Original Article

From Medication Alerts to Medication Sensemaking in Hospital Pharmacy

João Silva[1]*, Pedro Costa[1], Ana Beatriz[2]

[1]Department of Hospital Pharmacy and Clinical Alerts, Faculty of Pharmacy, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil.
[2]Department of Medication Sensemaking and Informatics, Faculty of Pharmacy, University of Campinas, Campinas, Brazil.

Abstract

Medication-related clinical decision support is commonly designed to detect potentially hazardous prescriptions, interactions, doses, monitoring gaps, or patient–drug mismatches. Detection, however, does not establish that an alert is valid, clinically relevant, sufficiently contextualized, understood by the appropriate professional, or translated into a defensible medication decision. This article develops an original Medication-Sensemaking Model for hospital pharmacy. The model distinguishes signal production from the collaborative construction of meaning and positions alerts as inputs to professional reasoning rather than completed decisions. It proposes a sequence comprising signal-validity assessment, patient–drug–workflow context assembly, relevance adjudication, uncertainty appraisal, rationale formation, collaborative interpretation, action or monitored deferral, escalation, and governance feedback. Human and digital contributions are treated as interdependent but non-equivalent: computational systems may identify patterns and retrieve information, whereas professional authority, accountability, contextual judgment, and responsibility for escalation remain organizationally assigned. Evaluation should therefore extend beyond alert frequency, acceptance, and override rates to include contextual completeness, rationale quality, medication-task performance, workflow consequences, safety, equity, escalation reliability, generalizability, and lifecycle governance. The model is an evidence-informed conceptual synthesis rather than a validated clinical pathway. Its components, transition rules, and escalation boundaries require prospective testing across alert classes, professional configurations, patient populations, and hospital infrastructures. The principal contribution is a structured account of how hospital pharmacy may move from detecting medication signals toward constructing reviewable, uncertainty-aware, and collaboratively actionable medication meaning without equating automation with professional judgment or technical performance with clinical benefit.

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

Introduction

Medication alerts are intended to identify situations in which a prescription, administration plan, laboratory result, interaction, or patient characteristic may require attention. Their presence can improve the visibility of possible medication problems, but visibility is not equivalent to clinical relevance. An alert may be technically correct yet unimportant for the present patient, clinically important but poorly timed, or actionable only after information unavailable to the alerting system has been considered. The unresolved issue is therefore not simply how to generate more signals, but how those signals become defensible medication decisions within hospital work.

Alert burden is shaped by workload and repeated exposure, the tradeoff between hazardous-order detection and nuisance alerting, and the way alerts are presented and targeted to professional roles [1–3]. These findings challenge a linear assumption that adding alerts necessarily increases safety. More sensitive systems may identify additional potential hazards while also increasing interruption, repetition, and

Address for correspondence: João Silva, Department of Hospital Pharmacy and Clinical Alerts, Faculty of Pharmacy, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil.
joao.silva@ufrj.br

How to cite this article: Silva J, Costa P, Beatriz A. From Medication Alerts to Medication Sensemaking in Hospital Pharmacy. Arch Pharm Pract. 2025;16(1):60-8. https://doi.org/10.51847/9RIvyEeOnR

low-value review. Conversely, aggressive suppression may reduce burden while concealing uncommon but consequential problems. Neither alert volume nor acceptance rate alone resolves this tradeoff.

Hospitals have adopted multiple strategies to modify alert content, thresholds, interfaces, and delivery, yet optimization of alerting does not by itself demonstrate improved interpretation or safer medication decisions [4]. An alert may be accepted because it is easy to follow, dismissed because it lacks context, or overridden because a clinician possesses information that the system does not represent. Similar response patterns may therefore arise from very different combinations of alert quality, workflow pressure, patient complexity, professional expertise, and organizational norms.

This article proposes that hospital-pharmacy systems should distinguish detection from sensemaking. Detection refers to the identification and presentation of a possible medication-related discrepancy or risk. Sensemaking refers to the construction of an actionable interpretation by connecting the signal to patient state, drug and regimen characteristics, treatment goals, temporal information, workflow conditions, uncertainty, and professional responsibilities. This distinction does not diminish the value of alerts. It specifies that an alert is an input into medication work rather than evidence that the work has been completed.

The proposed contribution is an original clinical sensemaking model rather than an empirical effectiveness claim. It organizes the transitions through which an alert may be validated, contextualized, interpreted, challenged, documented, escalated, and returned to system governance. Its purpose is to make those transitions visible enough to be designed and evaluated without assuming that the same configuration will suit every medication task, hospital, patient population, or professional arrangement.

The Difference between Detection and Sensemaking

Clinical sensemaking involves developing a sufficiently shared and actionable understanding from information that may be incomplete, distributed, or differently interpreted across members of a care team [5]. Applied to medication alerts, this means that the central question is not only whether a signal appeared, but what the signal means for this patient, at this time, within the current treatment plan and organizational setting.

Detection is comparatively bounded. A rule, model, database, or monitoring process identifies a predefined pattern and presents it to a user. Sensemaking begins when the user assesses whether the underlying data are current, whether the rule applies, what contextual information is missing, and how the signal relates to treatment goals and competing risks. Context-sensitive decision support can improve access to information relevant to the user, patient, task, and point in the workflow [6]. Information access nevertheless remains distinct from interpretation: retrieving relevant material does not determine how conflicting evidence should be weighed or who should act.

A displayed alert may also be invalid or misleading because of failures in data, knowledge content, logic, configuration, presentation, workflow integration, or maintenance [7]. Detection should therefore not be treated as a neutral starting point. Before clinical meaning is constructed, the signal itself may require scrutiny. This introduces a signal-validity question that conventional alert-response measures often leave implicit.

Alert exposure, attention, acceptance, and action are also separable phenomena [8]. A user can see an alert without processing it, process it without agreeing, agree without acting immediately, or override it while taking an alternative risk-reducing action. Binary categories such as accepted and overridden consequently compress several cognitive and organizational transitions into a single recorded event. The proposed model treats those transitions as analytically distinct.

Figure 1 presents the original conceptual synthesis for medication-sensemaking model for hospital pharmacy, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

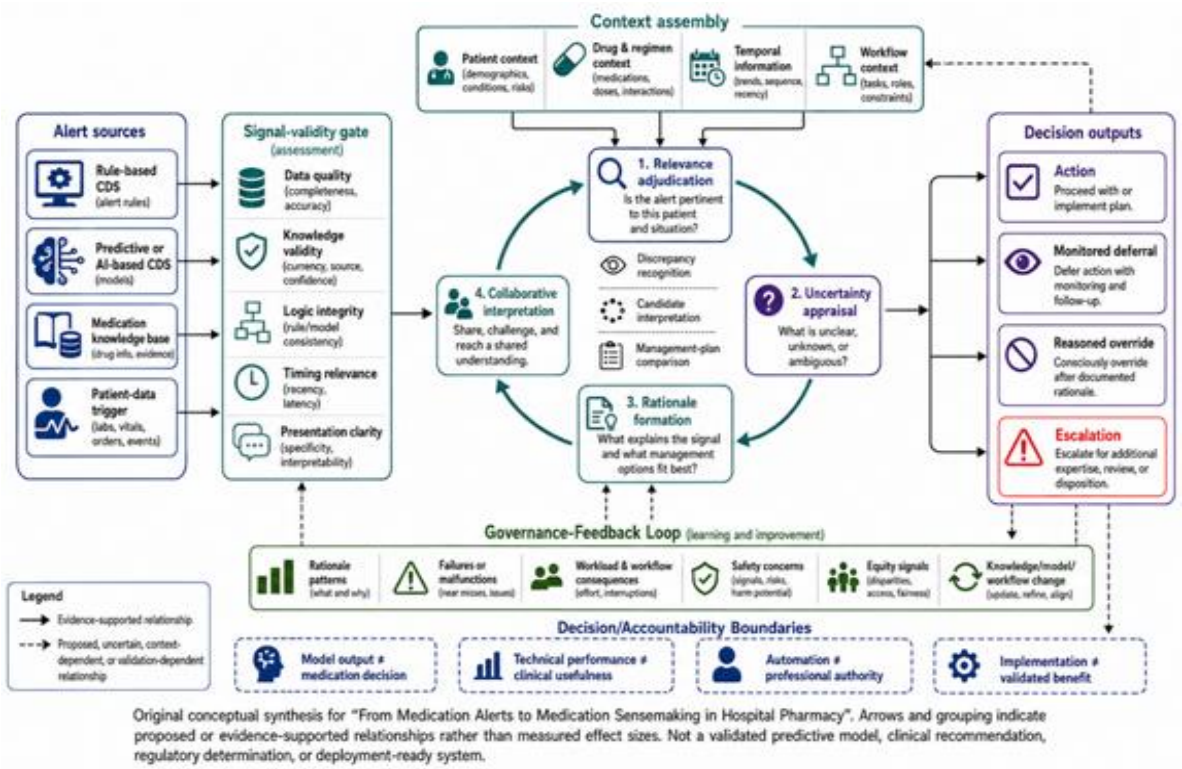

Figure 1. Medication-Sensemaking Model for Hospital Pharmacy

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article From Medication Alerts to Medication Sensemaking in Hospital Pharmacy.

Table 1. Definitions, components, relationships, and intended uses for the article From Medication Alerts to Medication Sensemaking in Hospital Pharmacy

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Alert signal	Possible medication problems may remain unnoticed.	Rules, models, medication knowledge, patient data, thresholds.	Detection presents a possible discrepancy for review; it does not determine its meaning.	Improves visibility of potentially relevant conditions.	Technical validity and task-specific detection performance.	Excessive sensitivity, omission, repetition, or poor timing [1, 2].	A signal is not a medication decision.
Signal-validity assessment	Displayed alerts may arise from defective data, logic, knowledge, or configuration.	Data currency, knowledge source, logic, interface, maintenance.	**Proposed:** examine signal integrity before clinical interpretation.	Prevents invalid alerts from being treated as authoritative.	Error analysis and malfunction surveillance [7].	Hidden defects may produce plausible but misleading alerts.	Technical availability does not establish validity.
Context assembly	Isolated alerts may omit information required for interpretation.	Patient, drug, regimen, laboratory, temporal, treatment-goal, and workflow information.	**Proposed:** assemble material context before relevance adjudication.	Supports patient- and task-specific interpretation.	Context completeness, accessibility, and timeliness; context-sensitive support evidence [6].	Missing, stale, conflicting, or inaccessible context.	More data do not automatically produce better understanding.
Attention and engagement	Alert exposure may be misclassified as cognitive review.	Workload, repetition, interruption, role, presentation.	Detection, attention, acceptance, and action are treated as separate stages [8].	Clarifies where the alert pathway fails.	Observation of user interaction and cognitive work.	Passive exposure recorded as meaningful engagement.	Acceptance and override are incomplete process measures.
Relevance adjudication	A technically correct alert may not matter for the present patient.	Clinical state, treatment goal, alternatives, reversibility, urgency.	**Proposed:** compare the signal with the patient's present management context.	Distinguishes potential hazard from actionable relevance.	Task-specific clinical review and prospective evaluation.	Inconsistent interpretation or overreliance on undocumented knowledge.	Relevance remains conditional on context.

Collaborative interpretation	Relevant knowledge may be distributed across professions and systems.	Pharmacist, prescriber, nurse, digital output, shared record.	**Proposed:** combine role-specific knowledge without equating contributions or authority.	Makes disagreement and missing information visible.	Team-process, communication, and responsibility evidence.	Automation bias, diffusion of responsibility, or unresolved disagreement.	Collaboration does not eliminate accountable authority.
Escalation and governance feedback	High-risk uncertainty and repeated failures may remain local and invisible.	Severity, uncertainty, disagreement, time sensitivity, recurring alert patterns.	**Proposed:** escalate cases beyond a defined boundary and return patterns to governance.	Links individual review with organizational learning.	Escalation reliability, safety review, and lifecycle monitoring.	Delayed escalation, unclear ownership, or unreviewed system drift.	No universal escalation threshold is established.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article From Medication Alerts to Medication Sensemaking in Hospital Pharmacy.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article From Medication Alerts to Medication Sensemaking in Hospital Pharmacy

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Detection layer	Identify a possible medication issue.	Patient and medication data to rule or model output.	Supplies the initial signal to validity review.	Task-specific detection and error characterization.	Missed hazard, excessive alerting, or unstable output [2].	Technical owners and clinical knowledge stewards.	Detection performance alone is insufficient.
Signal-validity gate	Assess data, logic, knowledge, timing, and presentation.	Alert provenance and supporting inputs to reviewer.	Determines whether contextual interpretation should proceed.	Traceability and malfunction detection [7].	Plausible alert generated from defective inputs.	Informatics, pharmacy, and system-governance teams.	Operational availability is not proof of reliability.
Context assembly	Gather information material to interpretation.	Patient, drug, treatment, temporal, and workflow information to reviewer.	Supports relevance and uncertainty assessment.	Availability, currency, usability, and task fit [6].	Fragmented, delayed, or misleading information.	Clinical teams and information-system owners.	Context sufficiency requires task-specific validation.
Engagement stage	Support attention and meaningful review.	Alert presentation to user response and rationale.	Connects detection with interpretation.	Distinguish exposure, attention, acceptance, and action [8].	Click-through behavior treated as understanding.	Interface designers, managers, and users.	Response rates do not establish decision quality.
Sensemaking cycle	Construct a patient- and task-specific interpretation.	Context and professional knowledge to rationale and candidate action.	Interacts with collaboration, uncertainty, and escalation.	**Proposed:** contextual completeness, rationale quality, and interpretive consistency.	Premature closure, anchoring, or undocumented assumptions.	Accountable medication professionals and care teams.	No validated sensemaking score or threshold exists.
Collaborative and escalation layer	Resolve disagreement and transfer unresolved risk.	Rationale, uncertainty, and requests across roles.	Connects interpretation with action and governance.	**Proposed:** timely acknowledgment, clear ownership, and closed-loop communication.	Suppressed concerns, responsibility gaps, or delayed response.	Hospital leadership and locally assigned professionals.	Authority arrangements remain locally governed.
Governance-learning layer	Review recurring failures, burden, and change.	Alert responses, malfunctions, workload effects, and safety concerns to oversight.	Modifies knowledge, design, workflow, or monitoring.	**Proposed:** transparent review and controlled change.	Unmonitored drift or optimization based only on alert counts.	Multidisciplinary governance body.	Implementation does not establish validated benefit.

Proposed Medication-Sensemaking Model

The proposed Medication-Sensemaking Model treats medication-alert work as a sociotechnical process rather than a direct transmission from algorithm to decision. Clinical decision support depends on data, knowledge resources, interfaces, users, workflows, and organizational arrangements [9]. The model therefore begins outside the alert itself. It asks how a signal was generated, what

information it represents, who receives it, what authority that person holds, and how subsequent reasoning becomes visible to others.

The first component is the signal-validity gate. Before relevance is considered, the system or reviewer examines whether the alert is supported by current data, valid knowledge, appropriate logic, and interpretable presentation. This gate is not intended to require a complete technical investigation during each medication review. It represents a design obligation: sufficient provenance and supporting information should be available for users to recognize when an alert may be unreliable.

The second component is context assembly. Relevant patient, medication, temporal, and workflow information is brought into the same reasoning space. Context assembly is intentionally selective rather than exhaustive. A model that displays every available datum may increase cognitive burden without improving interpretation. The proposed requirement is that information material to the alert's validity, relevance, urgency, reversibility, and available alternatives should be identifiable and accessible.

The third component is the sensemaking cycle, comprising relevance adjudication, discrepancy recognition, uncertainty appraisal, candidate interpretation, comparison with the treatment plan, and rationale formation. These operations may occur rapidly and iteratively rather than in a fixed sequence. Their separation is analytical: it makes possible the study of whether failures arise from missing information, misleading system output, unresolved uncertainty, or premature movement from signal to action.

The model is also bounded by implementation conditions. Adoption, abandonment, scale-up, and sustainability of health technologies depend on interactions among the technology, its users, organizational capacity, perceived value, and the wider system [10]. Accordingly, a medication-sensemaking process cannot be assessed only by examining individual clinicians. Staffing, workload, local medication policies, system maintenance, interoperability, and opportunities for adaptation may all shape whether the proposed functions can occur.

Human-factors evaluation is required because the same information can support or obstruct reasoning depending on its timing, layout, cognitive demands, and relation to actual tasks [11]. An explanation that is technically complete but difficult to interpret under hospital conditions may have little operational value. Conversely, a concise alert may be useful when its limits are visible and relevant context is easily obtainable. The model therefore distinguishes information quantity from usability and explanation availability from interpretive value.

Implementation experience may also diverge from anticipated technical benefit [12]. Users may distrust a system that interrupts work without explaining relevance, overtrust one that presents confident outputs, or develop workarounds when required actions conflict with clinical priorities. These responses should not automatically be attributed to resistance. They may reveal misalignment among system design, task requirements, professional responsibility, and organizational conditions.

The integrated architecture is proposed rather than validated. Its components may organize system design, workflow analysis, education, and prospective evaluation, but their ordering and relative importance are expected to vary across alert types. A dose alert, drug interaction, pharmacogenomic recommendation, renal-function warning, and therapeutic-monitoring signal may require different contextual inputs and escalation boundaries. The model consequently provides a common reasoning structure without claiming a single operating protocol.

Constructing Meaning from Patient, Drug, and Workflow Context

Medication meaning is produced by relationships among the patient, the drug, the regimen, the treatment objective, and the circumstances of care. A clinical-pharmaceutical assessment can distinguish the general appropriateness of alert content from its relevance to an individual patient [13]. This distinction is fundamental. A knowledge base may correctly identify a potential interaction or dosing concern while the patient's monitoring plan, treatment history, alternative options, or therapeutic priorities make a different response reasonable.

Patient context includes physiological state, diagnoses, organ function, laboratory trajectories, allergies, prior treatment response, vulnerability to harm, and preferences where relevant. It should also include time. A laboratory value may be recent but no longer representative; a dose may be inappropriate at initiation but reasonable during monitored titration; or an interaction may require surveillance rather than immediate discontinuation. Context therefore concerns relationships and trajectories, not merely the presence of individual data fields.

Renal medication alerts illustrate this requirement. Their interpretation may depend on the laboratory measure used, the direction and speed of renal change, the prescribed dose and interval, the indication, dialysis status, expected treatment duration, and whether monitoring or an alternative regimen is already planned [14]. A static alert may not represent these dependencies. The proposed model consequently treats renal function as part of a temporally situated medication decision rather than a single threshold value.

Drug and regimen context includes indication, dose, route, formulation, timing, treatment duration, co-medications, prior exposure, monitoring, and available alternatives. For

drug–drug interaction alerts, override occurrence should not be interpreted as inherently appropriate or inappropriate. Assessment requires attention to the interaction mechanism, patient vulnerability, intended management strategy, and possible adverse consequences [15]. A reasoned override accompanied by monitoring may represent active risk management, whereas acceptance without understanding may not indicate a high-quality decision.

Pharmacogenomic alerts further demonstrate that detection and meaning are different. A genotype or phenotype signal must be connected to the affected medication, strength of the recommendation, treatment stage, alternative therapies, and clinician response [16]. The genetic result does not independently determine the decision. Its significance is constructed through the relationship between molecular information and the patient's present therapeutic context.

Workflow context is equally important. The professional receiving an alert may lack authority to change therapy, access to the full record, time to investigate, or an immediate route for contacting the responsible prescriber. A clinically relevant alert can therefore fail because the workflow cannot support interpretation or action. Conversely, a lower-severity signal may become useful when routed to a pharmacist with the necessary medication knowledge, patient information, and communication pathway.

The model proposes a context-sufficiency question before an alert is converted into action: is the available information sufficiently current, coherent, and relevant to support a defensible response? "Sufficient" does not mean complete. It means that material uncertainty has been identified and that missing information is either unlikely to change the decision or has triggered further review, monitoring, or escalation. The answer will remain task- and setting-dependent.

Contextualization also has limits. More information can increase noise, delay urgent action, or create a false impression of certainty. Automated context assembly may reproduce documentation errors or privilege easily measured variables over clinically important but poorly structured knowledge. Professional interpretation may introduce inconsistency, bias, or undocumented assumptions. The proposed model therefore does not position either automation or unaided judgment as sufficient. It makes their respective contributions, uncertainties, and responsibility boundaries explicit so that they can be evaluated.

Collaborative Interpretation and Escalation

Medication-alert interpretation may require knowledge distributed among pharmacists, prescribers, nurses, patients, laboratory services, and digital systems. Collaboration should therefore not be reduced to forwarding an alert. It involves making the alert's basis, contextual interpretation, uncertainty, proposed action, and unresolved questions available to the professionals who hold relevant knowledge or authority.

Rationale capture provides the informational bridge between individual interpretation and collaborative review. Structured override reasons can make recurring causes of alert rejection visible and support later analysis [17]. Structured categories alone may, however, omit patient-specific reasoning. Free-text comments can preserve nuance, and computational methods may help organize those comments for review and alert redesign [17, 18]. Such summaries remain secondary representations of professional reasoning; they require human verification and must not be treated as authoritative clinical judgments.

Human–AI collaboration should be evaluated as a configured joint system rather than as the simple addition of a model to existing work. Evidence from clinical image interpretation indicates that combined performance depends on how outputs are communicated, how disagreements are handled, and which participant has the stronger information for the case [19]. In medication sensemaking, digital systems may identify patterns, retrieve knowledge, or expose inconsistencies, while professionals integrate undocumented context, treatment objectives, competing risks, and accountability. The proposed model therefore separates computational contribution from professional authority.

Escalation is required when the available context is insufficient, disagreement remains unresolved, potential harm is substantial, action is time-sensitive, or the receiving professional lacks authority to complete the medication decision. Speaking up is influenced not only by individual competence but also by safety climate and managerial response [20]. The model consequently proposes a closed-loop escalation process in which the concern, rationale, uncertainty, recipient, responsibility transfer, and response are visible. The appropriate threshold and responsible role must be defined locally and tested for the relevant medication task.

Evaluating Sensemaking rather than Alert Response

Evaluation should determine whether an alert supports the construction of a defensible medication interpretation, not merely whether it is displayed, accepted, or overridden. Explanation is one component, but its quality depends on audience, purpose, form, fidelity, and intended use [21]. A technically detailed explanation may remain unusable, while a persuasive explanation may conceal weak evidence or system error.

Evaluation should therefore preserve distinct outcome layers: signal validity, contextual completeness, relevance adjudication, rationale quality, medication-task performance, workflow burden, safety consequences, equity, implementation, and governance. AI-based decision support should be assessed for its effects on care processes and unintended organizational consequences as well as technical performance [22]. Improvements in alert acceptance should

not be interpreted automatically as improvements in decision quality or medication safety.

Early clinical evaluation should document user interaction, workflow integration, system errors, safety events, and modifications introduced during use [23]. Reports should also describe the intervention's required inputs, outputs, human–AI interaction, error handling, and failure analysis sufficiently clearly for interpretation and replication [24]. For the proposed model, evaluation would additionally examine whether material uncertainty was recognized, whether disagreement was resolved or escalated, and whether the recorded rationale was understandable and reviewable.

No validated composite measure of medication sensemaking is proposed. Candidate measures must be tested for specific alert classes, professional configurations, and clinical environments. Technical, clinical, workflow, safety, and equity outcomes should remain separately visible so that improvement in one dimension does not conceal deterioration in another.

Implications for Hospital-Pharmacy Systems

Artificial intelligence may help prioritize medication alerts, identify recurrent override patterns, assemble context, or route signals to appropriate reviewers. Current evidence remains heterogeneous, and many evaluations do not establish real-world clinical or organizational benefit [25]. AI should therefore be positioned as a configurable contribution to medication review rather than as evidence that interpretation has been completed.

Machine-learning alert systems may improve prioritization or reduce low-value interruption within defined environments [26]. Such performance is conditional on local data, alert definitions, workflow, and evaluation methods. It does not establish generalizable clinical sensemaking, professional substitutability, or safety across institutions.

Pharmacist-facing decision support may be most useful when aligned with a defined medication-review task, access to relevant patient information, and a clear intervention pathway [27]. This requires decisions about routing, response time, documentation, communication, and ownership. Systems should also allow pharmacists to identify misleading alerts, missing context, and recurrent workflow conflicts rather than treating all overrides as resistance.

Responsible AI-enabled decision support requires assigned accountability, documentation, monitoring, and mechanisms for responding to system failure and change [28]. Governance should review not only model performance but also alert burden, contextual mismatch, rationale patterns, escalation failures, safety concerns, and unequal effects across patient or professional groups.

Figure 2 presents the original conceptual synthesis for from alert detection to collaborative interpretation and escalation, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

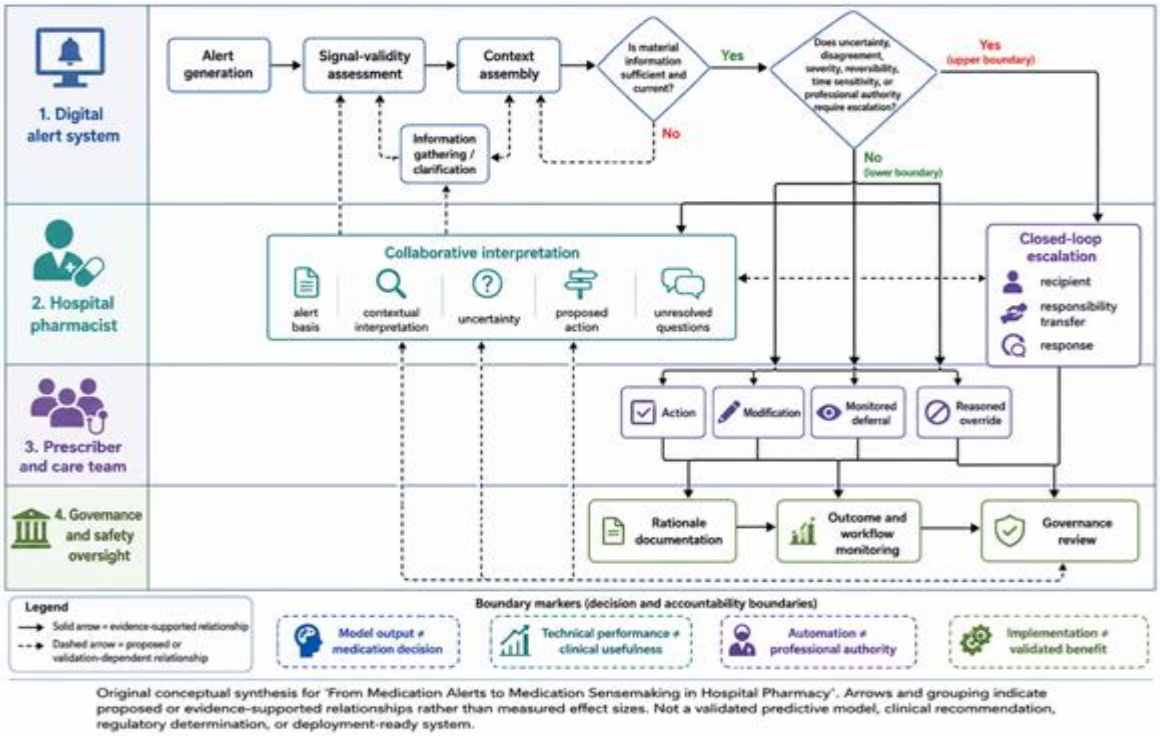

Figure 2. From Alert Detection to Collaborative Interpretation and Escalation

Boundary Conditions

The proposed model does not establish clinical benefit or deployment readiness. Responsible clinical-AI use requires representative evidence, clinically meaningful validation, workflow integration, stakeholder participation, and continued monitoring [29, 30]. Meeting these conditions reduces uncertainty but does not guarantee safety.

Explanations must not be treated as proof that a recommendation is correct or trustworthy, because post hoc methods may be unstable, incomplete, or persuasive without being faithful [31]. Similarly, apparently neutral targets can encode unequal access or utilization and reproduce inequity [32]. Applicability therefore remains conditional on the medication task, alert class, patient population, data-generating process, professional authority, organizational capacity, and local safety infrastructure.

Conclusion

Hospital-pharmacy alerting should be understood as the beginning of medication reasoning rather than its completion. The proposed Medication-Sensemaking Model distinguishes detection from the contextual and collaborative construction of actionable meaning. It connects signal validity, patient and drug context, temporal and workflow information, relevance adjudication, uncertainty appraisal, rationale formation, professional collaboration, escalation, evaluation, and governance feedback.

The model's scientific contribution is an explicit architecture for examining where medication-alert pathways succeed, fail, or conceal uncertainty. Its professional contribution is to preserve the pharmacist's role in contextual judgment while recognizing that medication knowledge and authority may be distributed across the care team. Its organizational contribution is to connect individual alert responses with system learning and accountable change control.

These contributions remain conceptual. The model does not prescribe a universal workflow, assign fixed professional authority, validate an AI system, or establish that greater automation improves medication safety. Empirical work is required to test its components and transitions across alert classes, patient groups, institutions, and professional arrangements. Evaluation must keep technical performance, medication-task quality, workflow consequences, safety, equity, and governance distinct. Under these boundaries, medication sensemaking provides a structured basis for designing hospital-pharmacy systems that support professional interpretation without mistaking alerts for decisions.

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

References

1. Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. BMC Med Inform Decis Mak. 2017;17(1):36. doi:10.1186/s12911-017-0430-8
2. Co Z, Holmgren AJ, Classen DC, Newmark L, Seger DL, Danforth M, et al. The tradeoffs between safety and alert fatigue: Data from a national evaluation of hospital medication-related clinical decision support. J Am Med Inform Assoc. 2020;27(8):1252-8. doi:10.1093/jamia/ocaa098
3. 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
4. Van Dort BA, Zheng WY, Sundar V, Baysari MT. Optimizing clinical decision support alerts in electronic medical records: A systematic review of reported strategies adopted by hospitals. J Am Med Inform Assoc. 2021;28(1):177-83. doi:10.1093/jamia/ocaa279
5. Leykum LK, O'Leary K. Web Exclusives. Annals for Hospitalists Inpatient Notes—Sensemaking—fostering a shared understanding in clinical teams. Ann Intern Med. 2017;167(4). doi:10.7326/M17-1829
6. Cook DA, Teixeira MT, Heale BSE, Cimino JJ, Del Fiol G. Context-sensitive decision support (infobuttons) in electronic health records: A systematic review. J Am Med Inform Assoc. 2017;24(2):460-8. doi:10.1093/jamia/ocw104
7. Wright A, Ai A, Ash J, Wiesen JF, Hickman TT, Aaron S, et al. Clinical decision support alert malfunctions: Analysis and empirically derived taxonomy. J Am Med Inform Assoc. 2018;25(5):496-506. doi:10.1093/jamia/ocx106
8. Murad DA, Tsugawa Y, Elashoff DA, Baldwin KM, Bell DS. Distinct components of alert fatigue in physicians' responses to a noninterruptive clinical decision support alert. J Am Med Inform Assoc. 2023;30(1):64-72. doi:10.1093/jamia/ocac191
9. 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
10. Greenhalgh T, Wherton J, Papoutsi C, Lynch J, Hughes G, A'Court C, et al. Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. J Med Internet Res. 2017;19(11). doi:10.2196/jmir.8775
11. Salwei ME, Hoonakker P, Carayon P, Wiegmann D, Pulia M, Patterson BW. Usability of a human factors-based clinical decision support in the emergency department: Lessons learned for design and implementation. Hum Factors. 2024;66(3):647-57. doi:10.1177/00187208221078625
12. Romero-Brufau S, Wyatt KD, Boyum P, Mickelson M, Moore M, Cognetta-Rieke C. A lesson in implementation: A pre-post study of providers' experience with artificial intelligence-based clinical decision support. Int J Med Inform. 2020;137:104072. doi:10.1016/j.ijmedinf.2019.104072
13. 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
14. Shah SN, Amato MG, Garlo KG, Seger DL, Bates DW. Renal medication-related clinical decision support (CDS) alerts and overrides in the inpatient setting following implementation of a commercial electronic health record: Implications for designing more effective alerts. J Am Med Inform Assoc. 2021;28(6):1081-7. doi:10.1093/jamia/ocaa222
15. Edrees H, Amato MG, Wong A, Seger DL, Bates DW. High-priority drug-drug interaction clinical decision support overrides in a newly implemented commercial computerized provider order-entry system: Override appropriateness and adverse drug events. J Am Med Inform Assoc. 2020;27(6):893-900. doi:10.1093/jamia/ocaa034
16. Nguyen JQ, Crews KR, Moore BT, Kornegay NM, Baker DK, Hasan M, et al. Clinician adherence to pharmacogenomics prescribing recommendations in clinical decision support alerts. J Am Med Inform Assoc. 2023;30(1):132-8. doi:10.1093/jamia/ocac187

17. Wright A, McEvoy DS, Aaron S, McCoy AB, Amato MG, Kim H, et al. Structured override reasons for drug-drug interaction alerts in electronic health records. J Am Med Inform Assoc. 2019;26(10):934-42. doi:10.1093/jamia/ocz033
18. Liu S, McCoy AB, Wright AP, Nelson SD, Huang SS, Ahmad HB, et al. Why do users override alerts? Utilizing large language model to summarize comments and optimize clinical decision support. J Am Med Inform Assoc. 2024;31(6):1388-96. doi:10.1093/jamia/ocae041
19. 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
20. Alingh CW, van Wijngaarden JDH, van de Voorde K, Paauwe J, Huijsman R. Speaking up about patient safety concerns: The influence of safety management approaches and climate on nurses' willingness to speak up. BMJ Qual Saf. 2019;28(1):39-48. doi:10.1136/bmjqs-2017-007163
21. Markus AF, Kors JA, Rijnbeek PR. The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies. J Biomed Inform. 2021;113:103655. doi:10.1016/j.jbi.2020.103655
22. Cresswell K, Rigby M, Magrabi F, Scott P, Brender J, Craven CK, et al. The need to strengthen the evaluation of the impact of artificial intelligence-based decision support systems on healthcare provision. Health Policy. 2023;136:104889. doi:10.1016/j.healthpol.2023.104889
23. 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
24. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
25. 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
26. Chen CY, Chen YL, Scholl J, Yang HC, Li YCJ. Ability of machine-learning-based clinical decision support system to reduce alert fatigue, wrong-drug errors, and alert users about look alike, sound alike medication. Comput Methods Programs Biomed. 2024;243:107869. doi:10.1016/j.cmpb.2023.107869
27. Robert L, Vidoni N, Gérard E, Chazard E, Odou P, Rousselière C, et al. Implementation of a clinical decision support system for the optimization of antidiabetic drug orders by pharmacists. Br J Clin Pharmacol. 2024;90(1):239-46. doi:10.1111/bcp.15898
28. Labkoff S, Oladimeji B, Kannry J, Solomonides A, Leftwich R, Koski E, et al. Toward a responsible future: Recommendations for AI-enabled clinical decision support. J Am Med Inform Assoc. 2024;31(11):2730-9. doi:10.1093/jamia/ocae209
29. 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
30. 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
31. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9
32. 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