Original Article

A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation

Erik Van der Meer[1]*, Lieke Jansen[1], Bram de Vries[2], Sophie Claes[3]

[1]Department of Medication Intelligence and Command Systems, Faculty of Pharmacy, Wageningen University, Wageningen, Netherlands.
[2]Department of Risk-Workload Coordination in Pharmacy, Faculty of Pharmacy, University of Groningen, Groningen, Netherlands.
[3]Department of AI Escalation and Clinical Oversight, Faculty of Pharmacy, University of Amsterdam, Amsterdam, Netherlands.

Abstract

Medication-related risk is produced across prescribing, verification, dispensing, administration, monitoring, transitions of care, and patient communication, whereas responsibility for recognizing and responding to that risk is distributed across professions, technologies, locations, and time. Existing dashboards, alerts, prediction models, prioritization tools, and operational command centres address parts of this problem but do not provide a unified mechanism for coordinating medication risk, patient urgency, workload, assignment, and escalation. This article proposes a Medication Intelligence Command-Centre Architecture as an original non-empirical sociotechnical synthesis. The architecture separates data provenance, medication-risk signals, patient urgency, uncertainty, workload and capacity, queue state, orchestration, professional adjudication, escalation, closure, and lifecycle governance. It treats model outputs as inputs to bounded coordination rather than medication decisions and preserves professional authority for clinical interpretation and action. A multidimensional orchestration process is proposed to determine whether work should be assigned, checked, monitored, deferred, escalated, or redirected through a fallback pathway. Evaluation is organized across data quality, technical validity, medication-task performance, safety, throughput, work redistribution, human–AI interaction, equity, implementation, and temporal stability. Advancement toward operational use would require local validation, prospective workflow evaluation, subgroup analysis, incident monitoring, change control, and evidence that apparent efficiency does not conceal missed risk or displaced work. The proposed architecture does not establish clinical effectiveness, regulatory acceptance, universal applicability, or deployment readiness. Its original contribution is to frame medication intelligence as a governed coordination function in which risk, urgency, workload, and authority remain distinguishable but operationally connected.

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

Introduction

Preventable harm occurs across health-care settings, while medication-related harm constitutes a substantial and cross-setting component of that burden [1, 2]. Medication risk is therefore not confined to a single prescription, alert, professional group, or information system. It can emerge from incomplete histories, prescribing choices, interactions, monitoring delays, dispensing processes, administration conditions, transitions, patient circumstances, and failures to act on previously recognized problems. The coordination problem is consequently temporal and organizational as well as clinical.

Health-care command centres have been developed to improve situational awareness, patient flow, capacity management, and operational coordination, but their purposes, structures, technologies, staffing arrangements, and evaluation methods remain heterogeneous [3]. Existing evidence does not define a medication-specific command-centre model capable of integrating clinical risk with workload and escalation. A dashboard displaying more

Address for correspondence: Erik Van der Meer, Department of Medication Intelligence and Command Systems, Faculty of Pharmacy, Wageningen University, Wageningen, Netherlands.
e.vandermeer@wur.nl
Received: 15 March 2026; **Accepted:** 07 July 2026

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

How to cite this article: Van der Meer E, Jansen L, de Vries B, Claes S. A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation. Arch Pharm Pract. 2026;17(3):1-9. https://doi.org/10.51847/z6mnNPLfLO

information would not by itself resolve that gap because visibility does not determine which signal is actionable, who should respond, how rapidly action is required, or what should happen when the designated service has insufficient capacity.

Medication-related clinical decision support illustrates this distinction. Increasing signal sensitivity may expose additional hazards while simultaneously increasing alert burden, desensitization, and inappropriate overrides [4]. The central design problem is therefore not maximal detection. It is the governed conversion of heterogeneous signals into proportionate, accountable, and executable work.

Pharmacy prioritization approaches similarly differ in their variables, scoring assumptions, intended users, and validation status [5]. A high medication-risk estimate may require immediate pharmacist review in one context, a laboratory check in another, or monitoring rather than interruption when evidence is incomplete. Patient deterioration may also increase urgency without changing the medication-specific nature of the concern. This article therefore proposes a Medication Intelligence Command Centre as a bounded sociotechnical coordination architecture. It is intended to connect risk sensing, uncertainty, patient urgency, workload, queues, assignment, escalation, professional judgment, and monitoring without treating any one model or metric as sufficient.

Command-Centre Functions and Decision Boundaries

A command centre is defined here by its coordination functions rather than by a physical room or visual display. Existing implementations combine systems engineering, operational data, visual management, designated roles, and structured response processes, while institutional configurations vary considerably [6, 7]. For medication intelligence, the proposed functions are: sensing medication-related conditions; checking data provenance and sufficiency; representing risk, urgency, workload, and queue state; routing work; escalating unresolved or time-sensitive concerns; recording acknowledgement and closure; and monitoring performance and unintended consequences.

These functions must be separated from medication authority. The command centre may identify, prioritize, route, request checking, recommend escalation, or expose an unresolved conflict. It should not independently prescribe, discontinue, verify, dispense, administer, or clinically resolve a patient-specific medication problem. Those acts remain with appropriately authorized professionals. The distinction is important because command-centre usefulness depends on data quality, operational interpretation, and integration with real work rather than the mere availability of consolidated information [8].

Decision boundaries should also constrain claims about safety. An observed association between a command centre and an operational or safety outcome does not establish that the architecture caused the outcome or that the result will transfer to medication management [9]. Alternative explanations may include concurrent staffing changes, revised documentation, altered escalation practices, increased managerial attention, or changes in the underlying patient population. The proposed command centre must therefore expose the origin of each signal, the basis for prioritization, the responsible actor, the permitted action, the fallback pathway, and the conditions requiring reassessment.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Table 1. Definitions, components, relationships, and intended uses for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Medication-risk sensing	Relevant hazards are distributed across records and workflows.	Orders, medication histories, allergies, laboratory data, administration and dispensing events, patient reports.	**Proposed:** combine traceable signals without assuming that signal quantity equals clinical importance.	Broader situational awareness.	Source validity, completeness, timeliness, false-positive and false-negative analysis. Medication-alert trade-offs are established [4].	Missing data, duplication, stale information, excessive alert volume.	A detected signal is not a medication decision.
Clinical prioritization	Pharmacy demand may exceed immediately available review capacity.	Drug-related-problem indicators, complexity, vulnerability, transition status, time sensitivity.	**Proposed:** use transparent dimensions to order review while preserving professional override.	More explicit queue ordering.	Local construct validity, inter-rater comparison, missed-risk analysis. Existing prioritization approaches are variable [5].	Hidden weighting, unstable rankings, systematic deprioritization.	Priority does not establish safety or eliminate review obligations.

Command-centre coordination	Information is visible but not consistently converted into accountable work.	Validated signals, operational roles, service availability, communication pathways.	**Proposed:** connect sensing, routing, acknowledgement, escalation, and closure. Command-centre functions are evidence-derived [6]; configuration variability is evidence-derived [7].	Closed-loop coordination.	Workflow observation, response-time analysis, failure-mode testing.	Dashboard use without ownership; ambiguous handoffs.	Coordination does not confer clinical authority.
Data provenance and quality	Decisions may be based on incomplete, delayed, or context-poor data.	Source, timestamp, transformation history, missingness, reconciliation status.	**Proposed:** attach provenance and sufficiency status to every actionable signal. Data-quality dependence is evidence-supported [8].	Better interpretation and auditable abstention.	Data-quality audits and comparison with authoritative sources.	False precision and unrecognized data latency.	Insufficient data should trigger checking or abstention, not fabricated certainty.
Workload, capacity, and queue state	Assignment may be nominally correct but operationally impossible.	Available staff, competencies, task volume, queue age, interruptions, deadlines.	**Proposed:** modify routing and escalation while preventing capacity constraints from erasing recognized risk.	Executable rather than merely theoretical assignments.	Prospective workload and redistribution measurement.	Risk transfer to saturated teams; invisible backlog.	Capacity may change timing or destination, not redefine a hazard as harmless.
Professional adjudication	Algorithmic or operational outputs may lack contextual meaning.	Patient-specific context, professional assessment, uncertainty, preferences, competing priorities.	**Proposed:** require authorized human interpretation before medication action.	Preservation of clinical accountability and contextual judgment.	Human-factors, concordance, override, and incident analysis.	Automation bias, over-reliance, unexplained override.	Prescribing, verification, and treatment decisions remain human-authorized acts.
Safety monitoring and closure	Completed routing may be mistaken for completed risk resolution.	Acknowledgement, action, rationale, unresolved status, outcome, incident reports.	**Proposed:** require explicit closure, re-entry, and review pathways. Safety effects remain context-dependent [9].	Auditable follow-through and learning.	Prospective safety surveillance and causal evaluation.	Premature closure, lost follow-up, misleading attribution.	Closure records workflow completion; it does not prove patient benefit.

Proposed Medication Intelligence Architecture

The proposed architecture contains six connected but non-interchangeable layers. The data and provenance layer receives clinical, medication-use, patient-reported, and operational information and records source, freshness, missingness, and reconciliation status. The analytic-services layer may contain rules, statistical models, machine-learning systems, or structured professional assessments. Clinical decision-support literature supports separating information inputs, knowledge services, presentation, workflow integration, and feedback rather than treating software output as a complete intervention [10].

The multidimensional state layer represents medication risk, patient urgency, uncertainty, workload, capacity, and queue conditions without forcing them into one universal score. The orchestration layer applies a proposed permission envelope defining eligible tasks, users, populations, information inputs, outputs, escalation limits, fallback arrangements, and monitoring obligations. The human action layer assigns work to authorized pharmacists or other professionals, captures acknowledgement and adjudication, and prevents routing logic from becoming medication authority. The governance and learning layer maintains ownership, documentation, auditability, incident review, model change control, and requalification. These safeguards reflect evidence that responsible clinical AI requires governance, lifecycle controls, workflow integration, and validation beyond technical model performance [11-13].

Figure 1 presents the original conceptual synthesis for medication intelligence command-centre architecture, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

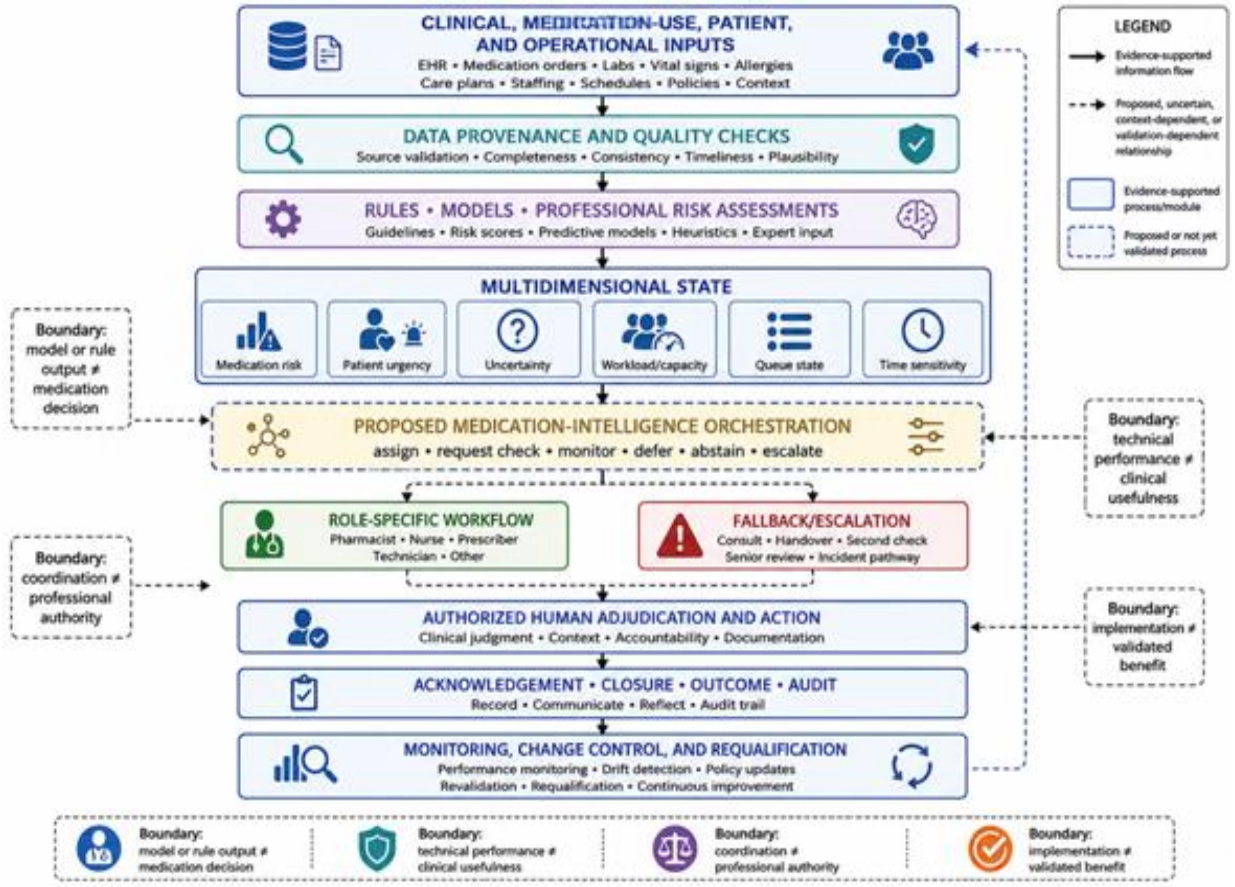

Figure 1. Medication Intelligence Command-Centre Architecture

Information should move through the architecture only when minimum provenance and sufficiency requirements are met. Where these conditions fail, the appropriate output may be a request for reconciliation, manual checking, abstention, or escalation rather than a calculated priority. The architecture therefore treats uncertainty as operational information. It also requires every assignment to contain an intended recipient, permitted action, expected response condition, escalation route, and closure requirement.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Data and provenance	Establish whether inputs are sufficiently complete, current, and interpretable.	Source data → provenance and quality status.	Conditions all downstream processing. Data-quality dependence is evidence-supported [8].	Completeness, timeliness, traceability, reconciliation accuracy.	Stale, missing, duplicated, or incorrectly linked data.	Data steward and clinical source owner.	No actionable routing from unverified critical inputs without fallback review.
Analytic services	Generate rules, risk estimates, classifications, or uncertainty information.	Validated inputs → model or rule outputs.	Supplies signals to the multidimensional state layer. CDS layering is evidence-supported [10].	Discrimination, calibration, error distribution, subgroup performance, reproducibility.	Model error, opaque transformation, false confidence.	Model owner with clinical and pharmacy oversight.	Technical validity does not authorize clinical use.

Multidimensional state	Keep risk, urgency, uncertainty, workload, capacity, and queue status distinguishable.	Analytic and operational signals → structured state representation.	Prevents one dimension from silently substituting for another.	Construct validity, interpretability, stability, missing-data behavior.	Invalid aggregation or hidden weighting.	Multidisciplinary design authority.	Combined weighting remains **proposed** until locally validated.
Orchestration	Select assignment, checking, monitoring, deferment, abstention, or escalation pathway.	State representation → bounded workflow recommendation.	Applies permission envelope and fallback rules. Governance requirements are evidence-supported [11].	Routing accuracy, missed-risk patterns, response feasibility, auditability.	Unsafe suppression, circular routing, assignment to unavailable services.	Pharmacy operations and clinical safety owners.	Orchestration output remains advisory until authorized acceptance.
Human adjudication	Interpret context and undertake authorized medication action.	Recommendation → acknowledgement, judgment, action, or override.	Feeds rationale and outcome information back to monitoring. Lifecycle safeguards are evidence-supported [12].	Usability, comprehension, reliance, override quality, response latency.	Automation bias, ignored signals, unrecorded workarounds.	Authorized professional and supervising service.	No autonomous prescribing, verification, or treatment modification.
Monitoring and requalification	Detect safety events, drift, workload displacement, and unintended effects.	Actions and outcomes → audit, review, change control.	May restrict, revise, suspend, or retire components. Translation challenges are evidence-supported [13].	Safety surveillance, workflow consequences, subgroup effects, change impact.	Monitoring blind spots, unreviewed drift, normalization of workarounds.	Governance committee, safety lead, model steward, human-factors lead.	Implementation cannot be equated with validated benefit or continuing fitness.

Integrating Risk, Capacity, Queues, and Patient Urgency

Medication risk, patient urgency, workload, capacity, and queue state should be integrated operationally but not treated as synonymous. Medication risk concerns the probability and potential consequence of a medication-related problem. A hybrid machine-learning and rule-based system has shown that prescriptions may be prioritized for medication review, but any rule-out or deprioritization function must be evaluated for missed drug-related problems [14]. In the proposed architecture, low estimated risk may alter queue position only within a defined permission envelope; it cannot certify that review is unnecessary.

Prioritization converts relevant dimensions into an ordered or categorized worklist. Existing outpatient pharmaceutical prioritization tools vary in inputs, scoring methods, intended uses, and validation approaches [15]. This heterogeneity argues against a universal score. The command centre should instead preserve the component dimensions and allow locally justified rules to determine how they interact. A complex regimen, recent transition, high-risk medicine, incomplete history, or unresolved laboratory abnormality may each contribute differently depending on the service and patient population.

Patient urgency concerns the time available before delayed assessment may become clinically unacceptable. It is not equivalent to medication risk. Methodological limitations in early-warning scores demonstrate that urgency tools require clearly defined outcomes, calibration assessment, appropriate thresholds, and external validation [16]. A deterioration signal may accelerate medication review or change the receiving team, but it should not automatically determine the content of a medication decision. Conversely, a medication problem can require urgent action even when a general deterioration score is low.

Workload and capacity describe whether the designated service can execute the required response within the necessary period. Staffing evidence in other clinical workforces indicates that capacity is associated with patient outcomes, supporting its treatment as a safety-relevant operational constraint rather than an administrative afterthought [17]. However, this evidence does not establish a pharmacy staffing ratio or universal workload threshold. The proposed system should expose capacity limitations, queue age, interruptions, competency requirements, and unacknowledged tasks. When the preferred assignee is saturated, the appropriate response may be redistribution, supervisory escalation, a narrower interim safety check, or activation of a fallback service.

The integration mechanism is therefore a proposed multidimensional orchestration gate. Risk indicates what may go wrong; urgency indicates how quickly attention may be required; uncertainty indicates whether the information is adequate; capacity indicates which response is feasible; and queue state indicates whether delay is accumulating. None of these dimensions should erase another. A capacity shortage must not relabel a recognized hazard as low risk, while a high model score must not create professional authority. The resulting assignment remains conditional on acknowledgement, contextual adjudication, documented

action, and closed-loop follow-up. This arrangement is conceptually plausible but requires local simulation, prospective evaluation, subgroup analysis, and testing under both normal and capacity-constrained conditions.

Prioritization, Assignment, and Escalation Workflows

Prioritization should determine the sequence and intensity of attention without implying that lower-ranked work is harmless. Presentation and routing should be tailored to professional role, because medication alerts differ in relevance, actionability, and interruption cost across pharmacists, technicians, prescribers, nurses, and operational managers [18]. The proposed command centre therefore converts each accepted signal into a bounded task containing its rationale, urgency, uncertainty, intended recipient, permitted response, acknowledgement requirement, and escalation route.

Filtering may reduce low-value work only when the acceptable false-negative boundary is specified and monitored [19]. Suppression should not be treated as an invisible technical function. The system should record what was filtered, why it was filtered, which population and conditions were covered, and how suppressed signals can re-enter review after new information, delay, deterioration, or model uncertainty. High uncertainty should favor checking or abstention rather than confident deprioritization.

Prescription-risk models may help order pharmacist worklists, but their outputs remain triage signals rather than medication decisions [20]. Assignment should consider professional competence, service remit, current workload, continuity, and required response time. Capacity may change the recipient or trigger escalation, but it should not redefine a recognized medication hazard as clinically unimportant.

Escalation is proposed as a closed-loop sequence: trigger, notification, acknowledgement, responder assignment, interim risk control, professional adjudication, handoff where required, closure, and subsequent review. Comparable rapid-response systems demonstrate the importance of defined triggers and response structures, although they do not validate this medication-specific adaptation [21]. Unacknowledged, overdue, rejected, or unresolved tasks should therefore return to an escalation gate rather than disappear from the queue.

Evaluation of Safety, Throughput, and Work Redistribution

Evaluation should begin by specifying the architecture's intended use, users, information requirements, human–AI interaction, fallback procedures, and foreseeable errors. These elements are necessary because reporting technical accuracy without describing how outputs reach and influence professionals cannot establish the performance of the complete intervention [22].

Early live evaluation should examine usability, workflow integration, response behavior, safety events, overrides, workarounds, and failure recovery before comparative effectiveness claims are attempted [23]. A staged pathway is proposed: technical verification; retrospective local validation; silent workflow testing; limited prospective use with close supervision; comparative evaluation; and continuing postimplementation monitoring. Progression should be conditional rather than automatic.

Outcome assessment should distinguish medication-task performance from downstream workflow consequences. Clinical decision-support effects vary across interventions and contexts, making task-specific comparators and absolute measures preferable to assumptions of general benefit [24]. Relevant domains include missed high-risk cases, inappropriate escalation, queue age, acknowledgement delay, time to professional review, task completion, interruption burden, duplicate work, cross-professional transfer, override patterns, incident reports, and patient-relevant consequences.

Independent local validation is essential because widely implemented prediction systems may perform differently outside their development setting [25]. Thresholds, calibration, subgroup behavior, staffing conditions, and workflow effects should therefore be re-examined before expanding the system's population, task, autonomy, or institutional scope. No single performance measure should compensate for unacceptable safety, equity, or human-factors failure.

Organizational and Human-Factors Implications

Implementation depends on infrastructure, workflow fit, stakeholder participation, professional acceptance, organizational resources, and credible evidence [26]. A command centre may otherwise create a technically functional but operationally unusable coordination layer. Its introduction should therefore include workflow observation, co-design, competency assessment, simulation, fallback preparation, and explicit protection against transferring unmanageable work to another professional group.

Figure 2 presents the original conceptual synthesis for risk, workload, patient urgency, assignment, and escalation workflow, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

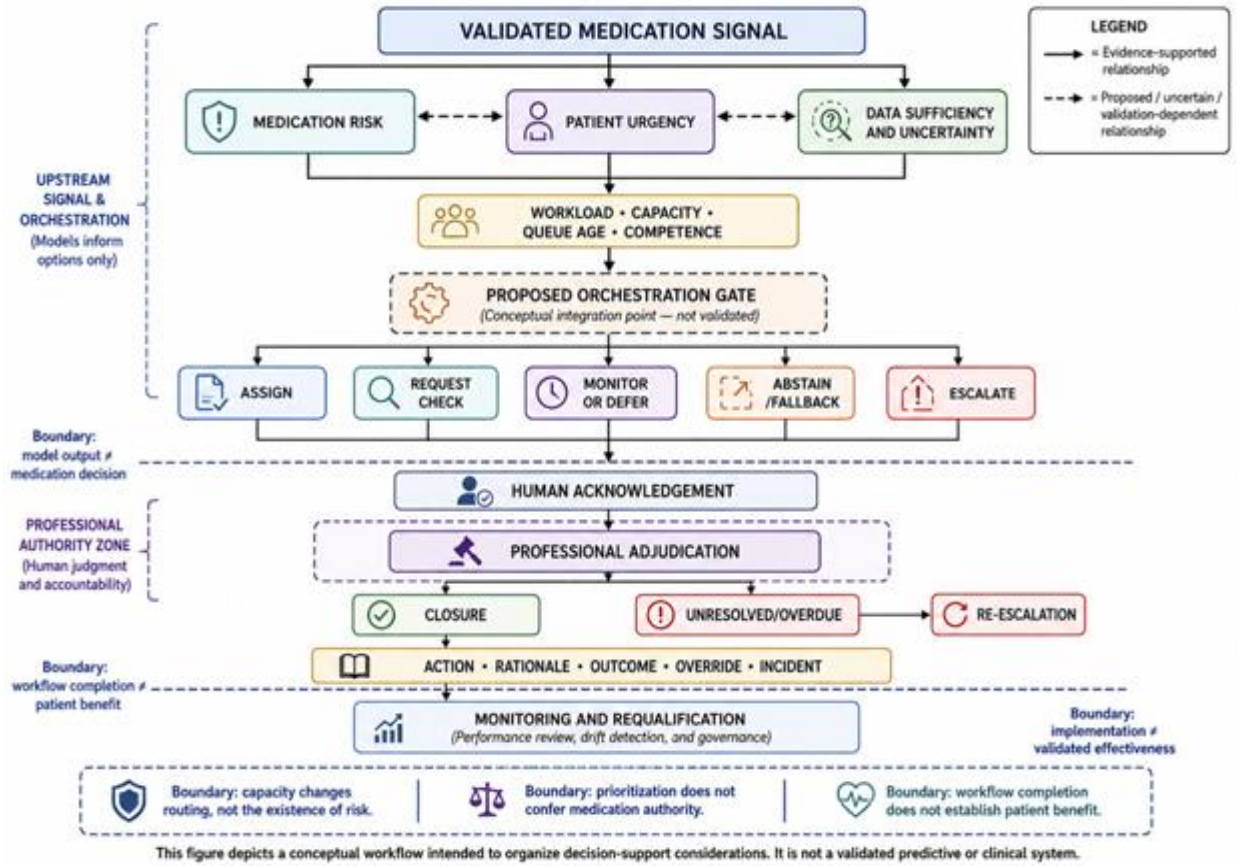

Figure 2. Risk, Workload, Patient Urgency, Assignment, and Escalation Workflow

Users also require accessible information about each analytic component's intended use, population, performance context, limitations, ownership, fallback, and maintenance history [27]. Such model facts may support calibrated reliance but cannot replace training, professional judgment, or empirical usability testing.

Work redistribution must be treated as an outcome. Digital systems can create stress through interruptions, documentation requirements, inbox burden, and inefficient interaction design [28]. Evaluation should therefore measure not only work completed by the command centre but also new coordination tasks, duplicated review, cognitive load, interruptions, and burden transferred to peripheral services.

Monitoring should actively seek unintended consequences such as automation bias, deskilling, gaming, normalization of workarounds, and organizational dependence [29]. **Table 3** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article A Medication Intelligence Command Centre for Coordinating Risk, Workload, and Escalation.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary
Clinical and medication authority	Authorized pharmacist and relevant prescriber or clinical team	Interpret and act on patient-specific medication concerns.	Assessment, decision, rationale, unresolved uncertainty.	Overrides, delayed review, disagreement, patient-safety events.	Serious risk, unresolved conflict, or action outside service remit.	Senior clinical review or transfer to the accountable team.	The command centre coordinates but does not prescribe, verify, or determine treatment.
Pharmacy operations	Pharmacy service lead	Define staffing, competencies, queues, fallback, and escalation coverage.	Service rules, coverage map, competency requirements.	Backlog, queue age, unacknowledged work, workload transfer.	Capacity threatens required response.	Redistribution, narrowed interim control, or managerial escalation.	Operational ownership does not replace individual professional judgment.

Data and model stewardship	Data steward and model owner	Maintain provenance, technical documentation, validation, and change control.	Data lineage, model facts, versions, validation and incident records [27].	Missingness, drift, calibration, technical failure.	Material data, model, population, or workflow change.	Restrict, recalibrate, revalidate, suspend, or retire.	Technical ownership does not confer clinical authority.
Human factors and implementation	Human-factors lead with frontline users	Assess usability, interaction design, training, and workflow fit.	Workflow maps, usability findings, training and workaround logs.	Interruption burden, task abandonment, cognitive load, user confusion [28].	Unsafe workaround or sustained burden.	Redesign interface, task routing, staffing, or training.	User acceptance alone does not establish safety.
Medication safety and incident response	Medication-safety lead	Investigate suspected harm and coordinate corrective action.	Incident analysis, contributing factors, actions, follow-up.	Missed risk, inappropriate escalation, duplicate or delayed action.	Serious event or recurring failure pattern.	Immediate containment and formal review.	Association with system use does not by itself establish causation.
Equity and patient impact	Equity lead with patient representation	Examine subgroup allocation, access, and consequences.	Subgroup analyses, exclusions, accessibility and complaint records.	Differential prioritization, delay, error, or service access.	Material unexplained disparity.	Restrict use, revise proxies, improve data, or redesign pathways.	Aggregate performance cannot establish equitable operation.
Lifecycle governance	Multidisciplinary governance body	Approve scope, monitor use, and restrict, suspend, or retire components.	Permission envelope, review decisions, changes, exceptions.	Safety, effectiveness, workload, equity, drift, unintended consequences [29].	Boundary violation or evidence no longer supports use.	Requalification, scope reduction, suspension, or retirement.	Governance authorization is local and does not imply regulatory approval or general readiness.

Limitations

Risk proxies can reproduce inequity when they reflect utilization or cost rather than patient need [30]. Subgroup allocation, missingness, access, and service availability must therefore be tested directly. Operational, clinical, and documentation changes may also shift the data-generating process and invalidate model or queue behavior [31].

Uncertainty cannot be eliminated and should instead support checking, abstention, delay, or escalation when appropriate [32]. The architecture may additionally create new burdens, dependencies, or failure pathways that are not visible during retrospective evaluation. Finally, validity cannot be presumed across institutions, populations, pharmacy services, staffing models, or information infrastructures [33]. The framework is a proposed conceptual architecture and does not establish effectiveness, universal applicability, legal responsibility, or deployment readiness.

Conclusion

The Medication Intelligence Command-Centre Architecture reframes medication intelligence as a governed coordination function rather than a prediction model, alert platform, dashboard, or substitute for professional judgment. Its central contribution is the explicit integration of medication risk, patient urgency, uncertainty, workload, capacity, queue state, assignment, escalation, closure, and lifecycle governance while preserving boundaries among these constructs.

The architecture proposes that information should become work only through traceable provenance checks, bounded orchestration, role-appropriate assignment, human acknowledgement, authorized adjudication, and closed-loop follow-up. It also requires evaluation to extend beyond model accuracy to safety, throughput, workload redistribution, equity, usability, implementation, temporal stability, and unintended consequences.

This contribution remains conceptual. Its scientific value depends on whether its constructs can be operationalized, tested, challenged, and revised in diverse pharmacy environments. Prospective validation must determine whether the architecture improves coordination without concealing missed risk, increasing inequity, shifting unsustainable work, or weakening professional accountability. Until such evidence exists, the proposed command centre should be understood as an evaluable architecture for organizing medication-risk coordination, not as a clinical recommendation, regulatory determination, universal standard, or deployment-ready system.

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

References

1. Panagioti M, Khan K, Keers RN, Abuzour A, Phipps D, Kontopantelis E, et al. Prevalence, severity, and nature of preventable patient harm across medical care settings: Systematic review and meta-analysis. BMJ. 2019;366. doi:10.1136/bmj.l4185
2. Hodkinson A, Tyler N, Ashcroft DM, Keers RN, Khan K, Phipps D, et al. Preventable medication harm across health care settings: A systematic review and meta-analysis. BMC Med. 2020;18(1):313. doi:10.1186/s12916-020-01774-9
3. Franklin BJ, Mueller SK, Bates DW, Gandhi TK, Morris CA, Goralnick E. Use of hospital capacity command centers to improve

patient flow and safety: A scoping review. J Patient Saf. 2022;18(6). doi:10.1097/PTS.0000000000000976
4. 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
5. Falconer N, Barras M, Cottrell N. Prioritising patients for hospital pharmacy services: Current approaches and the future of clinical prioritisation tools. Res Social Adm Pharm. 2019;15(10):1266-73. doi:10.1016/j.sapharm.2018.11.003
6. Kane EM, Scheulen JJ, Püttgen A, Martinez D, Levin S, Bush BA, et al. Use of systems engineering to design a hospital command center. Jt Comm J Qual Patient Saf. 2019;45(5):370-9. doi:10.1016/j.jcjq.2018.11.006
7. Franklin BJ, Yenduri R, Parekh VI, Fogerty RL, Scheulen JJ, High H, et al. Hospital capacity command centers: A benchmarking survey on an emerging mechanism to manage patient flow. Jt Comm J Qual Patient Saf. 2023;49(4):189-98. doi:10.1016/j.jcjq.2023.01.007
8. Mebrahtu TF, McInerney CD, Benn J, McCrorie C, Granger J, Lawton T, et al. The impact of hospital command centre on patient flow and data quality: Findings from the UK National Health Service. Int J Qual Health Care. 2023;35(4). doi:10.1093/intqhc/mzad072
9. Mebrahtu TF, McInerney CD, Benn J, McCrorie C, Granger J, Lawton T, et al. Effect of a hospital command centre on patient safety: An interrupted time series study. BMJ Health Care Inform. 2023;30(1). doi:10.1136/bmjhci-2022-100653
10. 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
11. Reddy S, Allan S, Coghlan S, Cooper P. A governance model for the application of AI in health care. J Am Med Inform Assoc. 2020;27(3):491-7. doi:10.1093/jamia/ocz192
12. 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
13. 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
14. Levivien C, Cavagna P, Grah A, Buronfosse A, Courseau R, Bézie Y, et al. Assessment of a hybrid decision support system using machine learning with artificial intelligence to safely rule out prescriptions from medication review in daily practice. Int J Clin Pharm. 2022;44(2):459-65. doi:10.1007/s11096-021-01366-4
15. Sousa MCVB, Fernandes BD, Foppa AA, Almeida PHRF, Mendonça SAM, Chemello C. Tools to prioritize outpatients for pharmaceutical service: A scoping review. Res Social Adm Pharm. 2020;16(12):1645-57. doi:10.1016/j.sapharm.2020.02.010
16. Gerry S, Bonnici T, Birks J, Kirtley S, Virdee PS, Watkinson PJ, et al. Early warning scores for detecting deterioration in adult hospital patients: Systematic review and critical appraisal of methodology. BMJ. 2020;369. doi:10.1136/bmj.m1501
17. McHugh MD, Aiken LH, Sloane DM, Windsor C, Douglas C, Yates P. Effects of nurse-to-patient ratio legislation on nurse staffing and patient mortality, readmissions, and length of stay: A prospective study in a panel of hospitals. Lancet. 2021;397(10288):1905-13. doi:10.1016/S0140-6736(21)00768-6
18. 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
19. Liu S, Kawamoto K, Del Fiol G, Weir CR, Malone DC, Reese TJ, et al. The potential for leveraging machine learning to filter medication alerts. J Am Med Inform Assoc. 2022;29(5):891-9. doi:10.1093/jamia/ocab292
20. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O, et al. A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. J Am Med Inform Assoc. 2020;27(11):1688-94. doi:10.1093/jamia/ocaa154
21. Lyons PG, Edelson DP, Churpek MM. Rapid response systems. Resuscitation. 2018;128:191-7. doi:10.1016/j.resuscitation.2018.05.013
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. 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. Kwan JL, Lo L, Ferguson J, Goldberg H, Diaz-Martinez JP, Tomlinson G, et al. Computerised clinical decision support systems and absolute improvements in care: Meta-analysis of controlled clinical trials. BMJ. 2020;370. doi:10.1136/bmj.m3216
25. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021;181(8):1065-70. doi:10.1001/jamainternmed.2021.2626
26. Strohm L, Hehakaya C, Ranschaert ER, Boon WPC, Moors EHM. Implementation of artificial intelligence applications in radiology: Hindering and facilitating factors. Eur Radiol. 2020;30(10):5525-32. doi:10.1007/s00330-020-06946-y
27. Sendak MP, Gao M, Brajer N, Balu S. Presenting machine learning model information to clinical end users with model facts labels. NPJ Digit Med. 2020;3:41. doi:10.1038/s41746-020-0253-3
28. Kroth PJ, Morioka-Douglas N, Veres S, Babbott S, Poplau S, Qeadan F, et al. Association of electronic health record design and use factors with clinician stress and burnout. JAMA Netw Open. 2019;2(8). doi:10.1001/jamanetworkopen.2019.9609
29. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA. 2017;318(6):517-8. doi:10.1001/jama.2017.7797
30. 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
31. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
32. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4:4. doi:10.1038/s41746-020-00367-3
33. Futoma J, Simons M, Panch T, Doshi-Velez F, Celi LA. The myth of generalisability in clinical research and machine learning in health care. Lancet Digit Health. 2020;2(9). doi:10.1016/S2589-7500(20)30186-2