Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

When Medication Recommendations Compete for Clinical Attention: A Sociotechnical Theory of Artificial Intelligence–Supported Medication Review, Pharmacist Judgment, and Action Prioritization in Complex Polypharmacy

Ingrid Berg[1]*, Thomas Müller[2], Naledi Mokoena[3], Kenji Sato[4]

1. ***Department of Pharmacy Practice and AI-Supported Medication Review, Norwegian School of Economics, Bergen, Norway.***
2. ***Department of Clinical Pharmacy and Sociotechnical Systems, University of Vienna, Vienna, Austria.***
3. ***Department of Pharmacy Practice and Health Informatics, University of Cape Town, Cape Town, South Africa.***
4. ***Department of Clinical Pharmacy and Decision Support, Keio University, Tokyo, Japan.***

ARTICLE INFO

Received:
18 June 2026
Received in revised form:
12 October 2026
Accepted:
14 October 2026
Available online:
28 October 2026

Keywords: Artificial intelligence, Medication review, Polypharmacy, Pharmacist judgment, Clinical decision support, Professional trust

ABSTRACT

Artificial intelligence is increasingly used to identify medication-related problems, generate alerts, and support medication review, yet the clinical value of a recommendation depends on more than whether it is technically correct. In complex polypharmacy, pharmacists may confront multiple plausible recommendations while simultaneously managing incomplete information, competing tasks, verification demands, and interprofessional constraints. This article develops a sociotechnical theory of recommendation prioritization that treats clinical attention as a limited, dynamically allocated resource. The theory distinguishes recommendation generation from patient relevance, attention capture, verification, professional judgment, and eventual action. It proposes that recommendation priority emerges from interactions among expected clinical consequence, patient specificity, actionability, verification burden, temporal urgency, workflow position, and calibrated professional trust. These interactions explain why increasing recommendation volume or predictive accuracy can fail to improve medication review and may, under some conditions, displace attention from higher-value problems. Pharmacist expertise is positioned as an interpretive and coordinating capability rather than a terminal approval step. AI-supported medication review should therefore be evaluated not only by predictive performance, but also by whether it allocates attention appropriately, preserves independent verification, fits clinical workflow, and supports accountable action. The resulting account provides testable conditions under which recommendation ranking should improve, fragment, or distort medication-review work.

*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, and build upon the work non commercially.*

To Cite This Article: Berg I, Müller T, Mokoena N, Sato K. When Medication Recommendations Compete for Clinical Attention: A Sociotechnical Theory of Artificial Intelligence–Supported Medication Review, Pharmacist Judgment, and Action Prioritization in Complex Polypharmacy. Pharmacophore. 2026;17(5):130-9. https://doi.org/10.51847/805azELVHD

Introduction

Artificial intelligence is moving from peripheral automation toward medication-related decision support in pharmacy practice. Contemporary applications include identifying potential medication-related problems, supporting pharmacist interventions, screening, and workflow functions, although evidence remains heterogeneous and the relationship between technical performance and patient benefit is still incompletely established [1]. This distinction matters most in complex polypharmacy, where the practical problem is rarely the absence of possible recommendations. A single patient may generate competing concerns involving indication, dose, interactions, duplication, adherence, monitoring, deprescribing, and treatment burden. The central difficulty is deciding which of several defensible possibilities should receive attention first, which should be verified, which can wait, and which should alter care.

Medication review is itself a multicomponent clinical intervention rather than a single act of error detection. Evidence on pharmacist-led medication review identifies combinations of information gathering, evaluation, patient involvement,

Corresponding Author: Ingrid Berg; Department of Pharmacy Practice and AI-Supported Medication Review, Norwegian School of Economics, Bergen, Norway. E-mail: ingrid.berg@nhh.no.

collaboration, follow-up, and implementation, while consensus work on medication-review definitions emphasizes the need to specify purpose, setting, actors, and documentation [2, 3]. The consequence is that an AI recommendation enters an existing clinical process with its own division of labour, temporal sequence, and professional responsibilities. It does not arrive in an empty decision space. Whether it becomes useful depends on how it interacts with those arrangements.

This interaction is difficult to infer from the medication-review literature because interventions and outcomes are substantially heterogeneous. Community pharmacist-led studies vary in populations, intervention composition, and outcome measurement, limiting any assumption that a single form of review or a single endpoint captures effectiveness [4]. A theory concerned with AI-supported review therefore has to operate below the level of "medication review works" and above the level of isolated algorithmic predictions. It must explain how recommendations are transformed as they pass through human reasoning and clinical organization.

Pharmacist clinical reasoning provides one starting point. Medication review has been represented as a dynamic, multipath reasoning process in which medication information, patient characteristics, clinical knowledge, and problem formulation interact rather than being processed as a simple rule sequence [5]. This makes pharmacist judgment more than a final check after an algorithm has produced an answer. Judgment determines whether a recommendation is relevant to the particular patient, what additional evidence is needed, whether another medication problem has greater consequence, and whether a proposed action is feasible in the current setting. In this article, judgment therefore refers to situated interpretation and comparison, not discretionary rejection of otherwise complete machine outputs.

The unresolved issue is how that reasoning operates when AI expands or reorders the stream of recommendations competing for attention. Existing work can evaluate detection, alert appropriateness, acceptance, or implementation, but these are different events. A recommendation may be accurate yet clinically marginal, relevant yet costly to verify, important yet poorly timed, accepted yet not implemented, or highly ranked yet distracting from a more consequential problem. Treating all of these states as "decision support performance" obscures the mechanism through which AI affects medication-review work.

This article develops a sociotechnical theory of recommendation prioritization for complex polypharmacy. Its unit of analysis is not an isolated prediction but the recommendation-in-workflow: a medication-related proposition whose practical priority is shaped by patient relevance, anticipated consequence, verification demand, timing, competing tasks, pharmacist expertise, professional trust, and the action environment. The theory does not assume that greater recommendation accuracy or volume produces better review. Instead, it asks when recommendation ranking helps pharmacists allocate scarce attention and when the same ranking can fragment attention, increase verification burden, or induce inappropriate reliance. The first step is to specify why medication recommendations constitute an attention-prioritization problem before considering the professional and organizational mechanisms that determine action.

Medication Recommendations as an Attention-Prioritization Problem

AI systems can identify medication-related problems that overlap with those recognized by pharmacists, but agreement at the level of problem detection does not establish which recommendation deserves priority in a particular review. In one validation study of an artificial pharmacology intelligence system for polypharmacy, algorithmic outputs were compared with pharmacist judgments and expert adjudication, illustrating the possibility of machine-supported identification while also preserving the need for clinical assessment [6]. For the present theory, such systems expand the candidate set from which attention must be allocated. The relevant question is therefore not only whether a recommendation is detectable, but whether attending to it now is more valuable than attending to another plausible recommendation.

This distinction becomes clearer in the medication-alert literature. A scoping review of artificial intelligence used to optimize clinical decision-support medication alerts found approaches intended to reduce alert burden or improve targeting, but relatively little real-world implementation and no external validation among the included studies [7]. Algorithmic prioritization can consequently alter what reaches clinicians without demonstrating that the resulting ordering improves medication-review decisions across settings. Ranking is an intervention on attention before it is evidence of clinical utility.

Patient relevance provides a second filter. Clinical-pharmaceutical assessment of medication decision-support alerts has shown that content appropriateness and patient relevance are not equivalent [8]. An alert may correctly encode a pharmacological rule yet be unimportant for the individual patient because of current monitoring, treatment intent, previous clinical decisions, or contextual information unavailable to the system. In complex polypharmacy, this gap is especially consequential because numerous individually plausible recommendations can coexist. A system that ranks technically valid but context-poor recommendations too highly can consume verification time that might otherwise be directed to a less obvious but more consequential problem.

Priority is also temporally unstable. Reviews of post-implementation medication-alert optimization show that alert configurations often require modification after deployment as local patterns of use, override, and burden become visible [9]. The practical significance of a recommendation can likewise change during a medication review: new laboratory information can resolve an apparent interaction; a planned discharge can make an adherence or supply problem urgent; clarification of patient goals can reverse the relative importance of deprescribing options. Priority should therefore be understood as a provisional ordering that is updated as evidence and workflow state change, not as a fixed score attached permanently to a recommendation.

Attention capture has a cost. Alert-fatigue research demonstrates that interruptive design and poor alignment with clinical roles can reduce acceptance, while role tailoring and interaction design may mitigate some burden [10]. The relevant cost is broader

than annoyance. Every recommendation imposes some combination of reading, interpretation, data retrieval, reconciliation, and communication. Even a recommendation that is quickly dismissed occupies a portion of cognitive and temporal capacity. When several recommendations compete, the system influences care not merely by what it says, but by determining which claims must be processed before other work can proceed.

Pharmacy practice is already interruption-rich. A recent scoping review found that interruptions and distractions in pharmacy settings were associated with workload, performance, waiting time, and practitioner well-being, although effects and tested interventions varied across studies [11]. AI recommendations become one more claimant on this contested attention environment. This makes recommendation volume clinically relevant even when each recommendation has acceptable standalone accuracy. If low-consequence recommendations repeatedly demand verification, they can displace attention from high-consequence medication problems, patient communication, or coordination work. Attention displacement is therefore a plausible failure mode of otherwise competent decision support.

Actionability modifies this competition. In a simulation trial of ABiMed, pharmacists using visual clinical decision support identified more relevant drug-related problems without an increase in medication-review time [12]. The finding does not establish real-world outcome benefit, but it shows why actionability should be separated from simple alert generation: decision support can change what pharmacists can recognize within a bounded review episode. A recommendation becomes more actionable when the information required to assess it is accessible, the implied action lies within or near the pharmacist's sphere of influence, the reason for prioritization is intelligible, and the timing allows a response. Conversely, a clinically important recommendation may remain low in immediate action priority if essential data are unavailable or if action depends on coordination that cannot occur at that point.

Taken together, these findings support a distinction among at least five properties of a medication recommendation: expected clinical consequence, patient relevance, temporal urgency, verification demand, and actionability. These properties can diverge. A highly consequential recommendation may have high verification demand; an urgent alert may have low patient relevance; a readily actionable recommendation may concern a relatively minor optimization; and a recommendation with strong pharmacological logic may depend on contextual information outside the AI system. The theory therefore defines recommendation priority as a provisional allocation of clinical attention produced by comparing these attributes under current workload and workflow constraints. It is not equivalent to severity, confidence, predicted risk, or algorithmic rank.

The resulting competition can be described as an attention field rather than a queue. Recommendations may enter simultaneously, gain or lose salience as evidence is checked, and compete with non-AI tasks such as medication reconciliation, patient counselling, and communication with prescribers. Some recommendations require immediate interruption; others can be batched; still others should remain visible without demanding current action. The key design problem is not to maximize the number of recommendations noticed. It is to preserve attention for the combination of medication problems whose verification and action are most consequential in the current clinical state.

This distinction also changes how recommendation quality should be evaluated. Precision-oriented measures ask whether the system selects correct or relevant items. An attention-oriented evaluation additionally asks what work those items create, which other problems are delayed, whether the ranking remains appropriate after contextual information is added, and whether pharmacists can recover from an incorrect or poorly timed priority assignment. These are sociotechnical properties because their value depends on relations among system output, pharmacist reasoning, patient state, and workflow capacity. They cannot be inferred from model discrimination alone.

The recommendation attributes developed above provide a practical basis for distinguishing attention capture from clinical priority. **Table 1** organizes these attributes by the kind of verification work they generate and by the conditions under which they can raise, lower, or defer action priority.

Table 1. Recommendation attributes that transform an AI-identified medication issue into a competing claim on clinical attention

Recommendation attribute	Question posed during pharmacist review	Typical verification work	How the attribute can raise priority	How the attribute can lower or defer priority	Failure if ranking ignores the attribute
Expected clinical consequence	What could happen if the medication issue is not addressed?	Establish severity, reversibility, and plausible patient consequence	Potential for substantial or time-sensitive harm	Consequence is minor, remote, or already controlled	Minor optimizations displace consequential problems, or severe problems are buried
Patient relevance	Does the recommendation fit the current patient, treatment intent, and available clinical information?	Reconcile diagnoses, laboratory values, prior decisions, goals, and monitoring	Strong patient-specific concordance	Rule is technically correct but contextually inapplicable	Generic pharmacological correctness is mistaken for clinical importance
Temporal urgency	When would delayed verification or action	Determine timing of exposure, monitoring,	Delay could remove an opportunity to	Issue can safely be reviewed later or batched	Interruptions are triggered by non-urgent work while

	materially change the decision?	transitions, and planned care	prevent harm or complete care		urgent problems wait
Verification demand	What evidence must be checked before the recommendation can be trusted or rejected?	Retrieve data, reconcile records, inspect trends, and resolve conflicting information	Low verification burden when consequence and relevance are already clear	High burden with weak expected informational gain	Attention is consumed by difficult-to-resolve low-value recommendations
Actionability	Is there a feasible action pathway if the recommendation is confirmed?	Identify responsible actor, available intervention, and coordination requirements	Clear, feasible, and authorized response	Necessary action is unavailable, blocked, or dependent on missing actors or data	Systems repeatedly surface issues that cannot progress to care
Dependency on other recommendations	Does acting on this issue alter the relevance of another medication problem?	Compare interacting recommendations and anticipated downstream effects	Recommendation resolves or conditions several related problems	Action should await resolution of a more upstream issue	Independent ranking fragments clinically linked medication decisions
Reassessment sensitivity	What new information would change this recommendation's position?	Identify pending tests, consultations, patient preferences, or treatment changes	Stable priority despite plausible updates	Priority is likely to change when imminent information arrives	Static ranks persist after their clinical basis has changed

The table treats these attributes as analytic dimensions rather than a numerical scoring system. No universal weighting is proposed. Their importance depends on the patient and on the review environment, and some may conflict. The implication is that AI-supported prioritization should preserve the pharmacist's ability to reinterpret relationships among recommendations rather than compressing all medication problems into an opaque ordinal list.

The same argument can be represented as a dynamic field in which recommendations compete for attention while their priority changes with verification and context. **Figure 1** depicts this proposed synthesis. Evidence-derived elements are the existence of alert burden, patient-relevance filtering, workflow interruption, iterative optimization, and decision-support effects on problem identification; the relationships among these elements are integrated here as a proposed sociotechnical account rather than as a validated causal model.

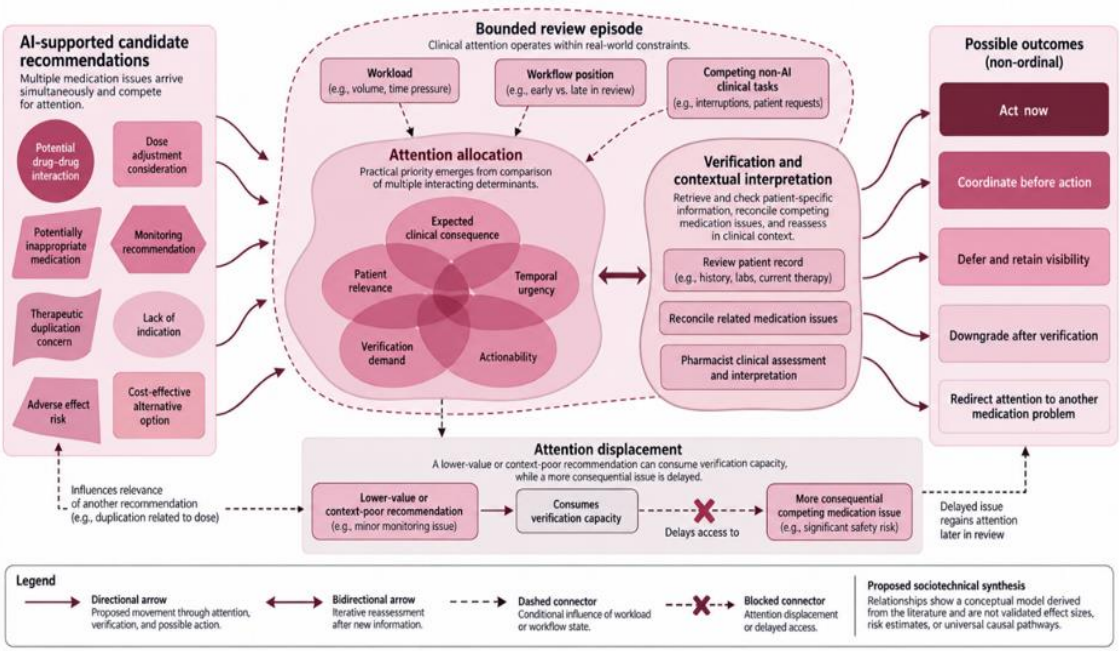


Figure 1. How competing medication recommendations acquire, lose, and redirect clinical attention during complex polypharmacy review

Pharmacist Judgment, Workflow and Professional Trust

A recommendation does not become clinically important merely because it appears high in a ranked list. It must be interpreted by a professional who decides whether the underlying medication problem is credible, consequential, sufficiently resolved to support action, and compatible with the patient's current care plan. Evidence from older adults receiving pharmacist

recommendations shows that prescriber acceptance varies substantially across studies and is influenced by contextual and collaborative factors [13]. This variation is important theoretically because it separates the informational quality of a recommendation from its capacity to mobilize action. A recommendation can be pharmacologically defensible while still failing to progress if its rationale is poorly aligned with the receiving clinician's priorities, arrives at an unsuitable time, or depends on information that has not yet been verified.

Acceptance itself should not be treated as the endpoint. Deprescribing work in frail older adults demonstrates that recommendations can pass through distinguishable stages of proposal, acceptance, and implementation [14]. For AI-supported medication review, this means that a system evaluated solely on recommendation acceptance may overstate its practical contribution. An accepted recommendation may remain unimplemented because the relevant prescriber is unavailable, the patient declines the change, monitoring cannot be arranged, competing treatment priorities intervene, or the care setting lacks authority to execute the proposed action. Conversely, a recommendation initially rejected may later become appropriate after new information emerges. The pathway from algorithmic output to medication change is therefore conditional and reversible.

The broader pharmaceutical-care literature reinforces this distinction. Outcomes associated with pharmacist-led services vary across clinical, humanistic, and economic domains [15]. A recommendation that is useful for medication safety may create treatment burden; a change that improves guideline concordance may be inconsistent with patient preference; and an intervention that appears efficient at the level of recommendation processing may generate downstream coordination work. Pharmacist judgment must therefore compare not only medication-related risks but also the consequences of acting, not acting, or waiting. In complex polypharmacy, prioritization is partly the management of these competing consequences.

Relational context matters because medication review is conducted through interactions among patients, pharmacists, prescribers, nurses, caregivers, and other professionals. Qualitative work on patient-centred medication review indicates that intervention effects depend on mechanisms activated within particular relationships and contexts [16]. The pharmacist's interpretation of an AI recommendation may consequently be altered by knowledge of the patient's goals, previous discussions with the prescriber, anticipated willingness to change therapy, or the reliability of available records. These forms of contextual knowledge are not necessarily represented in the algorithmic input. Their absence does not make the recommendation wrong, but it changes what can legitimately be inferred from it.

Workflow integration further shapes whether expertise can be exercised. Realist evidence concerning pharmacist participation in interprofessional ward rounds shows that professional contribution depends on how pharmacists are integrated into team processes and on the contexts and mechanisms that enable participation [17]. The implication for AI-supported review is that pharmacist expertise cannot compensate indefinitely for a system that delivers recommendations outside decision windows, duplicates work performed elsewhere, or separates verification from the clinical interaction in which action becomes possible. Workflow fit is therefore not a convenience added after model development. It determines whether the information can enter the decision process at a useful point.

Polypharmacy management also resists a single standardized response. Comparative evidence across medication-management interventions in older people with chronic disease demonstrates heterogeneity in outcomes and intervention performance [18]. Accordingly, an AI system that consistently gives one category of medication problem high priority may function well in one setting and poorly in another because the distribution of preventable harm, patient goals, available services, and professional responsibilities differs. Prioritization should be expected to vary across institutional, population, and temporal contexts.

Professional authority defines a further boundary. Systematic evidence on pharmacist prescribing shows that the structure of pharmacist authority and access to medicines can influence what pharmacists are able to accomplish [19]. A recommendation that can be independently enacted in one jurisdiction or service may require several handoffs in another. This alters actionability even when the pharmacological issue is identical. The same algorithmic recommendation may therefore deserve different operational priority depending on whether confirmation leads directly to action or initiates a costly coordination sequence.

These relationships suggest that professional trust must be treated as calibrated reliance, not simple willingness to accept AI output. Pharmacists require enough trust to allocate attention to system recommendations, but also enough independence to challenge them when patient context, competing evidence, or workflow realities indicate that the recommendation should be modified or rejected. Under-trust can cause useful recommendations to be ignored; over-trust can allow high-ranked outputs to displace independent verification. The relevant question is thus whether trust is proportionate to the quality of evidence and the consequences of relying on the system in that situation.

The interaction among judgment, workflow, trust, and action can be summarized as a set of sociotechnical conditions rather than a linear adoption sequence. **Table 2** distinguishes configurations in which an AI recommendation is likely to progress, stall, be redirected, or require additional verification.

Table 2. Sociotechnical conditions that transform ranked AI medication recommendations into pharmacist action, modification, deferral, or rejection

Sociotechnical condition	Pharmacist interpretive task	Workflow requirement	Trust requirement	Likely recommendation trajectory	Characteristic failure if the condition is absent

Patient-context concordance	Determine whether the recommendation fits goals, diagnoses, monitoring, and treatment intent	Relevant clinical information available at review	Reliance limited to information the system can legitimately represent	Progresses if context confirms relevance; otherwise modified or downgraded	Context-poor output is treated as patient-specific judgment
Resolvable verification demand	Test the recommendation against records and pharmacological knowledge	Required data can be retrieved without disproportionate delay	Confidence remains provisional until verification is complete	Progresses after confirmation or is redirected after contradiction	Verification cost consumes attention without improving the decision
Workflow opportunity	Determine whether action is possible at the current decision point	Recommendation arrives before the relevant prescribing, discharge, monitoring, or counselling window closes	System is trusted to surface issues at usable moments	Act now, coordinate, or deliberately defer	Correct recommendations arrive after the opportunity for action
Professional authority and role fit	Identify who can authorize, communicate, and implement change	Responsibility is clear and handoffs are feasible	Trust does not substitute for accountability	Progresses through direct action or explicit coordination	Recommendation remains informational because no actor owns the next step
Interprofessional credibility	Translate the rationale into a form acceptable to other decision makers	Pharmacist participation is integrated into team communication	Trust is distributed across pharmacist, system, and receiving clinician	Accepted, negotiated, or revised	System ranking is mistaken for interprofessional agreement
Calibrated reliance	Decide how much independent checking is warranted	Time and access permit challenge when stakes are high	Reliance rises or falls with evidence quality and context	Efficient verification with preserved professional oversight	Under-trust causes neglect; over-trust causes uncritical acceptance
Recommendation interdependence	Determine whether one proposed action changes another medication problem	Related recommendations can be reviewed together	Trust applies to relationships among recommendations, not isolated outputs	Recommendations are sequenced, bundled, or reprioritized	Independent ranking fragments clinically connected decisions
Reassessment capacity	Revisit the recommendation when patient or treatment information changes	System and workflow permit deferred issues to re-enter attention	Prior decisions are not treated as permanently settled	Deferred recommendation returns when conditions change	Stale priorities persist after their rationale has expired

The table makes two distinctions central to the proposed theory. First, recommendation priority is produced jointly by the informational properties of the recommendation and by the conditions under which a pharmacist can evaluate and act on it. Second, rejection, deferral, or modification should not automatically be interpreted as evidence that the decision-support system failed. These trajectories may represent appropriate professional responses to context that the system could not observe.

A Sociotechnical Theory of Recommendation Prioritization

Trust in AI-supported clinical decision making is shaped by characteristics of the technology, the user, the evidence supporting the system, and the organizational environment [20]. In medication review, this means that professional trust should be linked to the specific task being delegated. A pharmacist may reasonably trust a system to identify a potential interaction while withholding trust in its assessment of whether that interaction is clinically important for a particular patient. Trust can therefore vary across components of the same recommendation.

Empirical and normative work on trustworthy clinical AI further indicates that accuracy alone is insufficient. Validity, relevance of the underlying data, appropriateness of the clinical context, and the relationship between users and the system contribute to judgments about whether reliance is justified [21]. These considerations support a central proposition of the theory developed here: the priority assigned to an AI recommendation should be proportional not only to its predicted clinical importance but also to the confidence with which its patient-specific meaning can be established within the available workflow. Explainability modifies this relationship but does not solve it. Systematic evidence indicates that explanations can increase, decrease, or leave clinician trust unchanged, and experimental evidence shows that additional explanation does not necessarily improve diagnostic performance [22, 23]. Explanation therefore has an attentional cost as well as a potential verification

benefit. A detailed rationale may facilitate checking when uncertainty is clinically important, but excessive explanation can enlarge the amount of material that must be processed before action. The appropriate amount of explanation should depend on what uncertainty remains, the consequences of error, and whether the explanation changes the pharmacist's ability to verify the recommendation.

Automation bias represents the corresponding risk of insufficiently critical reliance. Empirical work demonstrates that exposure to erroneous AI advice can induce agreement with incorrect recommendations under some conditions [24]. Applied to medication review, this suggests that high algorithmic rank may acquire persuasive force independent of patient-specific validity. The danger is not simply that a pharmacist "trusts AI too much." The deeper problem is that ranking itself directs verification effort. Once a recommendation is given privileged attention, competing medication problems may receive less scrutiny.

Clinical pharmacology provides an important corrective. Reviews of AI in clinical pharmacology emphasize the continued need for domain expertise, suitable validation, and integration with the clinical decision context [25]. Medication recommendations often depend on mechanistic knowledge, dose–exposure relationships, organ function, indication, timing, formulation, monitoring, and treatment goals. AI can reorganize evidence within this space, but it cannot make these contextual requirements disappear. Recommendation prioritization is therefore best understood as allocation of analytical work between system and professional rather than transfer of the entire decision to the model.

Evidence from controlled human–AI medication-safety scenarios also indicates that combined performance can differ from that of either the pharmacist or AI operating alone [26]. The importance of such evidence for the present theory is not that combination is invariably superior. It is that the configuration of work matters. Human–AI complementarity should occur when the system directs attention toward problems that benefit from pharmacist interpretation while preserving enough independence for contextual correction. Complementarity fails when the system adds redundant verification work, suppresses alternative hypotheses, or directs attention toward recommendations that cannot be acted upon.

The theory therefore proposes four interacting operations. Selection determines which candidate medication problems enter active attention. Verification tests whether those candidates remain relevant when patient-specific and pharmacological context is added. Coordination determines whether a validated recommendation can be translated across professional roles and workflow boundaries. Reassessment allows recommendations to rise, fall, or return as information and circumstances change. These operations are not fixed stages. A recommendation can move backward from coordination to verification, bypass extensive verification when urgent danger is already clear, or remain deferred until new information alters its relevance.

The strongest prediction generated by this account is that improving recommendation accuracy will have diminishing or even adverse practical returns when the additional recommendations impose greater verification and coordination costs than the attention they save. Conversely, a system with modest predictive improvement may yield substantial benefit if it suppresses low-value interruptions, bundles interdependent recommendations, presents uncertainty at the point of verification, and returns deferred issues when action becomes feasible. This prediction distinguishes the theory from a simple better-ranking model.

A second prediction concerns errors. The clinical effect of an AI error should depend on where in the attention architecture the error occurs. A false-positive recommendation with low attention capture may be inexpensive; the same false positive placed at the top of an interruptive ranking can displace consequential work. A false negative may be partly recoverable when pharmacists conduct independent medication review but more dangerous when workflow evolves toward exclusive dependence on AI triage. Safety evaluation should therefore examine the interaction between model error and attention allocation rather than counting errors without reference to workflow.

Figure 2 integrates these relationships as a proposed sociotechnical architecture. It depicts AI output, pharmacist interpretation, workflow conditions, calibrated trust, coordination, and clinical action as interacting components with explicit feedback rather than as a linear automation pathway.

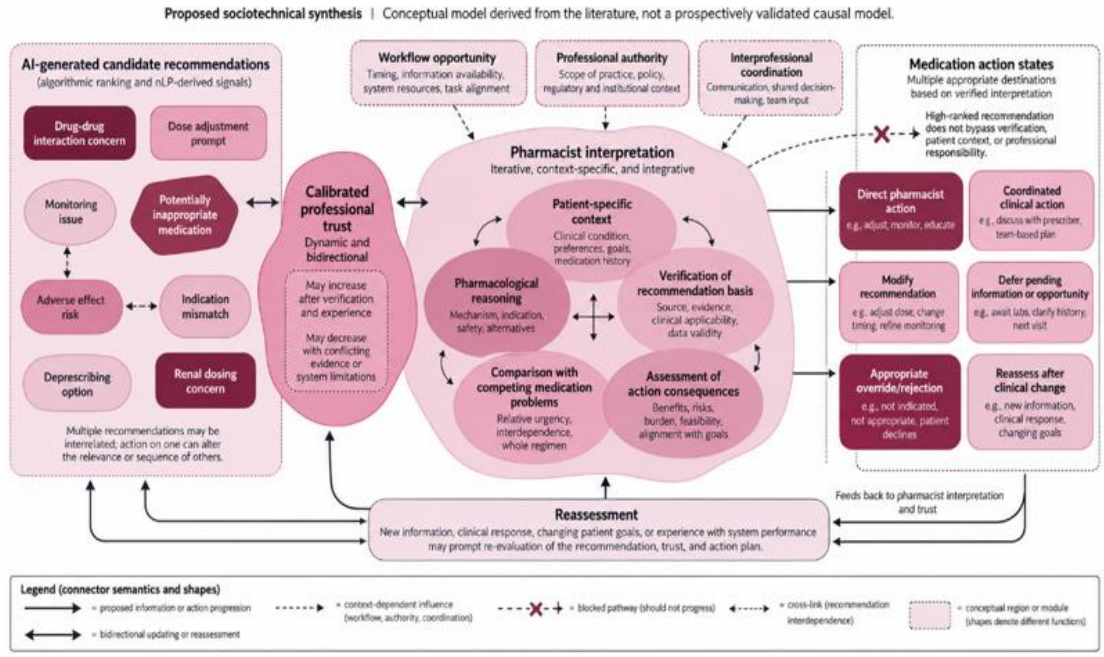


Figure 2. When AI-ranked medication recommendations become actionable, remain provisional, or are appropriately overridden

Implications for AI-Supported Medication Review

The theory changes what should be measured when AI-supported medication review is evaluated. Pharmacist prescribing research demonstrates that professional roles and authority structures influence health-related outcomes and the capacity to implement medication changes [27]. Evaluation should therefore follow recommendations beyond generation and ask whether relevant issues reach an actor who can verify and act on them. Useful endpoints include progression from recommendation to verification, reasons for deferral or override, time spent resolving recommendations, delayed competing work, and completion of clinically appropriate action. None of these measures should be interpreted in isolation as proof of patient benefit, but together they reveal whether the system improves the work of medication review.

Implementation also has an organizational cost. Economic analyses of pharmacist services show that value depends on service configuration, setting, and assumptions about resources and outcomes [28]. AI systems likewise consume pharmacist time through review, verification, troubleshooting, documentation, and communication. A system that reduces nominal alert volume could still increase workload if remaining recommendations require extensive reconciliation. Conversely, a tool that introduces additional visible recommendations could be useful if it substantially reduces search effort or coordinates related problems. Resource evaluation should therefore measure verification work displaced or created, not only the number of recommendations produced.

Context-specificity is equally important. Specialist pharmacist-prescribing literature demonstrates that roles and responsibilities differ substantially across service environments [29]. An AI prioritization strategy developed in one context should not be transferred unchanged to another simply because the medication knowledge base is similar. Local authority, information access, staffing, patient characteristics, monitoring capacity, and communication channels can all change which recommendation is actionable. Deployment should consequently include a workflow-specific account of who receives each recommendation, what information is needed to verify it, and what happens after acceptance.

Pharmacists retain a particular responsibility within this architecture because medication-quality work depends on expertise embedded within broader clinical systems [30]. AI should therefore support professional accountability rather than diffuse it. Systems should make it possible to identify why a recommendation was prioritized, what relevant information remains unknown, and what clinical assumption would cause its priority to change. They should also make deferral and override visible as reasoned states rather than treating non-acceptance as a simple failure of adherence to the system.

A practical design implication follows: recommendation interfaces should distinguish why an issue deserves attention from what action should follow. A high-consequence issue with missing contextual data should remain prominent while explicitly signalling verification need. An actionable but low-consequence optimization can be retained without demanding immediate interruption. Interdependent recommendations should be presented together when acting on one changes the meaning of another. Deferred recommendations should have defined re-entry conditions so that postponement does not become disappearance.

The theory also suggests a research agenda that compares systems with similar predictive performance but different attentional architectures. Such studies could examine whether prioritization reduces unnecessary verification, whether pharmacists detect important problems outside the AI-ranked set, how frequently priorities change after patient-specific information is added, and whether explanation length alters verification quality. Studies should also distinguish appropriate override from automation

resistance and appropriate acceptance from overreliance. These distinctions are particularly important because recommendation acceptance, team integration, and trust are each context-sensitive phenomena rather than universal indicators of success [13, 17, 21].

The most informative evaluations will therefore treat medication review as joint human–technical work. Model discrimination and calibration remain necessary, but they address only whether a system can identify or rank candidate problems. Clinical usefulness additionally requires evidence that the resulting distribution of attention supports pharmacist reasoning, preserves recovery from algorithmic error, fits workflow opportunities, and enables accountable action. The decisive question is not whether the AI generates the "right" recommendation in isolation, but whether its participation changes the overall medication-review process in a clinically defensible direction.

Conclusion

AI-supported medication review creates value only when useful recommendations acquire appropriate attention without displacing more consequential clinical work. In complex polypharmacy, priority is not an intrinsic property of an algorithmic recommendation. It emerges from the interaction of expected clinical consequence, patient relevance, temporal urgency, verification demand, actionability, pharmacist expertise, workflow opportunity, professional authority, and calibrated trust.

This sociotechnical account explains why more accurate or more numerous recommendations do not necessarily produce better medication review. Recommendations must survive patient-specific verification, compete successfully but appropriately for attention, enter a workable coordination pathway, and remain open to reassessment. Pharmacist override, modification, or deferral can therefore represent competent clinical reasoning rather than failure to follow decision support.

The central implication is conditional: AI improves medication review when it helps allocate scarce professional attention to medication problems whose verification and action are most consequential in the current clinical context while preserving independent judgment and recovery from error. When ranking increases verification burden, fragments interdependent decisions, arrives outside actionable workflow windows, or encourages uncritical reliance, the same technology can weaken the review process it is intended to support.

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. Craske ME, Hardeman W, Steel N, Twigg MJ. Components of pharmacist-led medication reviews and their relationship to outcomes: A systematic review and narrative synthesis. BMJ Qual Saf. 2024;33(12):808-22. doi:10.1136/bmjqs-2024-017283
3. Lias N, Lindholm T, Holmström AR, Uusitalo M, Kvarnström K, Toivo T, et al. Harmonizing the definition of medication reviews for their collaborative implementation and documentation in electronic patient records: A Delphi consensus study. Res Social Adm Pharm. 2024;20(6):52-64. doi:10.1016/j.sapharm.2024.01.016
4. Vogt CJ, Moecker R, Jacke CO, Haefeli WE, Seidling HM. Exploring the heterogeneity in community pharmacist-led medication review studies – A systematic review. Res Social Adm Pharm. 2024;20(8):679-88. doi:10.1016/j.sapharm.2024.03.012
5. Guignard B, Crevier F, Charlin B, Audétat MC. A graphical model to make explicit pharmacist clinical reasoning during medication review. Res Social Adm Pharm. 2024;20(12 Pt A):1142-50. doi:10.1016/j.sapharm.2024.09.005
6. Dil-Nahlieli D, Ben-Yehuda A, Souroujon D, Hyam E, Shafran-Tikvah S. Validation of a novel Artificial Pharmacology Intelligence (API) system for the management of patients with polypharmacy. Res Social Adm Pharm. 2024;20(7):633-9. doi:10.1016/j.sapharm.2024.04.003
7. 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
8. Bauer J, Busse M, Koch S, Schmid M, Sommer J, Fromm MF, et al. Clinical-pharmaceutical assessment of medication CDSS alerts: Content appropriateness and patient relevance in clinical practice. Front Pharmacol. 2025;16:1510425. doi:10.3389/fphar.2025.1510425

9. Ledger TS, Brooke-Cowden K, Coiera E. Post-implementation optimization of medication alerts in hospital computerized provider order entry systems: A scoping review. J Am Med Inform Assoc. 2023;30(12):2064-71. doi:10.1093/jamia/ocad193
10. 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
11. Ayanaw M, Lim A, Khera H, Vu T, Goordeen D, Malone D. How do interruptions and distractions affect pharmacy practice? A scoping review of their impact and interventions in dispensing. Res Social Adm Pharm. 2025;21(9):667-78. doi:10.1016/j.sapharm.2025.05.001
12. Mouazer A, Dubois S, Léguillon R, Boudegzdame N, Levrard T, Le Bars Y, et al. A simulation trial evaluating ABiMed, a clinical decision support system for medication reviews and polypharmacy management. Res Social Adm Pharm. 2025;21(11):904-13. doi:10.1016/j.sapharm.2025.06.006
13. Ramsey NC, Peterson GM, Mirkazemi C, Salahudeen MS. Factors influencing medical prescribers' acceptance of pharmacists' recommendations in non-hospitalized older adults: A systematic review and meta-analysis. J Am Med Dir Assoc. 2025;26(3):105462. doi:10.1016/j.jamda.2024.105462
14. Hurley E, Dalton K, Byrne S, Foley T, Walsh E. Pharmacist-led deprescribing using STOPPFrail for frail older adults in nursing homes. J Am Med Dir Assoc. 2024;25(9):105122. doi:10.1016/j.jamda.2024.105122
15. Fares R, Marjorie B, Chaballe C, Crunenberg R. The outcomes of pharmacist-led pharmaceutical care within community pharmacies: An overview of systematic reviews. Res Social Adm Pharm. 2025;21(5):332-9. doi:10.1016/j.sapharm.2025.01.015
16. Kaae S, Holst SS, Hansen JM, Vermehren C. New program theory in pharmacist-led patient-centered medication review in general practice: A qualitative pilot study. Res Social Adm Pharm. 2025;21(12):1070-8. doi:10.1016/j.sapharm.2025.07.002
17. Babu D, Luetsch K, Kalisch Ellett L, Harmon J, Marotti S, Wisdom A, et al. A realist review of pharmacists' integration and participation in interprofessional ward rounds: What works, for whom, why and in what circumstances. Res Social Adm Pharm. 2025;21(11):831-41. doi:10.1016/j.sapharm.2025.06.102
18. Wang Q, Zhang J, Li K, Miao X, Wang Y, Zhang L. Effectiveness of different medication management measures in older patients with chronic diseases and polypharmacy: A systematic review and network meta-analysis. Res Social Adm Pharm. 2025;21(10):753-64. doi:10.1016/j.sapharm.2025.05.011
19. Walpola RL, Issakhany D, Gisev N, Hopkins RE. The accessibility of pharmacist prescribing and impacts on medicines access: A systematic review. Res Social Adm Pharm. 2024;20(5):475-86. doi:10.1016/j.sapharm.2024.01.006
20. Rojas JC, Teran M, Umscheid CA. Clinician trust in artificial intelligence: What is known and how trust can be facilitated. Crit Care Clin. 2023;39(4):769-82. doi:10.1016/j.ccc.2023.02.004
21. Kostick-Quenet K, Lang BH, Smith J, Hurley M, Blumenthal-Barby J. Trust criteria for artificial intelligence in health: Normative and epistemic considerations. J Med Ethics. 2024;50(8):544-51. doi:10.1136/jme-2023-109338
22. Rosenbacke R, Melhus Å, McKee M, Stuckler D. How explainable artificial intelligence can increase or decrease clinicians' trust in AI applications in health care: Systematic review. JMIR AI. 2024;3. doi:10.2196/53207
23. Rezaeian O, Asan O, Bayrak AE. The impact of AI explanations on clinicians' trust and diagnostic accuracy in breast cancer. Appl Ergon. 2025;129:104577. doi:10.1016/j.apergo.2025.104577
24. Kücking F, Hübner U, Przysucha M, Hannemann N, Kutza JO, Moelleken M, et al. Automation bias in AI-decision support: Results from an empirical study. In: Röhrig R, Grabe N, Hübner UH, Jung K, Sax U, Schmidt CO, et al., editors. German medical data sciences 2024: Health - thinking, researching and acting together. Proceedings of the 69th Annual Meeting of the German Association of Medical Informatics, Biometry, and Epidemiology; 2024 Sep 8-11; Dresden, Germany. Amsterdam: IOS Press; 2024. p. 298-304. (Studies in Health Technology and Informatics; vol. 317). doi:10.3233/SHTI240871
25. Shahin MH, Barth A, Podichetty JT, Liu Q, Goyal N, Jin JY, et al. Artificial intelligence: From buzzword to useful tool in clinical pharmacology. Clin Pharmacol Ther. 2024;115(4):698-709. doi:10.1002/cpt.3083
26. Ong JCL, Jin L, Elangovan K, Lim GYS, Lim DYZ, Sng GGR, et al. Large language model as clinical decision support system augments medication safety in 16 clinical specialties. Cell Rep Med. 2025;6(10):102323. doi:10.1016/j.xcrm.2025.102323
27. Akintunde AO, Karaxhija A, Laaksonen R. Effects of pharmacist prescribing on health-related outcomes in secondary care compared with medical prescribing or no treatment: A systematic review. Res Social Adm Pharm. 2026;22(2):254-62. doi:10.1016/j.sapharm.2025.11.003
28. Silva Almodovar A, Blankenship B, Miller J, Trombetta M, Murphy EM. Return on investment of pharmacists' services among non-hospitalized patients: A scoping review. Res Social Adm Pharm. 2025;21(5):321-31. doi:10.1016/j.sapharm.2025.01.012
29. Johnson L, Khattab S, Strawbridge J, Cadogan C, Stewart D, De Frein AM, et al. Pharmacist prescribing in cancer services: A scoping review. Res Social Adm Pharm. 2025;21(12):951-74. doi:10.1016/j.sapharm.2025.05.004
30. Bajorek B. Ensuring the quality use of medicines in clinical trials: A review and perspective on optimising the role of pharmacists. Res Social Adm Pharm. 2026;22(3):363-9. doi:10.1016/j.sapharm.2025.11.004