Original Article

# The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care

**Kenji Tanaka[1]*, Yui Kobayashi[1], Takeshi Nakamura[2]**

[1]Department of Pharmacy Practice and Digital Ethics, Faculty of Pharmaceutical Sciences, University of Tokyo, Tokyo, Japan. [2]Department of Clinical Judgment and AI Accountability, Faculty of Pharmacy, Kyoto University, Kyoto, Japan.

## Abstract

Digital pharmacy systems increasingly make medication tasks visible as data fields, alerts, rankings, recommendations, and documented actions. Yet these representations can conceal the work required to make digital outputs clinically meaningful, including reconstructing incomplete context, repairing medication information, interpreting uncertainty, negotiating feasible care, and retaining responsibility for decisions and follow-up. This conceptual practice article proposes the Hidden Work Preservation Framework to explain how artificial intelligence may support, remove, shift, create, delay, or obscure such work. The framework distinguishes visible digital task completion from five interdependent practice domains: context-recovery, information-repair, interpretive, relational, and responsibility-bearing work. It further proposes a preservation loop through which patient-specific context is recovered, information is repaired, outputs are interpreted with patients and care teams, decisions are made or escalated, ownership and rationale are documented, and consequences are monitored. Evaluation should therefore extend beyond model discrimination or task time to context fidelity, repair burden, calibrated reliance, relational continuity, responsibility traceability, medication safety, equity, workflow redistribution, and organizational sustainability. These domains require direct observation, workflow mapping, record and incident review, patient and staff inquiry, subgroup analysis, and prospective evaluation in the intended setting. The framework is not a validated clinical algorithm, staffing model, regulatory standard, or autonomous decision pathway. Its original contribution is to make preservation of context, judgment, and responsibility an explicit design and evaluation problem for AI-supported pharmacy care.

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

## Introduction

### Digital Systems Reveal Tasks but Conceal Work

Digital systems render medication care as structured inputs, computable risks, alerts, recommendations, overrides, and completed transactions. Artificial intelligence may extend analytic capability, but predictive performance does not by itself show whether care has preserved the patient's circumstances, professional judgment, or accountable decision closure [1]. A technically successful output can therefore coexist with clinically important work that remains unrecorded, displaced, or attributed to no one.

Clinical-pharmacy AI has been applied to discrete activities such as prescription review, risk identification, prioritization, and decision support. However, the available pharmacy-specific literature remains concentrated in bounded applications and provides less certainty about real-world service benefit, workflow consequence, and the organization of professional responsibility [2].

The central conceptual problem is not whether algorithms can perform useful subtasks, but what must happen around those subtasks for medication care to remain interpretable, relational, safe, and accountable.

Health information technologies can reorganize coordination, gatekeeping, checking, and exception handling, while

**Address for correspondence:** Kenji Tanaka, Department of Pharmacy Practice and Digital Ethics, Faculty of Pharmaceutical Sciences, University of Tokyo, Tokyo, Japan. tanaka.kenji@eng.u-tokyo.ac.jp



**How to cite this article:** Tanaka K, Kobayashi Y, Nakamura T. The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care. Arch Pharm Pract. 2025;16(1):43-9. https://doi.org/10.51847/93LGPx91jn

medication alerts can transfer filtering and prioritization work to clinicians [3, 4].

This article calls that incompletely represented activity hidden work. "Hidden" does not mean secret, unnecessary, or inherently burdensome; it means that the activity is clinically consequential yet poorly captured by task lists, throughput measures, interface logs, or automation claims. The proposed framework treats digital pharmacy as a sociotechnical practice in which outputs acquire meaning only through interaction among data, patients, pharmacists, other professionals, local workflows, and governance arrangements. It does not assume that preserving all existing work is desirable. Instead, it asks which work protects context, judgment, relationships, and responsibility; which work reflects preventable system defects; where work is redistributed; and what evidence would be needed before claiming benefit, safety, equity, or readiness.

## Defining Hidden Work in Pharmacy Practice

Hidden work is proposed here as necessary activity that makes medication data, digital outputs, and formally specified tasks clinically usable but is incompletely represented in routine measures of work. Digital clinical activity commonly extends beyond visible encounters into documentation, inbox management, reconciliation, and after-hours completion [5]. The definition is functional rather than moral: some hidden work protects patients, some compensates for poor design, and some may be redundant, inequitable, or unsafe.

Medication management is also networked rather than confined to a single professional action. Pharmacy practice depends on distributed communication, boundary spanning, clarification, negotiation, and handoff across patients, prescribers, nurses, technicians, caregivers, and organizations [6]. A completed prescription check may therefore conceal the calls, messages, record searches, interpretive exchanges, and follow-up arrangements that made completion defensible.

The framework separates five proposed domains. Context-recovery work reconstructs medication history, chronology, goals, setting, and constraints. Information-repair work resolves omissions, contradictions, stale entries, terminology mismatches, and unclear provenance. Interpretive work converts evidence or model output into patient-specific meaning; relational work elicits goals and negotiates feasible plans; responsibility-bearing work assigns ownership, escalation, documentation, monitoring, and closure. Medication decision-support value is conditional on workflow integration, current information, and usable interaction [7]. When those conditions are poor, interaction friction may become professional burden rather than a minor interface inconvenience [8]. These domains are analytic distinctions, not empirically validated categories, and they may overlap in practice.

**Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care.

**Table 1.** Definitions, components, relationships, and intended uses for the article The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care.

| Component or scientific dimension | Problem addressed | Inputs or determinants | Proposed mechanism or relationship | Expected contribution | Evidence required | Failure or uncertainty risk | Boundary statement |
|---|---|---|---|---|---|---|---|
| **Hidden work** | Visible task completion omits clinically consequential activity. | Workflow, records, messages, interruptions, and local task definitions. | Proposed umbrella construct linking unmeasured work to usable medication decisions; informed by redistributed digital work [3] and work extending beyond visible encounters [5]. | Makes omitted work available for analysis and design. | Observation, task analysis, logs, interviews, and case tracing. | May label inefficiency or routine effort as clinically necessary. | The construct does not establish prevalence, value, or inevitability. |
| **Context-recovery work** | Digital data may lack chronology, goals, setting, or treatment constraints. | Patient account, medication history, prior records, and team knowledge. | Proposed reconstruction of decision-relevant context before interpretation. | Identifies contextual gaps and recovery burden. | Case review comparing available and recovered context. | Recovered information may remain incomplete or conflicting. | No universal minimum context dataset is claimed. |
| **Information-repair work** | Entries may be stale, duplicated, contradictory, or provenance-poor. | Medication lists, allergy records, laboratory data, terminology, and source metadata. | Proposed correction or qualification of information used by professionals and systems. | Improves traceability of input quality. | Record audit, discrepancy classification, and correction-pathway analysis. | A discrepancy is not automatically an error. | Repair does not guarantee a correct decision. |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| **Interpretive and relational work** | Scores or alerts do not determine patient-specific meaning or feasibility. | Evidence, uncertainty, patient goals, preferences, literacy, and team perspectives. | Proposed translation of outputs through judgment, explanation, and negotiation; conditioned by workflow and usability [7]. | Preserves meaning, trust, and feasible planning. | Observed consultations, rationale review, and patient and team inquiry. | Communication volume may be mistaken for relational quality. | Interpretation is not proof of appropriate action. |
| **Responsibility-bearing work** | Ownership may be unclear after an AI-supported recommendation. | Authority, escalation routes, documentation duties, and follow-up capacity. | Proposed assignment of verification, decision, communication, monitoring, and closure. | Makes accountability and handoffs inspectable. | Governance audit, incident tracing, and follow-up review. | Responsibility may silently diffuse to the end user. | The framework does not determine legal liability. |
| **Preservation loop** | Components may occur as disconnected tasks without closure. | All preceding domains plus organizational resources and governance. | Proposed sequence: detect need, recover context, repair information, interpret and negotiate, decide or escalate, document ownership, monitor, and feed back. | Organizes workflow mapping and evaluation. | Prospective testing across settings and decision types. | A formal loop may add burden or become performative documentation. | It is not a validated clinical algorithm or mandatory workflow. |

### Context-Recovery and Information-Repair Work

AI-supported medication decisions are conditional on what the system and pharmacist can know about the patient at the relevant time. Context-recovery work includes reconstructing what medicines are actually used, why they were started or stopped, what happened previously, which goals matter now, and which clinical, social, and access constraints shape feasible options. Medication reconciliation exemplifies this work across care transitions, although evidence about its downstream effects is heterogeneous and should not be converted into a claim of uniform benefit [9]. The conceptual point is narrower: a medication list is not equivalent to a trustworthy account of treatment.

Information repair begins when available records cannot safely support interpretation. Pharmacists may need to resolve duplicated therapies, ambiguous allergy labels, conflicting doses, outdated laboratory results, missing indications, or unclear source dates. Repair is not merely data cleaning; it is an organizational process involving role clarity, training, communication, workflow, and technology. Multicomponent reconciliation improvement supports this bundled view, but does not establish one universally transferable implementation package [10]. Accordingly, the proposed framework evaluates who detects a problem, who can amend the record, how uncertainty is communicated, and whether corrected information reaches downstream users and systems.

Context recovery and information repair are related but distinct. Recovery seeks missing meaning; repair changes or qualifies the information representation. Both may be triggered by an AI output, but both may also expose that the output should not be used. Drug-allergy decision support illustrates dependence on documentation and terminology quality, while machine-learning prioritization of prescriptions illustrates that risk ranking can guide pharmacist attention without independently confirming an error or closing the medication decision [11, 12]. Appropriate non-use, escalation, or deferral is therefore an intended possible outcome of the framework. Evaluation should measure unresolved gaps, provenance access, correction cycles, duplicated effort, interruptions, and downstream propagation rather than assuming that faster review reflects less work or better care.

### Interpretive, Relational, and Responsibility-Bearing Work

Interpretive work begins where a digital output stops. A pharmacist must determine whether an alert, score, explanation, or recommendation is relevant, sufficiently grounded, timely, and actionable for this patient. Relational care adds empathy, compassion, trust, explanation, and attention to lived circumstances that are not reducible to predictive accuracy or task completion [13]. In the proposed framework, these are not decorative additions to a technical decision; they influence whether information is understood, preferences are elicited, and a plan is feasible.

Interpretation must also remain value-flexible. Medication choices can involve competing goals, uncertain evidence, burdens, affordability, adherence constraints, and different tolerances for risk; AI advice therefore requires connection to patient-specific values rather than mechanical execution [14]. This does not authorize unconstrained discretion. It requires explicit separation of factual evidence, model uncertainty, professional judgment, and preference-sensitive choice.

**Figure 1** presents the original conceptual synthesis for hidden work in ai-supported pharmacy care, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

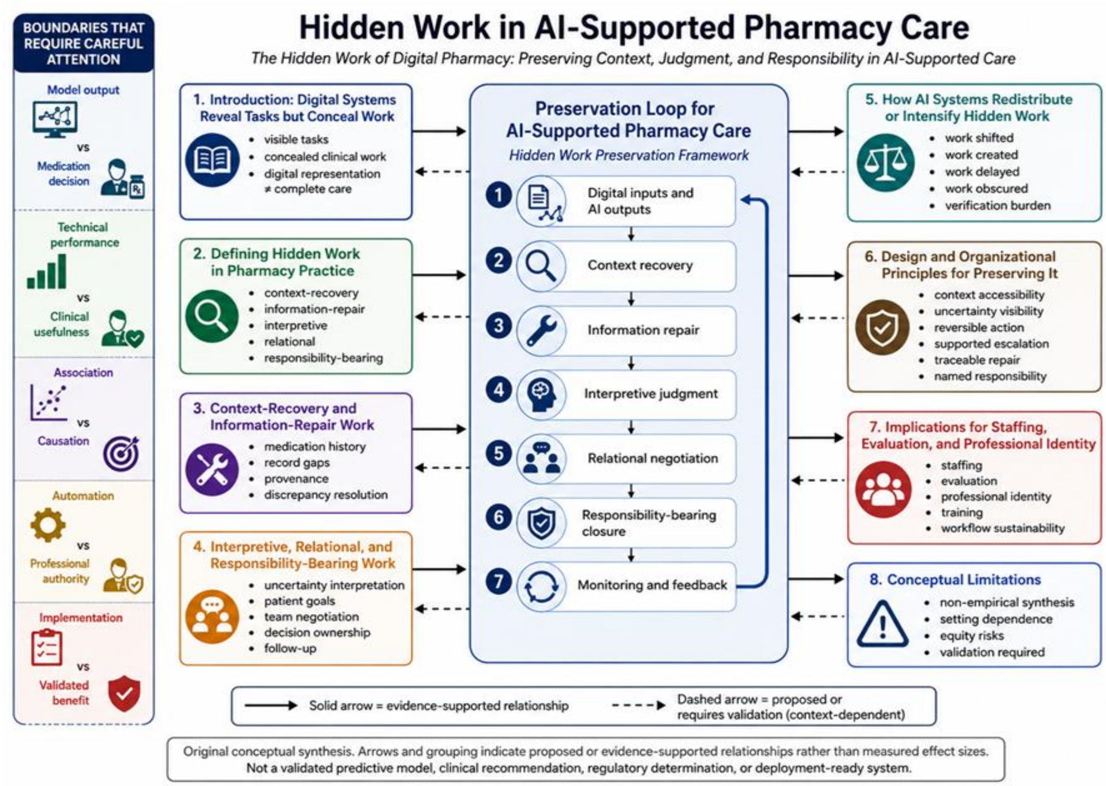


**Figure 1.** Hidden Work in AI-Supported Pharmacy Care

The figure is an original conceptual synthesis developed for “The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care.” 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.

Responsibility-bearing work converts interpretation into accountable action. Because clinicians can be influenced by erroneous AI advice, verification must include active checking, calibrated reliance, and the capacity to reject or escalate an output rather than treating human presence as a sufficient safeguard [15]. Responsibility also extends beyond the individual decision to data stewardship, transparency, fairness, privacy, safety, organizational oversight, and learning from incidents [16]. The proposed responsibility contract therefore asks who owns the data, model monitoring, final medication decision, rationale record, escalation, patient communication, follow-up, and incident review. It does not determine jurisdiction-specific liability or make the pharmacist solely responsible for system defects.

**Table 2** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care.

**Table 2.** Evaluation criteria, evidence requirements, and failure modes for the article The Hidden Work of Digital Pharmacy: Preserving Context, Judgment, and Responsibility in AI-Supported Care.

| Architecture layer or process stage | Core function | Information flow | Interaction with other components | Evaluation criterion | Potential failure mode | Human or organizational responsibility | Readiness boundary |
|---|---|---|---|---|---|---|---|
| **Context intake and recovery** | Establish decision-relevant history, goals, timing, and constraints. | Patient, record, team, and source metadata enter a qualified context account. | Precedes repair and interpretation. | Context fidelity and visibility of unresolved gaps. | Correct-looking output derived from stale or incomplete context. | Ensure access, provenance, discrepancy escalation, and patient verification. | Do not broaden decision use while material context gaps are undetectable or incommunicable. |
| **Information repair** | Correct or qualify contradictions, omissions, and | A detected discrepancy leads to a corrected entry, uncertainty label, | Changes the input conditions for AI and professional judgment. | Repair burden, correction traceability, and | A local correction fails to reach other records or systems. | Assign correction authority, communication | Efficiency claims require end-to-end measurement rather |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | terminology problems. | and downstream communication. | | propagation completeness. | | routes, and closure checking. | than local task time. |
| **AI output and interpretive judgment** | Assess relevance, evidence, uncertainty, and actionability. | Model output and recovered context inform acceptance, rejection, deferral, or escalation. | Depends on context and informs relational discussion. | Calibrated reliance, rationale quality, and error recovery [15]. | Automation influence, unjustified confidence, or reflexive rejection. | Provide training, verification resources, and escalation authority. | Autonomy cannot be inferred from model performance or nominal human oversight. |
| **Relational negotiation** | Elicit goals, explain uncertainty, and form a feasible plan. | Interpreted options are exchanged among the patient, pharmacist, and care team. | May revise interpretation or trigger further context recovery. | Understanding, concordance, accessibility, and continuity; relational relevance is evidence-informed [13]. | Throughput gains are achieved by truncating explanation or preference elicitation. | Protect communication time and accessible alternatives. | Task speed cannot substitute for direct evidence of patient-relevant quality. |
| **Responsibility-bearing closure** | Assign decision ownership, documentation, monitoring, and follow-up. | Decisions and rationales pass to named owners, handoffs, surveillance, and audit trails. | Closes the preservation loop and supports organizational learning. | Responsibility traceability and closure completeness [16]. | Diffuse ownership, silent handoff, or unmonitored consequences. | Specify a responsibility contract across professional and organizational actors. | No scale-up where ownership, escalation, or incident review remains ambiguous. |
| **Whole-workflow evaluation** | Determine what work is removed, shifted, created, delayed, or concealed. | Cross-role workflow, logs, observations, patient evidence, and incident evidence support system-level assessment. | Integrates all layers and tests the proposed framework. | Medication safety, workflow redistribution, equity, and sustainability. | Apparent local efficiency is accompanied by downstream repair, interruption, or access burden. | Fund prospective evaluation, maintenance, and feedback processes. | The framework remains conceptual until tested in the intended setting and population. |

## How AI Systems Redistribute or Intensify Hidden Work

AI does not simply remove or retain work. It may eliminate a visible task, shift work to another role, create new verification activity, delay work until an exception occurs, or conceal work from routine measurement. Machine learning can introduce silent errors, altered dependence, and unanticipated practice effects even when technical performance appears satisfactory [17]. The proposed framework therefore treats redistribution as a set of hypotheses to be tested. A faster review screen may reduce manual searching while increasing investigation of ambiguous outputs; an automated prioritization tool may concentrate attention while allowing low-ranked cases to receive less scrutiny; and a recommendation may shorten selection while expanding explanation, documentation, or follow-up. None of these consequences should be assumed from the presence of AI alone.

Clinical decision support also requires continuing data, rule, interface, and governance work. Maintenance, monitoring, override handling, error recovery, updating, and integration are distributed across pharmacists, technicians, informatics personnel, vendors, managers, and other clinicians [18]. This work can occur before use, during a patient encounter, or after a decision has apparently been completed. Evaluation should therefore trace workload across roles and time rather than reporting only the duration of the supported task. The relevant comparison is not "human work versus automation," but the total configuration of human, technical, and organizational work needed to produce and sustain a defensible medication decision.

Explainability may itself redistribute work. Post-hoc explanations can require clinicians to interpret another technical artifact, defend acceptance or rejection, and reconcile an apparently plausible explanation with contradictory patient information without assuring causal understanding or reliability [19]. Explanations may still support inquiry, but their value must be evaluated through actual decision use, error recognition, patient communication, and burden distribution. The proposed preservation test is therefore whether an AI feature makes material uncertainty, provenance, alternatives, and responsibility more inspectable without transferring unsupported justification work to the pharmacist or patient.

## Design and Organizational Principles for Preserving It

Preservation should be designed at the level of the whole work system and patient journey rather than the AI interface alone [20]. The proposed principles are context accessibility, uncertainty visibility, interruptibility, reversible action, traceable repair, supported escalation, and named responsibility. Systems should permit pharmacists to inspect relevant sources, qualify uncertain inputs, disagree with outputs without punitive friction, and reopen decisions when new information emerges. Organizations should distinguish valuable judgment from compensatory rework: the goal is not to preserve every existing task, but to retain protective functions while redesigning avoidable defects.

Implementation decisions should also examine the technology, value proposition, intended adopters, organizational capacity, wider institutional environment, and adaptation over time [21]. A tool that functions in a controlled evaluation may become unsustainable when data ownership is unclear, local workflows differ, maintenance resources are absent, or benefits accrue to one group while repair work is transferred to another. The framework therefore proposes organizational ownership for data correction, model surveillance, training, patient communication, incident response, and periodic reassessment of work redistribution.

Before expansion, early live evaluation should examine actual workflow, human factors, errors, safety signals, iterative modifications, and differences between intended and observed use [22]. Proposed transition gates are adequate context and provenance, recoverable errors, calibrated human reliance, explicit decision and follow-up ownership, acceptable patient and subgroup consequences, and sustainable maintenance capacity. Failure at a gate should trigger redesign, restricted use, or further evaluation rather than automatic progression toward scale. These gates organize research and governance; they are not validated thresholds, regulatory criteria, or evidence of clinical benefit.

### *Implications for Staffing, Evaluation, and Professional Identity*

Staffing analyses should count net distributed work, not merely tasks apparently automated. A decision-support function may be analytically useful while adding navigation, interpretation, and correction work when usability or task–technology fit is poor [23]. Workforce evaluation should therefore measure time, interruptions, unresolved cases, cross-role transfers, after-hours work, escalation, and follow-up alongside safety, quality, and patient experience. The framework supports neither staffing reduction nor expansion without direct local evidence.

Pharmacy AI applications and evaluation approaches remain heterogeneous, so automation of selected activities cannot be equated with established workforce substitution [24]. Evaluation should compare complete medication-care pathways, including ordinary cases, exceptions, non-use, rejected recommendations, record repair, patient explanation, and downstream monitoring. It should also identify whose work is reduced and whose work increases. Apparent productivity gains are incomplete when they depend on unmeasured technician labor, patient self-navigation, informal pharmacist checking, or deferred correction elsewhere in the system.

Professional identity should consequently develop around accountable digital practice rather than protection of residual tasks. Digital pharmacy requires competence in data quality, technology appraisal, uncertainty communication, interdisciplinary coordination, ethics, and critical use of automated outputs [25]. Education should prepare pharmacists to recognize when context is missing, challenge inappropriate recommendations, document reasoning, negotiate patient-specific plans, and participate in system governance. This proposed identity is neither technological enthusiasm nor automation resistance; it treats stewardship of medication decisions as a professional responsibility that persists as task boundaries change.

### *Conceptual Limitations*

The framework is an original conceptual synthesis and may omit forms of hidden work specific to community, hospital, ambulatory, remote, or resource-constrained practice. Its categories may overlap, and activities classified as protective in one setting may be wasteful or harmful in another. Clinical and machine-learning systems cannot be assumed to generalize across populations, organizations, jurisdictions, or changing conditions without local validation [26].

Equity is also not secured by aggregate performance. Operational proxies can encode structural inequity, and overall accuracy can conceal subgroup or intersectional under-service [27, 28]. Pharmacy evaluation must therefore examine target validity, access, error recovery, communication burden, and who bears context-repair work. These sources establish an evaluation principle, not pharmacy-specific evidence that the proposed framework prevents inequity.

The framework does not predict implementation success or sustainability [21]. Reporting a live evaluation does not itself establish safety, effectiveness, or readiness [22]. Local validation does not guarantee transportability or future stability after data, workflows, products, or populations change [26]. The framework specifies no safe autonomy threshold, staffing ratio, legal liability allocation, or regulatory status; empirical work may need to revise, reject, or extend its constructs.

## Conclusion

AI-supported pharmacy care cannot be evaluated adequately by asking only whether a model predicts well or a task is completed faster. Medication decisions depend on work that digital representations may conceal: recovering context, repairing information, interpreting uncertainty, negotiating feasible care, assigning responsibility, and monitoring consequences. This article proposes the Hidden Work Preservation Framework to make those functions visible as objects of design, staffing, governance, education, and evaluation.

The framework's principal claim is conditional. AI may remove avoidable effort, but it may also shift, create, delay, or obscure work and responsibility. Preserving professional judgment does not require retaining every manual task; it requires identifying which functions protect medication safety, patient agency, equity, relational continuity, and accountable closure, then testing whether a new work configuration sustains them. Evaluation must therefore connect technical performance to workflow consequence, decision quality, burden distribution, patient experience, organizational capacity, and responsibility traceability.

The framework remains non-empirical. It is not a clinical recommendation, validated workflow, staffing formula, regulatory standard, or deployment-ready architecture. Its value lies in offering testable distinctions and decision boundaries for examining what AI-supported pharmacy care requires beyond visible task completion. Future studies should determine where the proposed categories hold, how they interact, which forms of hidden work should be preserved or eliminated, and whether explicit preservation improves care without creating new inequities or unsustainable burdens.

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. Ranchon F, Chanoine S, Lambert-Lacroix S, Bosson JL, Moreau-Gaudry A, Bedouch P. Development of artificial intelligence powered apps and tools for clinical pharmacy services: A systematic review. Int J Med Inform. 2023;172:104983. doi:10.1016/j.ijmedinf.2022.104983
3. Bergey MR, Goldsack JC, Robinson EJ. Invisible work and changing roles: Health information technology implementation and reorganization of work practices for the inpatient nursing team. Soc Sci Med. 2019;235:112387. doi:10.1016/j.socscimed.2019.112387
4. 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
5. Arndt BG, Beasley JW, Watkinson MD, Temte JL, Tuan WJ, Sinsky CA, et al. Tethered to the EHR: Primary care physician workload assessment using EHR event log data and time-motion observations. Ann Fam Med. 2017;15(5):419-26. doi:10.1370/afm.2121
6. Kjos AL, Bryant GA. Communication networks of medication management in an ambulatory care setting. Res Social Adm Pharm. 2019;15(2):182-92. doi:10.1016/j.sapharm.2018.04.010
7. 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
8. Melnick ER, Dyrbye LN, Sinsky CA, Trockel M, West CP, Nedelec L, et al. The association between perceived electronic health record usability and professional burnout among US physicians. Mayo Clin Proc. 2020;95(3):476-87. doi:10.1016/j.mayocp.2019.09.024
9. Redmond P, Grimes TC, McDonnell R, Boland F, Hughes C, Fahey T. Impact of medication reconciliation for improving transitions of care. Cochrane Database Syst Rev. 2018;8(8):CD010791. doi:10.1002/14651858.CD010791.pub2
10. Schnipper JL, Mixon A, Stein J, Wetterneck TB, Kaboli PJ, Mueller S, et al. Effects of a multifaceted medication reconciliation quality improvement intervention on patient safety: Final results of the MARQUIS study. BMJ Qual Saf. 2018;27(12):954-64. doi:10.1136/bmjqs-2018-008233
11. Légat L, Van Laere S, Nyssen M, Steurbaut S, Dupont AG, Cornu P. Clinical decision support systems for drug allergy checking: Systematic review. J Med Internet Res. 2018;20(9):e258. doi:10.2196/jmir.8206
12. 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
13. Kerasidou A. Artificial intelligence and the ongoing need for empathy, compassion and trust in healthcare. Bull World Health Organ. 2020;98(4):245-50. doi:10.2471/BLT.19.237198
14. McDougall RJ. Computer knows best? The need for value-flexibility in medical AI. J Med Ethics. 2019;45(3):156-60. doi:10.1136/medethics-2018-105118
15. 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
16. Morley J, Machado CCV, Burr C, Cowls J, Joshi I, Taddeo M, et al. The ethics of AI in health care: A mapping review. Soc Sci Med. 2020;260:113172. doi:10.1016/j.socscimed.2020.113172
17. 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
18. 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(1):17. doi:10.1038/s41746-020-0221-y
19. 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
20. Carayon P, Wooldridge A, Hoonakker P, Hundt AS, Kelly MM. SEIPS 3.0: Human-centered design of the patient journey for patient safety. Appl Ergon. 2020;84:103033. doi:10.1016/j.apergo.2019.103033
21. 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
22. 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
23. Marcilly R, Colliaux J, Robert L, Pelayo S, Beuscart JB, Rousselière C, et al. Improving the usability and usefulness of computerized decision support systems for medication review by clinical pharmacists: A convergent, parallel evaluation. Res Social Adm Pharm. 2023;19(1):144-54. doi:10.1016/j.sapharm.2022.08.012
24. 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
25. Silva ROS, de Araújo DCSA, dos Santos Menezes PW, Neves ERZ, de Lyra DP Jr. Digital pharmacists: The new wave in pharmacy practice and education. Int J Clin Pharm. 2022;44(3):775-80. doi:10.1007/s11096-021-01365-5
26. 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):e489-92. doi:10.1016/S2589-7500(20)30186-2
27. 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
28. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82. doi:10.1038/s41591-021-01595-0