Original Article

Keeping the Patient Visible when Pharmacy Services Become Data-Driven

Youssef Hariri[1]*, Hassan Nasser[1], Fatima Zahra[2]

[1]Department of Patient-Centered Pharmacy Informatics, Faculty of Pharmacy, Mohammed V University, Rabat, Morocco. [2]Department of Data-Driven Care and Patient Visibility, Faculty of Pharmacy, University of Casablanca, Casablanca, Morocco.

Abstract

Data-driven pharmacy services can combine electronic medication records, patient-generated information, clinical decision support, and artificial intelligence to improve the organization and interpretation of medication-related data. However, a technically comprehensive record may still obscure the person whose treatment is being considered. Important explanations, preferences, practical capabilities, and social constraints may be absent, outdated, poorly interpreted, or disconnected from the eventual decision. This article develops an original Patient-Visibility Framework for examining whether the patient remains meaningfully represented when pharmacy services become data-driven. Patient visibility is proposed as a decision property arising when relevant context is captured with provenance, interpreted for the current medication task, open to patient correction, discussed when uncertainty or discordance exists, connected transparently to professional action, and updated over time. The framework organizes visibility through four interacting layers: narrative, preference, capability, and social context. These layers enter a sociotechnical pathway comprising context capture, interpretation, pharmacist–patient dialogue, medication action, follow-up, and accountability. Proposed failure modes include patient–record substitution, narrative flattening, preference freezing, capability misclassification, context staleness, unequal missingness, workflow bypass, and responsibility diffusion. Evaluation therefore requires more than model performance. Construct validity, representational fidelity, workflow feasibility, medication-safety consequences, patient experience, equity, governance, and temporal validity must be examined separately. The framework is a conceptual and evaluative contribution rather than a validated instrument, clinical protocol, or deployment standard. Its categories, relationships, transition gates, and proposed indicators require empirical testing across pharmacy tasks, populations, technologies, and care settings.

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

Introduction

The Risk of Replacing the Patient with the Record

Digital pharmacy increasingly organizes medication decisions through electronic histories, coded diagnoses, dispensing records, laboratory results, adherence indicators, risk scores, alerts, and algorithmic recommendations. These representations can improve access to relevant information, but they can also encourage a subtle substitution: the available record becomes treated as though it were the patient. Computationally useful outputs may produce unintended cognitive, organizational, and clinical consequences when their limitations are overlooked or when prediction is mistaken for understanding [1].

The central problem is therefore not simply whether more patient data are available, but whether the person remains recognizable and influential within the decision process.

Patient visibility includes a relational dimension. Attention, preparation, listening, agenda elicitation, and connection are identifiable practices through which professionals recognize

Address for correspondence: Youssef Hariri, Department of Patient-Centered Pharmacy Informatics, Faculty of Pharmacy, Mohammed V University, Rabat, Morocco.
youssef.hariri@um5.ac.ma
Received: 16 February 2025; **Accepted:** 28 May 2025

How to cite this article: Hariri Y, Nasser H, Zahra F. Keeping the Patient Visible when Pharmacy Services Become Data-Driven. Arch Pharm Pract. 2025;16(2):53-61. https://doi.org/10.51847/KDzAs5zVzk

patients as participants rather than data sources [2]. In pharmacy practice, this recognition matters because medication use occurs outside the consultation and is shaped by experiences, concerns, routines, resources, competing responsibilities, and willingness to accept particular trade-offs. A service may be technically individualized while remaining insufficiently person-centred if these factors cannot be expressed, questioned, or connected to action.

Electronic records can improve information availability while simultaneously altering eye contact, conversational flow, information sharing, and patient–professional communication [3]. The problem is not inherent to digitization: records may support continuity, recall, and coordination. The risk arises when structured representations acquire unwarranted authority, incomplete entries are treated as absence of need, or documentation activity displaces clarification. The patient may then become visible as medication history, risk category, or utilization pattern but invisible as an interpreter of her own circumstances.

Artificial intelligence and clinical decision support intensify this distinction because they can transform recorded variables into classifications or recommendations that appear decisive. Such systems remain aids to clinical reasoning whose value depends on the decision context, data quality, workflow, and human interpretation [4]. Model output is therefore not equivalent to a medication decision, technical performance is not equivalent to clinical usefulness, and automation does not transfer professional authority or accountability.

This article proposes a Patient-Visibility Framework for data-driven pharmacy services. It does not reject structured data, automation, or artificial intelligence. Instead, it asks what contextual information must remain interpretable, who may correct it, how uncertainty and disagreement should be handled, and how patient context should connect to pharmacy action without producing indiscriminate data collection. The contribution is an original non-empirical architecture whose constructs and evaluation requirements remain subject to validation.

What Patient Visibility Means in Digital Pharmacy

Person-centred care requires attention to individual needs, values, participation, communication, and the environment in which care is delivered [5]. Building on this foundation, patient visibility is proposed here as the extent to which information required to understand a medication decision remains meaningfully connected to the person concerned. Visibility is not established merely because a patient can be identified, a record is populated, or an algorithm uses individual-level variables.

Preference is one component of visibility, but it cannot be represented reliably as a single permanent choice. Shared decision-making requires discussion of options, priorities, and the trade-offs that fit the person's situation [6]. A preference recorded during one encounter may change with symptoms, treatment experience, affordability, caregiving responsibilities, or new information. Preference visibility therefore requires revisability and an opportunity for deliberation rather than passive retrieval of a prior selection.

Patient-generated health data further demonstrate why availability and visibility are different. Patients and professionals may differ in what they expect data to represent, how much information should be collected, and what should follow from its review [7]. Data may be accurate yet irrelevant to the immediate task, clinically important yet difficult to interpret, or burdensome to review without an assigned action pathway. Visibility consequently depends on semantic meaning, provenance, task relevance, and responsibility, not volume alone.

Electronic records also differ in whether their language and structure preserve the person behind the documented facts [8]. A medication list may omit why treatment was stopped. An adherence indicator may conceal intolerable adverse effects, unstable housing, or inability to obtain refills. A documented refusal may conceal an unaddressed preference or communication barrier. The record should therefore be treated as a contestable representation rather than an authoritative replacement for the patient.

The framework proposes that patient visibility is present when decision-relevant context is captured, interpretable, current, correctable, and connected to accountable action. These conditions are task-specific. A routine supply transaction may require less contextual depth than initiating a high-risk medicine or resolving contradictory medication histories. More information is not automatically better: excessive capture can create intrusion, stigma, workload, and false confidence.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Keeping the Patient Visible when Pharmacy Services Become Data-Driven.

Table 1. Definitions, components, relationships, and intended uses for the article Keeping the Patient Visible when Pharmacy Services Become Data-Driven.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement

Patient visibility	A populated record may not preserve the person's meaning or influence.	Needs, goals, experiences, medication task, opportunity to participate	**Proposed:** visibility arises when relevant context is captured, interpreted, current, correctable, and linked to action.	Distinguishes person representation from data availability.	Construct testing with patients and pharmacy professionals; comparison across tasks and settings	Visibility may be confused with data completeness or patient satisfaction.	This is a proposed construct, not a validated score.
Relational visibility	Screen and documentation activity may displace listening or connection.	Professional attention, consultation design, communication conditions	Relational practices support recognition of the patient as a participant [2].	Preserves dialogue alongside digital processing.	Observational and qualitative studies of consultations	Presence may be reported but not reflected in decisions.	Relational quality cannot be inferred from interface use alone.
Narrative	Structured entries may omit explanations for medication use, non-use, concerns, or change.	Patient account, treatment experience, chronology, provenance	**Proposed:** narrative supplies explanatory context for interpreting records and signals.	Supports clarification of apparent discrepancies.	Patient-confirmed semantic-fidelity assessment	Narrative may be incomplete, inaccurate, excessive, or flattened during summarisation [7].	Narrative should not automatically override corroborated evidence.
Preference	Static fields may freeze a choice outside its original context.	Goals, priorities, acceptable burdens, desired involvement	Deliberation links available options to what matters to the patient [6].	Supports preference-sensitive medication decisions.	Decision-specific assessment of concordance and revisability	Recorded preference may be stale, constrained, or misunderstood.	Not every pharmacy task requires extensive preference elicitation.
Capability	Access may be mistaken for practical ability to use a service or medicine.	Skills, language, cognition, sensory or physical ability, available support	**Proposed:** capability moderates whether a medication or digital pathway is feasible.	Identifies implementation barriers that may otherwise appear as nonadherence.	Accessibility, usability, and task-performance evaluation	Capability may be reduced to a deficit attributed to the patient.	Capability is contextual and should not be treated as a fixed trait.
Social context	Treatment feasibility may depend on circumstances absent from clinical records.	Affordability, housing, transport, work, caregiving, service access	**Proposed:** social context modifies the feasibility and consequences of medication options.	Supports recognition of constraints affecting medication use.	Consent, provenance, currency, and actionability studies	Sensitive information may be collected without an available response.	Pharmacists cannot be assumed to resolve every identified social need.
Record correctability	Recorded information may acquire authority despite error or staleness.	Patient access, correction route, documentation governance	Person-centred record content should remain open to clarification [8].	Reduces patient–record substitution.	Evaluation of correction uptake, disputes, and downstream changes	Corrections may not propagate across systems.	Patient correction does not remove the need for professional verification.
Context-to-action connection	Context may be collected but never influence care.	Decision rationale, pharmacist interpretation, documented response	**Proposed:** the workflow records whether context changed, confirmed, deferred, or did not affect action.	Makes patient influence and professional responsibility auditable.	Prospective workflow and documentation evaluation	Documentation may become burdensome or defensive.	A recorded trace does not itself demonstrate decision quality.

Proposed Patient-Visibility Framework

Patient visibility should be located across the pharmacy work system rather than assigned to one screen, profession, or data field. Human-centred work-system analysis shows that patient journeys arise through interactions among people, tasks, technologies, physical and organizational environments, and transitions between settings [9]. A failure of visibility may therefore originate during data capture, system integration, professional review, patient communication, or follow-up even when each isolated component appears functional.

Implementation is also shaped by interacting forms of complexity. Technologies can be abandoned, poorly scaled, or sustained unevenly because of features of the condition, technology, adopters, organization, wider system, and adaptation over time [10]. The proposed framework consequently treats visibility as conditional and dynamic. A contextual representation that works in one service cannot be presumed transferable to another medication task, patient population, interface, or organizational setting.

Clinical decision-support systems convert data and encoded knowledge into prompts, classifications, or recommendations, but their practical value depends on data quality, representation, workflow integration, and user response [11]. Within the proposed architecture, decision support occupies an interpretive position: it may organize or

prioritize information, but it does not determine what the patient values, whether the information remains current, or which action should be taken.

Trustworthy use further requires attention to transparency, usability, explainability, validation, and intended context [12]. These characteristics are necessary considerations but are not individually sufficient. A transparent system may still omit important context; an explainable output may still be based on an inappropriate proxy; and a usable interface may accelerate an inadequately framed decision.

The proposed framework contains four visibility layers—narrative, preference, capability, and social context—connected to five process stages: provenance-aware capture, contextual interpretation, pharmacist–patient dialogue, medication action, and follow-up updating. Three proposed gates organize movement through these stages. A visibility-sufficiency gate asks whether the available context is proportionate to the consequence and reversibility of the decision. A context-discordance gate identifies conflicts among the record, patient account, system output, and professional interpretation. An actionability gate asks whether identified information can be addressed, communicated, referred, monitored, or explicitly bounded.

Figure 1 presents the original conceptual synthesis for patient-visibility framework for data-driven pharmacy services, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

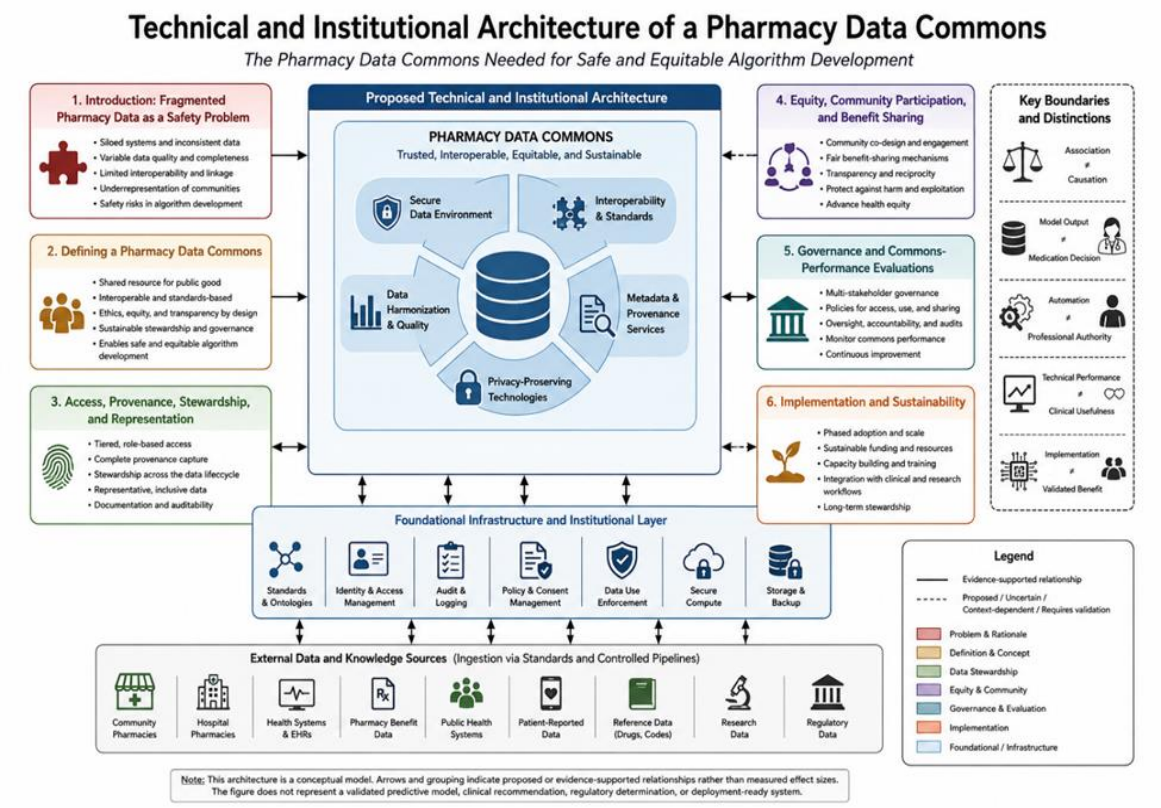

Figure 1. Patient-Visibility Framework for Data-Driven Pharmacy Services

The architecture treats visibility failures as more than missing-data problems. Patient–record substitution occurs when documented information is privileged over a current patient account without adequate reconciliation. Narrative flattening removes explanatory meaning during structuring or summarisation. Preference freezing treats an earlier choice as permanent. Capability misclassification interprets practical difficulty as unwillingness. Visibility debt accumulates when contextual information becomes stale, unverified, or disconnected from responsibility. Unequal invisibility occurs when some groups are less able to contribute, correct, or benefit from contextual data.

A visibility-preserving pharmacy service must be treated as a sociotechnical intervention whose safety depends on work-system fit, decision-support design, usability, transparency, and accountable human use [9, 11, 12]. These relationships support the architecture's rationale but do not validate its layers or gates. The framework is intended to organize design questions and empirical evaluation, not to prescribe a universal workflow.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes

for the article Keeping the Patient Visible when Pharmacy Services Become Data-Driven.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Keeping the Patient Visible when Pharmacy Services Become Data-Driven.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Provenance-aware capture	Obtain decision-relevant context without assuming completeness.	Patient, record, caregiver, and service information enter with source and time status.	Supplies all four visibility layers.	Semantic fidelity, provenance, currency, correction rate	Missingness hidden by a populated record; intrusive overcollection	Organization defines collection purpose; patient can clarify; pharmacist verifies when relevant.	**Proposed:** no progression when critical information lacks identifiable source or meaning.
Visibility-sufficiency gate	Judge whether context is proportionate to the medication task.	Current information is compared with decision consequence, uncertainty, and reversibility.	Determines whether interpretation can proceed.	Detection of decision-critical missingness; appropriate clarification	Routine decisions overburdened or consequential decisions under-contextualized	Pharmacist retains task-specific judgment.	Requires construct and inter-rater validation before operational use.
Contextual interpretation	Reconcile medication information, patient meaning, and system output.	Structured and narrative information are reviewed together.	Connects visibility layers with decision support [11].	Accuracy of interpretation; discrepancy recognition	Narrative flattening; automation bias; inappropriate proxy use	Pharmacist interprets; system exposes provenance and uncertainty.	Technical accuracy alone is insufficient.
Context-discordance gate	Identify material conflict among sources or interpretations.	Contradictions are returned for clarification rather than automatically resolved.	Activates dialogue, correction, deferral, or escalation.	Appropriate detection and resolution of discordance	Excessive false prompts or unnoticed contradictions	Organization defines escalation routes; pharmacist determines relevance.	**Proposed:** thresholds must be task- and setting-specific.
Pharmacist–patient dialogue	Clarify meaning, preferences, feasibility, and uncertainty.	Information moves bidirectionally and may correct the record.	Revises all four layers and professional interpretation.	Patient opportunity to correct; preference concordance; burden	Token consultation after the decision is effectively fixed	Pharmacist supports deliberation; patient controls personal meaning.	Dialogue quality cannot be inferred from completion of a field.
Actionability gate	Determine what can responsibly follow from identified context.	Context is linked to action, referral, monitoring, explicit non-action, or boundary documentation.	Prevents data collection without an accountable response.	Context-to-action trace; appropriate referral or follow-up	Sensitive need identified without capability or responsibility to respond	Service defines ownership and limits.	Identification does not create unlimited professional responsibility.
Medication action	Integrate evidence, patient context, professional judgment, and decision support.	Model output remains one input rather than final authority [4].	Produces a documented medication decision or service action.	Decision appropriateness, safety events, overrides, unintended effects	Model output treated as instruction; context used selectively	Pharmacist retains professional accountability.	No deployment claim follows from framework plausibility.
Follow-up and updating	Test whether context and decisions remain valid over time.	Outcomes, changes, and corrections return to the relevant layers.	Reduces visibility debt and supports continuity.	Staleness, update completion, propagation across interfaces	Outdated preferences or circumstances persist as authoritative	Organization assigns update and monitoring ownership.	Update frequency requires empirical and local justification.
Governance boundary	Control privacy, equity, monitoring, workload, and responsibility.	Audit information spans all stages.	Constrains every layer and gate [10].	Accountability, access control, equity differences, incident response	Responsibility diffusion or surveillance without benefit	Leadership, informatics teams, pharmacists, and patient representatives share defined roles.	Governance arrangements require local evaluation and cannot certify benefit.

Narrative, Preference, Capability, and Social-Context Layers

The narrative layer represents the patient's account of medication use, treatment experience, symptoms, concerns, reasons for change, and interpretation of what has happened. Patient-generated information can enrich care, but its integration raises questions about engagement, review, workload, actionability, and responsibility [13]. Narrative visibility therefore requires provenance, chronology, and preservation of meaning rather than indiscriminate accumulation. A summary should disclose what was compressed or remains uncertain, and the patient should have a route to correct consequential misrepresentation.

The preference layer represents goals, priorities, acceptable burdens, desired involvement, and willingness to accept trade-offs. Resistance to medical artificial intelligence may arise when people believe that a system cannot recognize their uniqueness [14]. This evidence does not establish how every patient will respond to pharmacy AI, but it supports a design concern: recommendations derived from population patterns should not be presented as though they already incorporate individual priorities. Preferences should remain decision-specific, revisable, and distinguishable from assumptions inferred from behaviour.

The capability layer concerns whether a patient can practically use a medicine, device, communication channel, or digital service. Digital inequality extends beyond physical access to differences in the quality, continuity, and diversity of available technological resources [15]. In pharmacy practice, capability may also depend on language, cognition, vision, dexterity, health literacy, time, caregiver support, connectivity, and the complexity of administration. The framework treats capability as relational: a person may manage one regimen or interface successfully but not another. Difficulty should not automatically be labelled nonadherence or resistance.

The social-context layer represents conditions that shape medication feasibility, including affordability, housing, food access, transport, work, caregiving, and access to services. Incorporating social-determinants data into electronic records requires decisions about sourcing, standardisation, consent, updating, interoperability, and actionability [16]. Capturing these conditions may improve interpretation of medication use, but collection without a response pathway can expose sensitive information while leaving the underlying constraint unchanged. The framework therefore requires an actionability boundary identifying what the pharmacy service can address, communicate, refer, or monitor.

The four layers are analytically distinct but practically interdependent. A stated preference may reflect cost or caregiving constraints. A capability problem may be interpreted through narrative. Social circumstances may explain apparent nonadherence, while a treatment experience may alter both preference and capability. These connections are proposed relationships rather than validated causal pathways. Their relevance should be evaluated for each medication task instead of assuming that every contextual element must be collected.

Proportionality is the principal boundary. Contextual depth should reflect foreseeable harm, decision reversibility, uncertainty, and the burden imposed on the patient and pharmacy team. The layers are not an exhaustive ontology of personhood, and they should not convert the patient into a larger record. Their purpose is narrower: to prevent consequential dimensions of the person from disappearing between data capture and medication action.

Embedding Visibility into Pharmacy Workflows

Patient visibility becomes operational only when contextual information can affect pharmacy work. Medication reconciliation illustrates the distinction between information availability and meaningful use: electronic tools can support comparison of sources and discrepancy detection, but their effectiveness depends on the completeness of medication information, professional review, and integration into the surrounding workflow [17]. A populated medication list should not therefore be presumed reconciled or current.

Embedding additional context also creates attention costs. Workload, task complexity, and repeated alerts influence responses to clinical decision support [18]. Visibility functions should prioritize decision-critical uncertainty rather than produce undifferentiated prompts. Fragmented records and incompatible systems can further disrupt medicines optimisation across care interfaces, even when individual services hold useful information [19]. Continuity consequently requires both technical interoperability and shared responsibility for interpretation.

Human–AI interaction is equally important. In simulated medication verification, the presentation of model uncertainty affected pharmacists' decision processes and response times [20]. This does not establish real-world safety benefit, but it shows that information design may alter reliance. The proposed workflow therefore places model output beside, rather than above, medication evidence, patient context, and professional judgment.

Six workflow functions are proposed: proportionate context intake; provenance review; contextual reconciliation; patient clarification; documentation of context-to-action relationships; and follow-up updating. Decision-critical missingness, contradiction, staleness, or poorly understood uncertainty should trigger clarification, deferral, or escalation rather than automatic progression. Workflow integration must address repetitive alert burden, fragmentation of medicines information, and the way uncertainty is communicated to pharmacists [18-20]. These functions require task-specific validation and should not become a

uniform documentation requirement for every pharmacy encounter.

Equity and Patient-Experience Evaluation

Patient visibility may be distributed unequally. An apparently neutral prediction target can reproduce inequity when it uses healthcare expenditure or utilization as a proxy for underlying need [21]. In pharmacy services, analogous risks may arise when prescription collection, portal activity, claims, or recorded adherence are interpreted without considering unequal access and opportunity.

Digital-health equity extends beyond device ownership to the interaction of technological, social, structural, and health determinants across design, implementation, and continued use [22]. Evaluation should therefore examine who can contribute contextual information, whose information is considered credible, who can correct the record, and who bears the burden of participation.

Users' responses to AI-enabled decision aids are also shaped by trust, transparency, usability, human involvement, and compatibility with shared decision-making [23]. Patient-experience evaluation should measure perceived recognition, correction opportunities, explanation of data use, preference concordance, and participation burden. Equity evaluation must examine biased proxies, digital determinants of participation, and user experience rather than relying on aggregate model accuracy alone [21-23]. Observed group differences should prompt investigation; they should not automatically be attributed to technology, professional behaviour, or patient choice.

Implications for Digital-Service Design

Responsible design should begin with the medication decision, intended users, data provenance, foreseeable harms, validation context, and monitoring plan rather than with the availability of a model [24]. Early clinical evaluation should report workflow integration, user roles, human–AI interaction, errors, adaptations, and contextual conditions [25]. These requirements are especially important because computerized decision support produces heterogeneous improvements in care processes, making adoption or technical performance insufficient evidence of meaningful benefit [26].

Figure 2 presents the original conceptual synthesis for embedding narrative, preference, capability, and social context into workflow, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

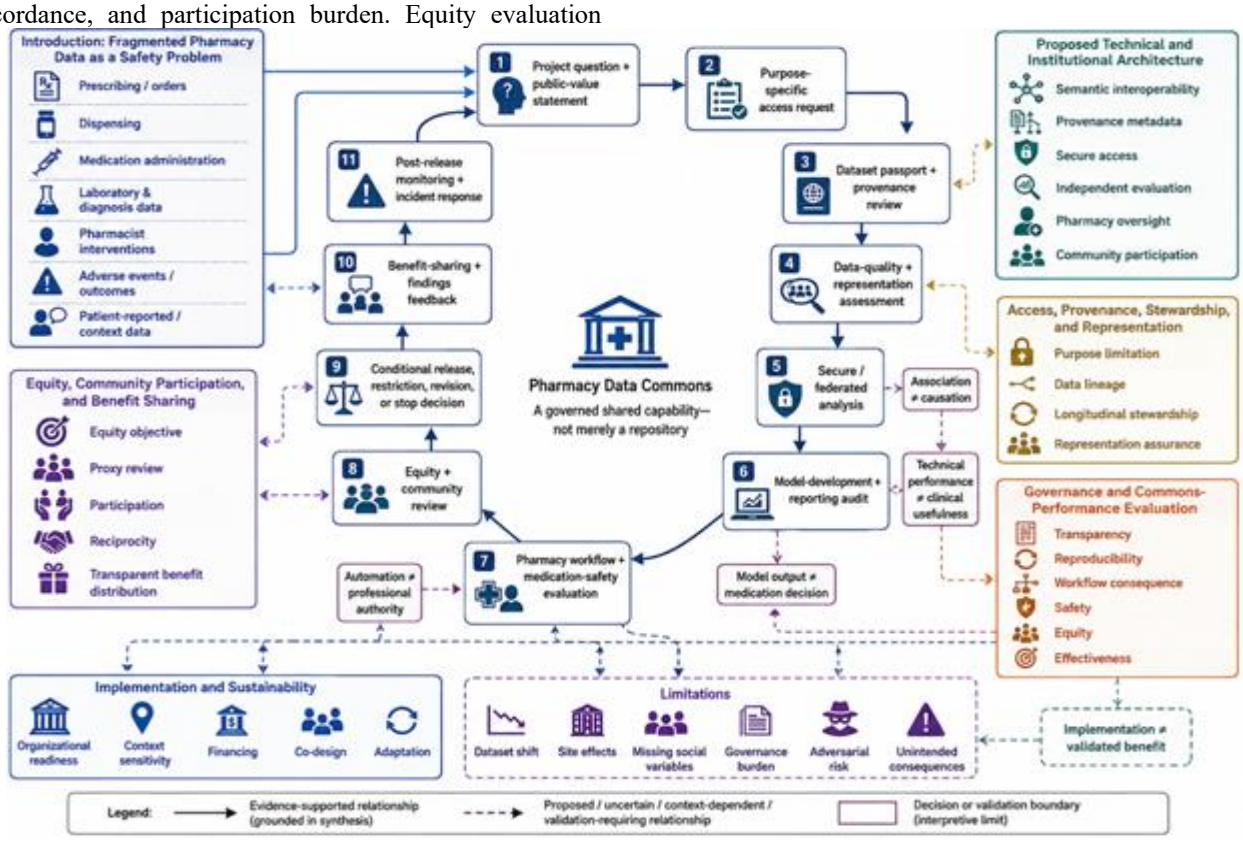

Figure 2. Embedding Narrative, Preference, Capability, and Social Context into Workflow

The figure proposes a context-discordance gate at which the record, patient account, professional interpretation, and system output are compared. Its use should be evaluated through construct testing, patient-confirmed representational fidelity, usability and workload assessment, medication-safety evaluation, equity analysis, and longitudinal

monitoring. Passing one evaluation stage should not imply that later stages or broader implementation requirements have been satisfied.

Limitations

The framework does not resolve the wider challenges of data quality, generalisability, workflow integration, human factors, monitoring, and local implementation that constrain the clinical impact of artificial intelligence [27]. Its categories may omit culturally specific, relational, emotional, or institutional dimensions.

Patient visibility should also not be equated with model explainability. Post-hoc explanations may be unreliable or misleading and cannot independently establish safety, fairness, or contextual validity [28]. Furthermore, changes in populations, technology, professional behaviour, and service settings can alter data relationships after implementation [29]. Narrative, preference, capability, and social-context representations may likewise become stale.

Additional contextual capture may introduce privacy, stigma, surveillance, workload, and responsibility risks. The framework should therefore be treated as a falsifiable organizing model whose usefulness and unintended consequences require comparative evaluation.

Conclusion

Data-driven pharmacy services risk replacing the patient with a technically useful but incomplete representation. The proposed Patient-Visibility Framework addresses this risk by distinguishing the person from the record and organizing contextual representation through narrative, preference, capability, and social-context layers. These layers are connected to provenance-aware capture, interpretation, pharmacist–patient dialogue, medication action, correction, follow-up, and accountable governance.

The framework preserves several boundaries. Data availability is not equivalent to understanding; model output is not a medication decision; technical performance is not clinical usefulness; explanation is not patient visibility; implementation is not validated benefit; and professional authority is not transferred automatically to an automated system. Visibility should be proportionate to the decision and should not justify unlimited contextual collection.

The framework's contribution is conceptual and evaluative. It identifies proposed failure modes, including patient–record substitution, narrative flattening, preference freezing, capability misclassification, visibility debt, unequal invisibility, workflow bypass, and responsibility diffusion. Empirical work must determine whether these constructs are valid, measurable, useful, and acceptable across pharmacy settings. Until such evidence exists, the framework should guide questions for design and evaluation rather than serve as a clinical protocol, performance score, or deployment standard.

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

References

1. 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
2. Zulman DM, Haverfield MC, Shaw JG, Brown-Johnson CG, Schwartz R, Tierney AA, et al. Practices to foster physician presence and connection with patients in the clinical encounter. JAMA. 2020;323(1):70-81. doi:10.1001/jama.2019.19003
3. Rathert C, Mittler JN, Banerjee S, McDaniel J. Patient-centered communication in the era of electronic health records: What does the evidence say? Patient Educ Couns. 2017;100(1):50-64. doi:10.1016/j.pec.2016.07.031
4. Shortliffe EH, Sepúlveda MJ. Clinical decision support in the era of artificial intelligence. JAMA. 2018;320(21):2199-200. doi:10.1001/jama.2018.17163
5. Santana MJ, Manalili K, Jolley RJ, Zelinsky S, Quan H, Lu M. How to practice person-centred care: A conceptual framework. Health Expect. 2018;21(2):429-40. doi:10.1111/hex.12640
6. Elwyn G, Durand MA, Song J, Aarts J, Barr PJ, Berger Z, et al. A three-talk model for shared decision making: Multistage consultation process. BMJ. 2017;359. doi:10.1136/bmj.j4891
7. Reading MJ, Merrill JA. Converging and diverging needs between patients and providers who are collecting and using patient-generated health data: An integrative review. J Am Med Inform Assoc. 2018;25(6):759-71. doi:10.1093/jamia/ocy006
8. Heckemann B, Chaaya M, Jakobsson Ung E, Olsson DS, Jakobsson S. Finding the person in electronic health records: A mixed-methods analysis of person-centered content and language. Health Commun. 2022;37(4):418-24. doi:10.1080/10410236.2020.1846275
9. 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
10. Greenhalgh T, Wherton J, Papoutsi C, Lynch J, Hughes G, A'Court C, et al. Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. J Med Internet Res. 2017;19(11). doi:10.2196/jmir.8775
11. 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
12. Cutillo CM, Sharma KR, Foschini L, Kundu S, Mackintosh M, Mandl KD, et al. Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency. NPJ Digit Med. 2020;3:47. doi:10.1038/s41746-020-0254-2
13. Adler-Milstein J, Nong P. Early experiences with patient generated health data: Health system and patient perspectives. J Am Med Inform Assoc. 2019;26(10):952-9. doi:10.1093/jamia/ocz045
14. Longoni C, Bonezzi A, Morewedge CK. Resistance to medical artificial intelligence. J Consum Res. 2019;46(4):629-50. doi:10.1093/jcr/ucz013
15. van Deursen AJAM, van Dijk JAGM. The first-level digital divide shifts from inequalities in physical access to inequalities in material access. New Media Soc. 2019;21(2):354-75. doi:10.1177/1461444818797082
16. Cantor MN, Thorpe L. Integrating data on social determinants of health into electronic health records. Health Aff (Millwood). 2018;37(4):585-90. doi:10.1377/hlthaff.2017.1252

17. Marien S, Krug B, Spinewine A. Electronic tools to support medication reconciliation: A systematic review. J Am Med Inform Assoc. 2017;24(1):227-40. doi:10.1093/jamia/ocw068
18. Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. BMC Med Inform Decis Mak. 2017;17(1):36. doi:10.1186/s12911-017-0430-8
19. Tolley C, Seymour H, Watson N, Nazar H, Heed J, Belshaw D. Barriers and opportunities for the use of digital tools in medicines optimization across the interfaces of care: Stakeholder interviews in the United Kingdom. JMIR Med Inform. 2023;11. doi:10.2196/42458
20. Lester C, Rowell B, Zheng Y, Co Z, Marshall V, Kim JY, et al. Effect of uncertainty-aware AI models on pharmacists' reaction time and decision-making in a web-based mock medication verification task: Randomized controlled trial. JMIR Med Inform. 2025;13. doi:10.2196/64902
21. 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
22. Richardson S, Lawrence K, Schoenthaler AM, Mann D. A framework for digital health equity. NPJ Digit Med. 2022;5:119. doi:10.1038/s41746-022-00663-0
23. Hassan N, Slight R, Bimpong K, Bates DW, Weiand D, Vellinga A, et al. Systematic review to understand users perspectives on AI-enabled decision aids to inform shared decision making. NPJ Digit Med. 2024;7:332. doi:10.1038/s41746-024-01326-y
24. 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
25. 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
26. 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
27. 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
28. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9
29. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med. 2021;385(3):283-6. doi:10.1056/NEJMc2104626