Original Article

Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions

James Anderson[1], Maria Rossi[2]*, William Clark[3]

[1]Department of Human-AI Coordination in Pharmacy, Faculty of Pharmacy, University of Manchester, Manchester, United Kingdom. [2]Department of Pharmacy Informatics and Decision Support, Faculty of Pharmacy, University of Milan, Milan, Italy. [3]Department of Digital Health and Clinical Workflow, Faculty of Pharmacy, University of Melbourne, Melbourne, Australia.

Abstract

Digital pharmacy increasingly places algorithmic functions within medication workflows, yet the pharmacy workbench is still commonly treated as a physical workstation or software interface. This framing is insufficient because medication decisions emerge from interactions among patient context, prescription information, professional judgment, technical outputs, temporal pressures, communication, and organizational authority. This article proposes a sociotechnical architecture that reconceptualizes the pharmacy workbench as a human–AI coordination space rather than an automation endpoint. The architecture comprises six connected layers: patient context; prescription and medication information; bounded algorithmic functions; pharmacist judgment and action; coordination-state and temporal control; and escalation, governance, and learning. Algorithmic functions may detect, retrieve, prioritize, predict, summarize, explain, estimate uncertainty, or abstain, but their outputs remain distinct from medication decisions and professional authorization. Proposed coordination states make incomplete context, human–AI discordance, interruption, deferral, uncertainty, escalation, and closure visible rather than allowing them to remain implicit. Evaluation is organized across technical validity, clinical-task relevance, human–AI team performance, workflow fit, medication safety, equity, governance, external validity, and implementation sustainability. The architecture additionally requires task ownership, information provenance, resumption support, escalation acknowledgement, auditability, and controlled system change. It does not establish clinical effectiveness, reduced workload, improved medication safety, or deployment readiness. Its components, transition rules, and decision boundaries require empirical evaluation across pharmacy settings, medication tasks, professional roles, populations, and organizational conditions. The original contribution is an evidence-bounded architecture for studying and designing coordinated medication work without treating any model, alert, professional action, or data source as sufficient by itself.

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

Introduction

Artificial intelligence and clinical decision support may augment professional decision work, but their value and safety depend on how model outputs, medication alerts, clinical context, and professional action are brought together [1–3]. In pharmacy practice, this interaction occurs within a work environment traditionally described through counters, dispensing stations, electronic records, verification screens, communication tools, and prescription queues. Such descriptions identify visible equipment but underrepresent the distributed reasoning through which medication information is assembled, questioned, reconciled, communicated, and converted into action.

Current pharmacy applications of artificial intelligence are heterogeneous, spanning operational, informational, clinical, and patient-facing functions, without yet establishing a unified architecture for coordinating medication decisions [4]. A predictive model may identify risk without determining what the pharmacist should do. An alert may detect a possible interaction without resolving indication, treatment priority, patient preference, monitoring capacity, or prescriber intent. Likewise, pharmacist review cannot be represented as a final binary approval step because professional judgment may involve seeking missing information, negotiating alternatives, documenting uncertainty, counselling a patient, or transferring responsibility.

Address for correspondence: Maria Rossi, Department of Pharmacy Informatics and Decision Support, Faculty of Pharmacy, University of Milan, Milan, Italy.
maria.rossi@unimi.it

How to cite this article: Anderson J, Rossi M, Clark W. Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions. Arch Pharm Pract. 2025;16(1):34-42. https://doi.org/10.51847/ng72yn5h2b

The problem is therefore not simply where to place an AI tool on an existing workstation. It is how to organize relationships among patient context, prescription objects, algorithmic functions, professional authority, communication, timing, and escalation. The proposed concept of a human–AI pharmacy workbench refers to this wider coordination space. It includes the visible interface but also the actors, information dependencies, task states, organizational rules, and accountability structures through which medication work proceeds.

This article develops an original sociotechnical architecture. Evidence is used to define established workflow, human-factors, implementation, and safety problems. The integration of those problems into layered architecture, coordination states, transition conditions, and escalation boundaries is proposed in this article. The framework does not assume that AI improves medication decisions, that pharmacists can invariably identify model error, or that one design will apply across community, hospital, ambulatory, specialist, or remote pharmacy services.

From Physical Workstation to Sociotechnical Coordination Space

A physical workstation is the material and digital location at which medication tasks are displayed and performed. Its components may include dispensing equipment, electronic records, verification interfaces, communication channels, reference resources, and decision-support notifications. However, medication-related decision support is accepted, ignored, or worked around partly according to its relevance, timing, disruption, and fit with professional work [5]. The analytic unit must therefore extend beyond the device or screen.

A sociotechnical coordination space is defined here as the configuration of people, technologies, information objects, routines, authority, communication, and temporal conditions through which medication decisions are produced. Medication-safety work has been shown to depend on interactions among social and technical conditions rather than on isolated technologies [6]. Under this interpretation, the workbench is not a passive site where information is viewed. It is an active boundary through which incomplete and sometimes conflicting information is transformed into situated professional action.

Electronic prescribing illustrates this distinction. Digitization may alter communication, role distribution, work visibility, and local practices rather than merely replace paper with an electronic order [7]. An apparently complete prescription can remain operationally incomplete when its indication, clinical priority, recent laboratory context, patient preference, or communication history is unavailable. Conversely, an interface may contain abundant data while failing to show which elements are current, relevant, disputed, or awaiting confirmation.

Human-factors evidence further indicates that usability and workflow integration must be evaluated in relation to actual users, tasks, and implementation environments [8]. A workbench may be technically functional yet poorly coordinated if information arrives too late, actions are distributed across disconnected systems, unresolved work has no owner, or interruption removes the reasoning context needed for safe resumption. The proposed shift from workstation to coordination space consequently changes the design question from "What can the system display?" to "What must the human–technology system preserve, communicate, and resolve before a medication decision can be closed?"

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions.

Table 1. Definitions, components, relationships, and intended uses for the article Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Physical workbench	Equipment-centred descriptions omit distributed reasoning.	Hardware, software, workspace, task queue	Provides the visible site of work but not its complete coordination structure.	Defines the material boundary from which analysis begins.	Observation of task and interface use [5].	Interface performance may be mistaken for workflow performance.	A workstation is not equivalent to the proposed coordination space.
Sociotechnical coordination space	Isolated components cannot explain medication-safety work.	People, technology, routines, organization, context	Interactions among components shape how medication work proceeds [6].	Supports system-level design and evaluation.	Multilevel workflow and organizational evidence	Relationships remain setting-dependent.	The construct organizes inquiry; it does not establish causation.

Digital workflow	Electronic orders may redistribute rather than eliminate work.	Prescribing system, messages, roles, local practices	Digitization changes visibility, communication, and task allocation [7].	Identifies hidden coordination consequences.	Comparative workflow and communication evidence	New work or workarounds may remain undocumented.	Electronic transmission does not establish coordination completeness.
Human-factors fit	A usable function may still conflict with actual tasks.	User characteristics, timing, workload, interface sequence	Design is aligned with real task demands and implementation conditions [8].	May reduce avoidable interaction burden.	Usability, workload, and workflow testing	Laboratory usability may not transfer to practice.	Usability does not demonstrate clinical benefit.
Patient-context layer	Prescription data alone may not represent the decision situation.	Clinical, experiential, social, access, and preference information	**Proposed:** context objects are shown with provenance, recency, and completeness.	Makes missing or disputed context visible.	Task-specific validation of required context	Relevant context may be unavailable or inaccurate.	Completeness requirements vary by task and patient.
Algorithmic function	"AI" can conceal materially different functions.	Data, task definition, model, output, uncertainty	**Proposed:** detection, prediction, retrieval, summarization, explanation, and abstention remain distinct.	Prevents one output from being mistaken for a decision.	Function-specific technical and clinical evidence [3].	Function boundaries may be blurred in an interface.	Model output does not confer medication authority.
Pharmacist judgment	Oversight may be represented as simple approval.	Evidence, expertise, context, communication, accountability	**Proposed:** pharmacists interpret, reconcile, question, document, and disposition outputs.	Preserves professional reasoning and responsibility.	Human–AI task-performance evidence	Human review may be biased, rushed, or over-reliant.	Human involvement alone does not guarantee safety [1].
Coordination state	Unresolved work may remain implicit.	Context sufficiency, agreement, risk, uncertainty, interruption	**Proposed:** explicit states represent alignment, clarification, discordance, deferral, escalation, and closure.	Supports ownership, continuity, and auditability.	Prospective workflow and safety evaluation	State labels may oversimplify complex work.	States are proposed constructs, not clinical classifications.

Actors, Information Objects, Tasks, and Temporal Dependencies

Medication decision support should be examined across a lifecycle rather than only at the moment an alert or prediction appears [9]. Relevant work begins when a medication-related need is recognized and may continue through context acquisition, prescription interpretation, technical screening, pharmacist assessment, communication, intervention, decision closure, follow-up, and system review. Different actors may enter and leave this lifecycle, while information and responsibility persist across time.

The proposed actor model includes patients, caregivers, pharmacists, pharmacy technicians, prescribers, nurses, specialists, organizational leaders, software teams, and governance bodies. Technical systems may act as informational agents by retrieving, sorting, detecting, predicting, or presenting information, but they are not treated as accountable clinical actors. Patient-centred decision support requires coordination among policies, processes, technologies, expertise, and workflow execution [10]. The architecture therefore distinguishes participation in a task from authority to make, authorize, communicate, or accept a medication decision.

An information object is proposed as a bounded unit of information that has content, provenance, time status, ownership, and a relationship to an unresolved task. Examples include a prescription, medication history, allergy record, laboratory value, interaction signal, model prediction, pharmacist note, patient preference, prescriber response, or escalation request. This definition differs from a raw data field because an information object has a coordination function. Its significance depends not only on its value but also on whether it is current, verified, contested, incomplete, or awaiting action.

Digitally encoded task structures may fail when their sequence or content does not correspond to frontline workflow [11]. The workbench should therefore avoid assuming that every medication decision follows a fixed linear pathway. A pharmacist may identify missing indication information after reviewing an interaction alert, discover a recent dose change during patient counselling, or receive a delayed prescriber response after another professional has assumed responsibility. Proposed dependency markers should show what information is required, what work has been completed, what remains unresolved, and whether another task must occur first.

Tasks are similarly differentiated. Technical tasks include data retrieval, consistency checks, calculation, classification, and alert generation. Interpretive tasks include evaluating relevance, reconciling conflicting evidence, assessing uncertainty, and comparing alternatives. Coordination tasks include assigning ownership, requesting clarification, confirming communication, and managing handoffs. Accountable actions include authorizing dispensing, communicating an intervention, documenting a decision, or escalating beyond the pharmacist's authority or expertise. These categories are proposed to prevent technical capability from being conflated with professional responsibility.

Time is not merely a timestamp attached to the prescription. It affects whether information remains valid, whether delay is clinically meaningful, whether a response is still pending, and whether work can safely be resumed. Community-pharmacy work is interruption-prone, with interruptions arising from people, communication systems, competing tasks, and environmental demands [12]. The architecture consequently treats interruption as a coordination state rather than a momentary inconvenience. A resumption record should preserve the task owner, completed actions, reasoning already undertaken, unresolved questions, information awaited, and time sensitivity.

Temporal visibility may also reveal alternative explanations for apparent system failure. An ignored alert may reflect poor relevance, but it may instead have arrived after the decision was effectively made. A delayed intervention may indicate workflow congestion rather than lack of clinical concern. A pharmacist override may represent informed contextual correction, inappropriate reliance on habit, or an incomplete model input. The workbench must therefore preserve enough sequence and provenance to permit later interpretation without assuming that any single action reveals decision quality.

The proposed actor, object, task, and time models do not establish which information should be mandatory for every medication decision. Requirements must be adapted to the medication, patient, service, professional role, and consequence of delay. Their purpose is to make coordination dependencies inspectable and testable rather than leaving them embedded invisibly within interfaces and local routines.

Proposed Human–AI Pharmacy Workbench Architecture

The proposed architecture organizes the workbench into six connected layers spanning context acquisition, medication information, algorithmic processing, professional judgment, coordination-state management, and organizational governance. End-to-end clinical AI frameworks indicate that implementation conditions extend beyond model development to workflow, users, evaluation, deployment, and continuing oversight [13]. The workbench accordingly treats the algorithm as one bounded component within a wider medication-work system.

Figure 1 presents the original conceptual synthesis for human–ai pharmacy workbench architecture, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

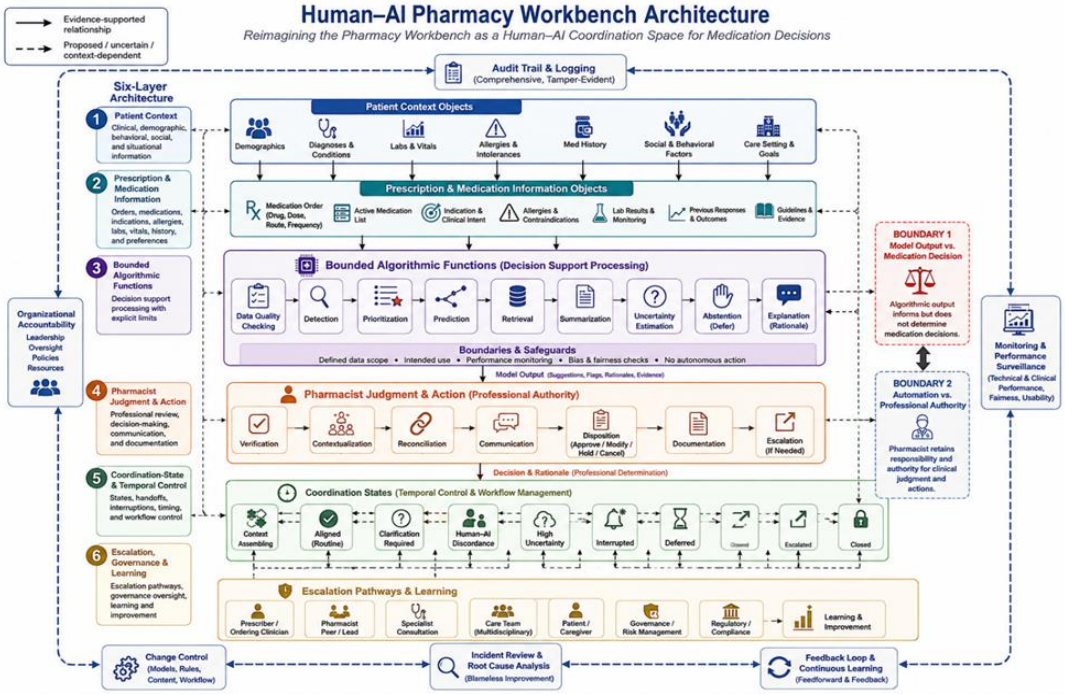

Figure 1. Human–AI Pharmacy Workbench Architecture.

The figure is an original conceptual synthesis developed for "Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

The first layer assembles patient context, including clinical characteristics, treatment goals, prior medication experience, access conditions, and expressed preferences. The second represents prescription and medication information, including medicine, formulation, indication, dose, route, duration, concurrent therapy, modification history, and communication status. Objects in both layers should display provenance, recency, completeness, and unresolved conflict where relevant.

The third layer contains bounded algorithmic functions. Responsible clinical AI requires attention to problem definition, data, evaluation, deployment context, and downstream consequences; explanation displays alone do not establish safe understanding, while intrinsically interpretable approaches should be considered where feasible for high-stakes tasks [14–16]. The architecture therefore separates detection, retrieval, prioritization, prediction, summarization, explanation, uncertainty estimation, and abstention. No function independently authorizes prescribing, dispensing, or treatment change.

The fourth layer represents pharmacist judgment and action. Here, algorithmic outputs are interpreted against patient context, prescription intent, alternative evidence, professional knowledge, and practical care conditions. Available dispositions may include acceptance, modification, rejection, clarification, deferral, communication, counselling, or escalation. Documentation should record the decision rationale and remaining uncertainty without implying that a recorded explanation proves the quality of reasoning.

The fifth layer maintains coordination state and temporal continuity. It identifies whether information is being assembled, review is pending, the human and system are aligned, clarification is required, disagreement remains unresolved, uncertainty is high, work has been interrupted, or escalation is underway. The sixth layer connects the encounter to governance and learning through accountability assignment, audit trails, incident review, monitoring, access control, change management, and restriction or withdrawal of unsafe functions.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Patient-context layer	Assemble decision-relevant context.	Patient, clinical, experiential, and access information enter with provenance.	Supplies algorithmic and pharmacist review.	Context relevance, completeness, recency, and subgroup coverage	Missing or stale context is treated as absence of risk.	Define required context and manage unresolved gaps.	No use when essential context cannot be obtained or represented.
Prescription-information layer	Represent medicine, indication, regimen, status, and intent.	Order data and communication history remain linked.	Constrains screening and professional interpretation.	Accuracy, reconciliation, version control, and intent visibility	Superseded or ambiguous instructions are acted upon.	Confirm discrepancies and preserve amendment history.	Transmission alone does not establish a valid medication decision.
Algorithmic-function layer	Perform a bounded computational task.	Inputs generate outputs, uncertainty, provenance, or abstention.	Informs but does not replace pharmacist disposition.	Task validity, calibration, robustness, uncertainty, and interpretability [14].	A technically plausible output is used outside its intended task.	Specify intended use, monitor performance, and restrict misuse.	Model performance does not establish clinical usefulness.
Explanation or representation process	Support examination of an output.	Features, rules, examples, or rationale are displayed.	May inform verification and challenge.	Fidelity, comprehensibility, decision relevance, and usability	Persuasive but invalid explanation increases reliance.	Test user understanding and preserve access to source evidence.	Explanation does not demonstrate causal validity or safety [15].
Pharmacist-judgment layer	Interpret and disposition information.	Context, evidence, model output, and	Updates coordination state	Appropriate reliance, rationale quality,	Automation bias, under-reliance, rushed approval,	Retain professional authority, training, and escalation access.	Human review alone is not

		uncertainty are compared.	and initiates communication.	workload, and task performance	or undocumented correction		evidence of safe oversight.
Coordination-state layer	Preserve task status, ownership, and temporal continuity.	State changes follow context, agreement, interruption, risk, and uncertainty.	Connects active work to clarification, deferral, escalation, or closure.	State accuracy, resumption completeness, and timely transition	Unresolved work appears complete or loses ownership.	Define state-entry, state-exit, and handoff responsibilities.	Proposed states require prospective workflow validation.
Escalation process	Transfer unresolved work to appropriate authority or expertise.	Risk, disagreement, uncertainty, or missing context accompanies the request.	Links pharmacist review to prescriber, specialist, patient, or governance actor.	Timeliness, acknowledgement, information completeness, and closure	Escalation is sent but not received, understood, or acted upon.	Assign recipient capacity, backup routes, and follow-up ownership.	Escalation capability must exist before the pathway is relied upon.
Governance and learning layer	Monitor, review, and control system change.	Decisions, incidents, overrides, failures, and adaptations enter oversight.	May modify any architecture layer through controlled change.	Accountability, auditability, change control, equity review, and monitoring [13].	Performance drift or unsafe workflow adaptation remains undetected.	Maintain multidisciplinary oversight and authority to restrict use.	Implementation does not establish validated benefit or continuing safety.

Coordination States, Handoffs, Interruptions, and Escalation

Human–AI performance cannot be inferred by comparing model accuracy with unaided professional accuracy because outcomes also depend on how advice is presented, interpreted, and incorporated into work [17]. The proposed workbench therefore represents medication work through explicit coordination states rather than treating every algorithmic output as a recommendation awaiting acceptance.

The initial state is context assembling, in which required patient and prescription objects are identified and checked. Work may then enter review pending, followed by aligned and routine when the pharmacist considers the available information sufficient and finds no unresolved conflict requiring escalation. Alignment does not establish that the decision is correct; it records only the current relationship among available context, algorithmic output, and professional assessment.

A human–AI discordance state is required because erroneous assistance can reduce human performance [18]. Discordance should not automatically privilege either the model or the pharmacist. Instead, it should initiate examination of input validity, intended use, alternative evidence, uncertainty, and the professional rationale for disagreement. This safeguard is necessary because users may be influenced by incorrect algorithmic advice even when they possess relevant expertise [19].

The workbench should separately represent high or unresolved uncertainty. Uncertainty may arise from missing context, conflicting information, weak model confidence, limited applicability, or ambiguity in the medication task. Communicating uncertainty and permitting abstention can support safer review than forcing a binary output [20]. However, uncertainty displays require evaluation because poorly designed estimates may confuse users or create false reassurance.

Handoffs must transfer more than a task notification. The proposed minimum handoff packet includes the task owner, completed actions, unresolved question, information awaited, rationale already developed, degree of urgency, and required acknowledgement. An escalation is not complete when a message is sent; responsibility must be accepted by a person or team with appropriate expertise, authority, and capacity.

Interruption may occur from any active state. A persistent resumption checkpoint should preserve the last verified state, pending dependencies, and time sensitivity. These mechanisms are proposed coordination controls rather than validated safety interventions. Their usefulness must be tested against alternative explanations, including added documentation burden, excessive state complexity, delayed action, and inappropriate escalation.

Safety Controls and Evaluation Requirements

Evaluation must distinguish transparent model reporting, evaluation of an AI-containing clinical intervention, and early live assessment of human factors, workflow, and safety [21–23]. High technical performance is therefore insufficient when the output is poorly timed, misunderstood, applied outside its intended task, or integrated into a workflow that creates new failure opportunities.

Evaluation should proceed across linked levels. Model-level assessment includes data validity, discrimination, calibration, robustness, uncertainty, and subgroup performance. Task-level assessment examines whether outputs answer a clinically relevant medication question. Human–AI assessment evaluates comprehension, appropriate reliance, disagreement resolution, workload, and performance of the combined team. Workflow assessment examines interruptions, handoffs, timing, communication closure, and

unintended redistribution of work. Safety assessment must identify both intercepted and newly introduced hazards.

Governance should include multidisciplinary responsibility for intended use, access, monitoring, equity, incident review, change control, and authority to restrict or withdraw a function [24]. Proposed progression gates should permit continuation, redesign, restricted use, additional evidence collection, or rejection. No universal threshold is proposed: evidentiary sufficiency must reflect the medication task, potential harm, population, setting, and availability of effective escalation.

Implementation evaluation must also examine whether the workbench preserves professional accountability without creating nominal oversight, in which a pharmacist remains formally responsible but lacks time, information, or practical ability to challenge the system. Reporting compliance, user acceptance, or continued use must not be interpreted automatically as evidence of clinical effectiveness or safety.

Implications for Pharmacy Design and Professional Practice

The architecture shifts pharmacy design from adding isolated functions toward coordinating complete medication work. Adoption and sustainability will depend on interactions among the technology, its users, organizational arrangements, wider systems, and adaptation over time [25]. Implementation analysis should therefore consider the characteristics of the innovation, participating individuals, inner and outer settings, and implementation process [26].

Professional implications should be framed as task reconfiguration rather than predetermined replacement. Digital transformation may alter how pharmacists retrieve information, verify routine elements, manage exceptions, communicate, and document decisions [27]. The proposed workbench places greater emphasis on context reconstruction, uncertainty interpretation, disagreement management, patient communication, and escalation. These roles require organizational time and authority; they cannot be added safely as invisible work.

Productivity should not be reduced to prescription throughput, alerts dismissed, or time saved. Pharmacy technology must also be evaluated against the quality and value of professional activity, workforce configuration, and service contribution [28]. A faster pathway may be inferior if it suppresses relevant context or transfers unresolved work elsewhere.

Figure 2 presents the original conceptual synthesis for coordination-state and escalation pathway for medication decisions, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

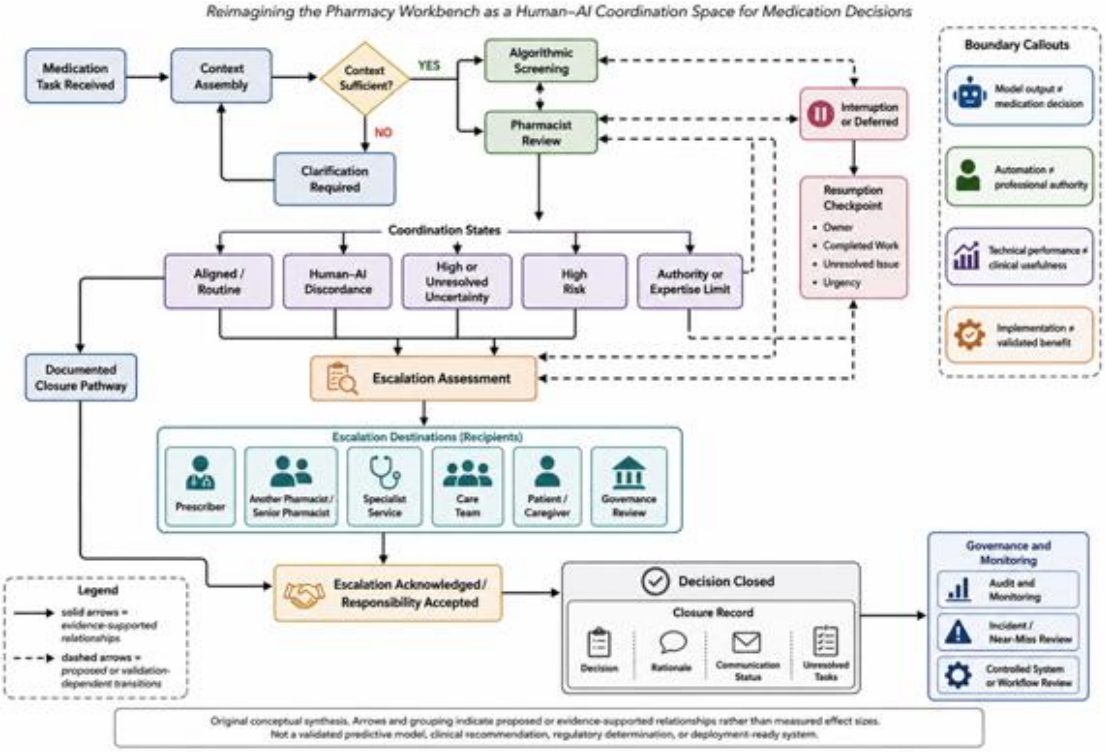

Figure 2. Coordination-State and Escalation Pathway for Medication Decisions.

The figure is an original conceptual synthesis developed for "Reimagining the Pharmacy Workbench as a Human–AI Coordination Space for Medication Decisions." Arrows and grouping indicate proposed or evidence-supported

relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

Boundary Conditions and Limitations

The architecture is vulnerable to inequity when target variables, proxy measures, or available data reproduce structural differences rather than the medication need the system is intended to address [29]. Performance may also vary across organizations because models can learn institution-specific associations that do not transfer to other populations or workflows [30].

Clinical translation remains conditional on evidence quality, generalizability, usability, workflow integration, governance, and continuing monitoring [31]. The architecture does not determine which medication tasks should be automated, what escalation thresholds are appropriate, or whether benefits outweigh burdens in a particular setting. Applicability is limited by interoperability, staffing, professional authority, data completeness, patient participation, and access to timely escalation. Its layers and states may require simplification or extension, and their empirical effects remain unknown.

Conclusion

The pharmacy workbench should be understood not only as a workstation but as a sociotechnical coordination space in which patient context, prescription information, algorithmic functions, professional judgment, communication, time, and organizational authority interact. The proposed architecture makes these relationships explicit through six connected layers and a coordination pathway that distinguishes routine alignment from missing context, discordance, uncertainty, interruption, deferral, escalation, and closure.

Its principal scientific contribution is to separate model output from medication decision, technical performance from clinical usefulness, and implementation from validated benefit. Its professional contribution is to represent pharmacist work as contextual interpretation, communication, responsibility management, and accountable disposition rather than residual approval after automation. Its organizational contribution is to connect individual decisions with monitoring, incident review, change control, and the capacity to restrict unsafe functions.

The architecture remains conceptual. It does not demonstrate improved medication safety, efficiency, equity, workload, or patient outcomes. Validation requires task-specific and setting-specific evidence across technical, clinical, human-factors, workflow, organizational, safety, equity, and implementation dimensions. The workbench should therefore be used as a framework for designing research, exposing hidden dependencies, and defining decision boundaries—not as a clinical protocol or deployment authorization.

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

References

1. Topol EJ. High-performance medicine: The convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
2. 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
3. Graafsma J, Murphy RM, van de Garde EMW, Karapinar-Çarkit F, Derijks HJ, Hoge RHL, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411-22. doi:10.1093/jamia/ocae076
4. 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
5. Van Dort BA, Zheng WY, Baysari MT. Prescriber perceptions of medication-related computerized decision support systems in hospitals: A synthesis of qualitative research. Int J Med Inform. 2019;129:285-95. doi:10.1016/j.ijmedinf.2019.06.024
6. Jeffries M, Salema NE, Laing L, Shamsuddin A, Sheikh A, Avery T, et al. Using sociotechnical theory to understand medication safety work in primary care and prescribers' use of clinical decision support: A qualitative study. BMJ Open. 2023;13(4):e068798. doi:10.1136/bmjopen-2022-068798
7. Farre A, Heath G, Shaw K, Bem D, Cummins C. How do stakeholders experience the adoption of electronic prescribing systems in hospitals? A systematic review and thematic synthesis of qualitative studies. BMJ Qual Saf. 2019;28(12):1021-31. doi:10.1136/bmjqs-2018-009082
8. Salwei ME, Hoonakker PLT, Carayon P, Hundt AS, Wiegmann D, Patterson BW, et al. 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
9. Sittig DF, Boxwala A, Wright A, Zott C, Desai P, Dhopeshwarkar R, et al. A lifecycle framework illustrates eight stages necessary for realizing the benefits of patient-centered clinical decision support. J Am Med Inform Assoc. 2023;30(9):1583-9. doi:10.1093/jamia/ocad122
10. Sittig DF, Boxwala A, Wright A, Zott C, Gauthreaux NA, Swiger J, et al. Patient-centered clinical decision support challenges and opportunities identified from workflow execution models. J Am Med Inform Assoc. 2024;31(8):1682-92. doi:10.1093/jamia/ocae138
11. Li RC, Wang JK, Sharp C, Chen JH. When order sets do not align with clinician workflow: Assessing practice patterns in the electronic health record. BMJ Qual Saf. 2019;28(12):987-96. doi:10.1136/bmjqs-2018-008968
12. Reddy A, Abebe E, Rivera AJ, Stone JA, Chui MA. Interruptions in community pharmacies: Frequency, sources, and mitigation strategies. Res Social Adm Pharm. 2019;15(10):1243-50. doi:10.1016/j.sapharm.2018.10.030
13. van der Vegt AH, Scott IA, Dermawan K, Schnetler RJ, Kalke VR, Lane PJ. Implementation frameworks for end-to-end clinical AI: Derivation of the SALIENT framework. J Am Med Inform Assoc. 2023;30(9):1503-15. doi:10.1093/jamia/ocad088
14. 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
15. 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):e745-50. doi:10.1016/S2589-7500(21)00208-9

16. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
17. 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
18. Kiani A, Uyumazturk B, Rajpurkar P, Wang A, Gao R, Jones E, et al. Impact of a deep learning assistant on the histopathologic classification of liver cancer. NPJ Digit Med. 2020;3:23. doi:10.1038/s41746-020-0232-8
19. Gaube S, Suresh H, Raue M, Merritt A, Berkowitz SJ, Lermer E, et al. Do as AI say: Susceptibility in deployment of clinical decision-aids. NPJ Digit Med. 2021;4(1):31. doi:10.1038/s41746-021-00385-9
20. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
21. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
22. 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
23. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385:e078378. doi:10.1136/bmj-2023-078378
24. 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):2730-39. doi:10.1093/jamia/ocae209
25. 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):e367. doi:10.2196/jmir.8775
26. Damschroder LJ, Reardon CM, Widerquist MAO, Lowery J. The updated Consolidated Framework for Implementation Research based on user feedback. Implement Sci. 2022;17(1):75. doi:10.1186/s13012-022-01245-0
27. Baines D, Nørgaard LS, Babar ZUD, Rossing C. The Fourth Industrial Revolution: Will it change pharmacy practice? Res Social Adm Pharm. 2020;16(9):1279-81. doi:10.1016/j.sapharm.2019.04.003
28. Baines DL, Bates I, Bader L, Hale C, Schneider P. Conceptualising production, productivity and technology in pharmacy practice: A novel framework for policy, education and research. Hum Resour Health. 2018;16(1):51. doi:10.1186/s12960-018-0317-5
29. 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
30. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. PLoS Med. 2018;15(11):e1002683. doi:10.1371/journal.pmed.1002683
31. 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