, or proxy-based information is treated as complete.	Pharmacy organization maintains data access and provenance.	Missing decision-critical context triggers Escalate or Stop.
Technical performance	Determine whether the system performs its specified computational or physical function.	Inputs are transformed into outputs, actions, or priorities.	Technical results inform but do not determine delegation.	Performance is task-specific, externally examined, calibrated where applicable, and tested under foreseeable variation.	Aggregate accuracy conceals severe error types or subgroup failures.	Developer validates function; organization verifies local fitness.	Technical performance alone cannot establish Proceed.
Uncertainty and abstention	Represent limitations in evidence or confidence.	Uncertainty accompanies the system output rather than being hidden.	Influences ambiguity and escalation gates.	Uncertain or out-of-scope cases are identified reliably and communicated meaningfully.	A single definitive output is produced despite inadequate evidence [10].	Developer designs uncertainty handling; pharmacist interprets its practical relevance.	Unmanaged material uncertainty is incompatible with unreviewed execution.
Human–AI interaction	Support calibrated reliance and independent review.	Outputs, rationale, context, and alternatives reach the responsible professional.	Determines whether escalation and verification are meaningful.	Users can detect seeded errors, challenge outputs, and act without undue reliance.	Persuasive presentation increases acceptance of an incorrect recommendation [11].	Organization provides training, time, interface governance, and escalation authority.	Nominal human oversight without effective checking capacity does not satisfy readiness.
Patient-vulnerability assessment	Identify circumstances that alter harm, communication, or recovery.	Patient-specific information modifies the delegation boundary.	Interacts with all six dimensions.	Subgroup performance, accessibility, communication needs, and recovery	Vulnerability is omitted, inferred from unreliable proxies, or used to exclude patients.	Pharmacist and organization ensure accessible human pathways.	Unmitigated vulnerability narrows or closes the Proceed region.

				capacity are evaluated.			
Proceed state	Permit bounded automated execution.	Authorized task output moves to execution and logging.	Remains subject to safeguards, monitoring, and reassessment.	Execution stays within intended use; errors are detectable; recovery and accountability are operational.	Automation expands beyond its approved task or suppresses exceptions.	Named professional and organizational owners retain oversight and remediation duties.	Proceed is conditional, local, and revocable rather than a general approval.
Escalate state	Require professional interpretation before action.	System output and relevant context are transferred to an authorized reviewer.	May return the task to Proceed, retain human control, or trigger Stop.	Escalations are timely, informative, actionable, and available when needed.	Excessive escalation creates overload, while insufficient escalation hides risk.	Organization ensures competent staffing, response pathways, and decision authority.	No readiness where escalation exists only nominally or cannot be completed safely.
Stop state	Prevent inappropriate delegation.	Automated execution is blocked or suspended.	Leads to human handling, additional information, redesign, or revalidation.	Stop conditions activate reliably and cannot be bypassed without accountable review.	Commercial, workload, or usability pressures normalize unsafe override.	Pharmacist and organization retain authority to suspend use.	High consequence combined with low detectability, absent recovery, or missing context supports Stop.
Monitoring and reassessment	Detect changes that invalidate the original boundary.	Incidents, overrides, drift, subgroup findings, and workflow changes return to governance.	May reassign any task from Proceed to Escalate or Stop.	Monitoring covers safety, calibration, human behaviour, equity, and operational context.	Stable technical metrics conceal changing workflow or patient consequences.	Shared responsibility across organization, professional users, developer, and governance body.	Continuing use is not justified when meaningful deterioration cannot be detected or acted upon [16].

Application across Dispensing, Verification, Counselling, and Follow-Up

The proposed boundary model applies to components of pharmacy services rather than assigning one delegation state to an entire service. A task may move among Proceed, Escalate, and Stop as its patient, medication, informational, or workflow conditions change.

Dispensing

Dispensing combines prescription intake, data entry, product selection, preparation, labelling, exception handling, final release, and patient communication. Structured transfer and preparation functions may enter the Proceed state when inputs are complete, identity controls are reliable, exceptions are detectable, and recovery remains possible. However, community-pharmacy work is frequently interrupted, and interruptions may disrupt task continuity, checking, and recovery [17]. A normally bounded dispensing function may therefore move to Escalate when records conflict, the intended product or dose is unclear, the patient reports new information, or workflow disruption weakens independent verification. Stop applies when clinically significant ambiguity cannot be resolved or the process cannot prevent release before professional assessment.

Verification

Automation may support verification by comparing structured fields, prioritizing alerts, identifying potential discrepancies, or suppressing low-value notifications. Evidence concerning artificial intelligence for medication-alert optimization nevertheless remains heterogeneous, with limited demonstration of clinical and workflow validity across settings [18]. Alert ranking should therefore not be equated with authorization to accept or reject a medication order. Proceed may apply to bounded consistency checks with transparent rules and reliable exception capture. Escalate is appropriate when clinical significance depends on treatment intent, comorbidity, recent laboratory information, adherence, or competing risks. Stop is warranted when the system cannot represent a known exception or its output cannot be independently examined.

Counselling

Digital tools may collect questions, prepare standardized information, identify likely educational needs, or support accessible communication. A community-pharmacy intervention has shown that digital preparation can alter how patients interact with pharmacists [19]. Such evidence supports augmentation rather than autonomous replacement of counselling. Proceed may be suitable for standardized preparation or reinforcement when content is accurate, comprehensible, and consistent with the prescribed plan. Escalate is required when the patient expresses uncertainty, fear, nonadherence, adverse effects, conflicting beliefs, or difficulty understanding the information. Stop applies when communication barriers, vulnerability, or unresolved clinical questions make automated delivery insufficient.

Follow-up

Follow-up may include reminders, structured symptom collection, adherence prompts, monitoring questions, and decisions about further intervention. Pharmacist-led interventions vary substantially in content, intensity, population, and reported outcomes [20]. Routine reminders or collection of predefined information may proceed when non-response and concerning answers trigger reliable escalation. Interpretation of deterioration, adverse effects, nonadherence, or changing treatment goals should remain escalated. Automation should stop when urgent risk cannot be detected, timely human response is unavailable, or continued contact could create false reassurance.

Across verification, counselling, and follow-up, the available evidence remains task- and setting-specific rather than transferable wholesale [18–20]. These examples are conceptual applications of the model and do not establish that any complete pharmacy service is suitable for autonomous operation.

Governance, Monitoring, and Boundary Reassessment

Delegation boundaries require governance before implementation and continuing review after use begins. Early evaluation of artificial-intelligence decision support should examine intended use, workflow integration, human factors, errors, safety, and interaction with professional judgment rather than technical performance alone [21]. Evaluation of the proposed model should consequently test whether users can distinguish the three states, identify missing context, recognize unsafe automation, and escalate without excessive delay or workload.

Governance should specify the task owner, system owner, responsible professional, permitted users, escalation pathway, monitoring responsibilities, incident response, and authority to suspend use. Trustworthy healthcare AI also requires attention to fairness, universality, traceability, usability, robustness, and explainability across the system lifecycle [22]. Within community pharmacy, these principles should be translated into local questions: whether the system works for the pharmacy's patients, whether outputs and overrides are traceable, whether staff can use the tool safely, and whether vulnerable groups experience different errors or access barriers.

A Proceed decision must remain revocable. Calibration may deteriorate over time even when other measures appear comparatively stable [23]. Reassessment should therefore be triggered by changes in patient populations, medications, data sources, software, workflow, staffing, professional responsibilities, or patterns of error. Incidents, near misses, unexplained overrides, rising escalation rates, subgroup disparities, unavailable reviewers, and recovery failures should also prompt review.

Evaluation, governance, and drift monitoring must be linked rather than treated as separate lifecycle activities [21–23]. Reassessment may retain the original state, narrow the intended use, require stronger safeguards, move a task from Proceed to Escalate, or stop delegation until deficiencies are resolved. The model does not specify universal monitoring intervals or thresholds; these require prospective, task-specific justification.

Implications for Community-Pharmacy Workflow

Automation should be evaluated according to how it redistributes work, attention, interruption, and responsibility. Community-pharmacist workload includes patient interaction, clinical activity, administrative work, and operational demands [24]. Removing one manual step does not necessarily create usable clinical capacity if automation generates new alerts, exception queues, documentation, monitoring, or recovery work.

Workflow redesign should therefore identify who receives automated outputs, when review occurs, how unresolved cases are handed over, and what happens when staffing or systems are unavailable. The model also requires protected authority to pause automation. Without such authority, workload pressure may convert conditional Proceed decisions into routine acceptance.

Real-world implementation depends on infrastructure, data integration, organizational processes, staff capability, leadership, and governance [25]. Although much implementation evidence originates in hospital settings, the underlying lesson is transferable: retrospective technical performance does not demonstrate safe integration into live community-pharmacy work. Local evaluation must account for pharmacy size, staffing, prescribing interfaces, access to clinical records, patient population, service model, and escalation resources.

Community pharmacists may express positive attitudes toward artificial intelligence while also identifying needs for human supervision, training, infrastructure, and financial support [26]. Acceptance should not be interpreted as readiness, and reluctance should not automatically be interpreted as resistance to innovation. Both may reflect realistic assessments of responsibility, workload, or insufficient safeguards. The proposed model may support implementation discussions by making the reasons for proceeding, escalating, or stopping explicit and reviewable.

Limitations and Context Dependence

The model is a conceptual synthesis and has not been empirically validated. Clinical impact cannot be inferred from internal technical performance and requires representative data, external evaluation, pathway integration, prospective assessment, and surveillance [27]. The dimensions may overlap, and their relative importance may

vary by medication, patient, jurisdiction, organization, and task.

Recorded health data may contain selection, measurement, missingness, and representation biases [28]. Relevant vulnerability or contextual information may therefore be absent even when a digital record appears complete. Human review may also fail when reviewers lack time, independent evidence, or authority.

Responsibility for AI-supported decisions may be distributed across professionals, employers, developers, vendors, and system owners [29]. The model does not determine legal liability or professional scope. Generalizability, equity, and accountability are independent constraints on decision readiness [27–29].

Conclusion

Automation in community pharmacy should not be governed by a simple choice between technological adoption and human control. Medication tasks differ in consequence, ambiguity, reversibility, contextual dependence, patient vulnerability, and error detectability. Moreover, these properties vary across individual encounters and operating conditions.

The proposed Stop–Proceed–Escalate Boundary Model reframes automation as a conditional delegation decision. Proceed permits bounded execution only when context, safeguards, recovery, monitoring, and responsibility are adequate. Escalate preserves automated support while requiring professional interpretation before action. Stop prevents delegation when uncertainty, potential harm, irreversibility, poor detectability, vulnerability, or absent accountability exceeds the acceptable boundary.

The model's scientific contribution is the integration of task decomposition, human–AI collaboration, medication safety, patient context, workflow, and lifecycle governance within one proposed architecture. Its professional contribution is to distinguish meaningful oversight from nominal human presence. Its organizational contribution is to make task authority, escalation, monitoring, and suspension explicit. Its equity contribution is to treat vulnerability and unequal representation as boundary conditions rather than secondary implementation concerns.

These contributions remain conceptual. The model requires construct validation, scenario testing, workflow simulation, prospective evaluation, subgroup analysis, and continuing reassessment. It is not a scoring instrument, clinical recommendation, legal standard, regulatory determination, or deployment-ready system. Its intended use is to structure more precise questions about what automation may do, under which conditions, and when community-pharmacy responsibility must remain decisively human.

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

References

1. Hatzimanolis J, Riley B, El-Den S, Aslani P, Zhou J, Chaar BB. Applications of artificial intelligence in current pharmacy practice: A scoping review. Res Social Adm Pharm. 2025;21(3):134-41. doi:10.1016/j.sapharm.2024.12.007
2. Baines D, Nørgaard LS, Babar ZUD, Rossing C. The Fourth Industrial Revolution: Will it change pharmacy practice? Res Social Adm Pharm. 2020;16(9):1279-81. doi:10.1016/j.sapharm.2019.04.003
3. 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
4. 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
5. 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
6. Nelson SD, Poikonen J, Reese T, El Halta D, Weir C. The pharmacist and the EHR. J Am Med Inform Assoc. 2017;24(1):193-7. doi:10.1093/jamia/ocw044
7. Lyell D, Magrabi F, Raban MZ, Pont LG, Baysari MT, Day RO, et al. Automation bias in electronic prescribing. BMC Med Inform Decis Mak. 2017;17(1):28. doi:10.1186/s12911-017-0425-5
8. Mouazer A, Léguillon R, Boudegzdame N, Levrard T, Le Bars Y, Simon C, et al. ABiMed: An intelligent and visual clinical decision support system for medication reviews and polypharmacy management. BMC Med Inform Decis Mak. 2025;25(1):173. doi:10.1186/s12911-025-03002-x
9. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
10. 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
11. Jabbour S, Fouhey D, Shepard S, Valley TS, Kazerooni EA, Banovic N, et al. Measuring the impact of AI in the diagnosis of hospitalized patients: A randomized clinical vignette survey study. JAMA. 2023;330(23):2275-84. doi:10.1001/jama.2023.22295
12. 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
13. Tschandl P, Rinner C, Apalla Z, Argenziano G, Codella N, Halpern A, et al. Human-computer collaboration for skin cancer recognition. Nat Med. 2020;26(8):1229-34. doi:10.1038/s41591-020-0942-0
14. Asan O, Bayrak AE, Choudhury A. Artificial intelligence and human trust in healthcare: Focus on clinicians. J Med Internet Res. 2020;22(6). doi:10.2196/15154
15. Amann J, Blasimme A, Vayena E, Frey D, Madai VI, Precise4Q Consortium. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. doi:10.1186/s12911-020-01332-6
16. 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
17. Reddy A, Abebe E, Rivera AJ, Stone JA, Chui MA. Interruptions in community pharmacies: Frequency, sources, and mitigation strategies. Res Social Adm Pharm. 2019;15(10):1243-50. doi:10.1016/j.sapharm.2018.10.030
18. 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

19. Qudah B, Chewning B. Exploring the impact of a digital health tool on patients' interaction with community pharmacists: A pilot randomized controlled study. Res Social Adm Pharm. 2024;20(10):986-94. doi:10.1016/j.sapharm.2024.06.011
20. Milosavljevic A, Aspden T, Harrison J. Community pharmacist-led interventions and their impact on patients' medication adherence and other health outcomes: A systematic review. Int J Pharm Pract. 2018;26(5):387-97. doi:10.1111/ijpp.12462
21. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. 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
22. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
23. Davis SE, Lasko TA, Chen G, Siew ED, Matheny ME. Calibration drift in regression and machine learning models for acute kidney injury. J Am Med Inform Assoc. 2017;24(6):1052-61. doi:10.1093/jamia/ocx030
24. Gregório J, Cavaco AM, Lapão LV. How to best manage time interaction with patients? Community pharmacist workload and service provision analysis. Res Social Adm Pharm. 2017;13(1):133-47. doi:10.1016/j.sapharm.2016.02.008
25. Rahimi AK, Pienaar O, Ghadimi M, Canfell OJ, Pole JD, Shrapnel S, et al. Implementing AI in hospitals to achieve a learning health system: Systematic review of current enablers and barriers. J Med Internet Res. 2024;26. doi:10.2196/49655
26. Jarab AS, Al-Qerem W, Alzoubi KH, Obeidat H, Abu Heshmeh S, Mukattash TL, et al. Artificial intelligence in pharmacy practice: Attitude and willingness of the community pharmacists and the barriers for its implementation. Saudi Pharm J. 2023;31(8):101700. doi:10.1016/j.jsps.2023.101700
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. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. JAMA Intern Med. 2018;178(11):1544-7. doi:10.1001/jamainternmed.2018.3763
29. Price WN 2nd, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. JAMA. 2019;322(18):1765-6. doi:10.1001/jama.2019.15064