Original Article

When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture

Pieter Botha[1]*, Anel Van Wyk[1], Johan Marais[2]

[1]Department of Medication Record Reconciliation and AI, Faculty of Pharmacy, Stellenbosch University, Stellenbosch, South Africa.
[2]Department of Multi-Agent Systems in Pharmacy, Faculty of Pharmacy, University of Pretoria, Pretoria, South Africa.

Abstract

Medication reconciliation is commonly represented as the production of one accurate medication list from several incomplete records. This representation obscures a more difficult problem: medication records are source-specific accounts produced for different purposes, at different times, and under different documentation conditions. Prescribing records, dispensing histories, administration records, clinical notes, and patient reports may therefore disagree without any source being uniformly authoritative. This article proposes a multi-agent medication-record reconciliation architecture that treats each record entry as a temporally situated assertion rather than an unquestioned fact. Bounded source-specific agents extract medication assertions while preserving provenance, documentation time, inferred validity intervals, uncertainty, and source limitations. Shared services normalize medication identity, classify disagreement, assess evidence suitability, and construct conflict sets without silently deleting competing assertions. A reconciliation orchestrator generates provisional medication states, while an independent challenge function examines unsupported assumptions and correlated agreement. Conflicts exceeding configured authority or evidentiary sufficiency are transferred to structured human adjudication, external verification, or safety escalation. The proposed architecture distinguishes semantic, temporal, intentional, provenance, omission, and uncertainty disagreements and separates technical reconciliation performance from workflow usefulness, medication safety, equity, and patient outcomes. Validation would require staged component testing, multisite retrospective assessment, silent prospective evaluation, human-factors studies, comparative workflow evaluation, subgroup analysis, clinical-outcome investigation, and postdeployment monitoring. The architecture is an original non-empirical synthesis. It does not establish that multi-agent reconciliation is superior to conventional software or pharmacist-led practice, and it should not be interpreted as a clinically validated, autonomous, or deployment-ready system.

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

Introduction

Medication Records as Conflicting Accounts

Medication reconciliation is often described as comparing medication lists, identifying discrepancies, and establishing the medicines a patient should receive. That formulation can imply that disagreement is an accidental defect surrounding an otherwise stable medication truth. In practice, the medication record is distributed across prescribing systems, dispensing histories, medication-administration records, clinical documentation, patient portals, and patient or caregiver accounts. Even information contained within one electronic health record may be incomplete or inconsistent across its medication-related fields [1]. The reconciliation problem is therefore not simply duplication removal. It is the interpretation of multiple assertions generated under different clinical, administrative, and temporal conditions.

Patients and caregivers add another evidentiary perspective. They may identify erroneous medications, incorrect doses, obsolete instructions, or misunderstandings that are not visible through record-to-record comparison [2]. Their reports are not automatically definitive, but neither should they be reduced to secondary commentary on a professionally generated list. A patient statement that a medicine was stopped, taken differently, never started, or obtained elsewhere may conflict with prescribing or dispensing data because the sources describe different stages of the medication-use process. Reconciliation must preserve that distinction before deciding whether the disagreement reflects

Address for correspondence: Pieter Botha, Department of Medication Record Reconciliation and AI, Faculty of Pharmacy, Stellenbosch University, Stellenbosch, South Africa.
pieter.botha@sun.ac.za
Received: 28 October 2025; **Accepted:** 30 January 2026

How to cite this article: Botha P, Van Wyk A, Marais J. When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture. Arch Pharm Pract. 2026;17(1):22-30. https://doi.org/10.51847/zrzLpexlrL

error, time, intention, adherence, access, or incomplete documentation.

The consequences of this conceptual distinction are important. A system may reduce documented discrepancies without demonstrating that it reduces adverse drug events, treatment interruption, or other patient-level harms [3]. Discrepancy closure is therefore a process outcome rather than sufficient proof of safer medication use. A reconciliation mechanism may also create new risks by accepting stale information, suppressing legitimate dissent, presenting uncertain outputs as verified facts, or transferring excessive authority to an automated recommendation.

Electronic tools have supported list comparison, data retrieval, discrepancy identification, and workflow documentation, but their functions, evaluation designs, and outcome evidence remain heterogeneous [4]. Most existing approaches operate as individual applications or workflow aids rather than as an explicit architecture for reasoning over conflicting sources. They rarely make clear how source-specific limitations, temporal changes, provenance, evidentiary suitability, agent disagreement, human authority, and escalation should interact.

This article therefore proposes a multi-agent reconciliation architecture in which the medication record is treated as a collection of contestable assertions. The architecture does not assume that one source is consistently authoritative or that computational agreement establishes truth. Its purpose is to preserve evidence, organize disagreement, support accountable adjudication, and define the validation required before stronger claims can be made. The proposal is conceptual and non-empirical; it does not establish clinical effectiveness, safety, regulatory acceptance, or readiness for autonomous medication decision-making.

Taxonomy of Medication-Record Disagreement

A medication discrepancy can involve omission, addition, dose, frequency, route, formulation, timing, or continuation status. Validated discrepancy taxonomies provide a foundation for classifying such differences systematically [5]. However, a list of discrepancy types does not fully explain why two records disagree or whether they are genuinely incompatible. The proposed taxonomy therefore treats disagreement as multidimensional. A medication assertion is defined as a source-specific statement about medication identity, strength, dose, route, frequency, status, intention, administration, dispensing, or reported use during an explicit or inferred time interval.

A semantic disagreement occurs when sources may refer to the same therapy using different names, codes, formulations, units, strengths, or levels of specificity. Some semantic differences disappear after justified normalization; others reveal clinically meaningful distinctions that must remain separate. A temporal disagreement occurs when apparently conflicting assertions describe different periods, such as a medication stopped after discharge but retained on an unrevised active list. A medication-intent disagreement distinguishes what was prescribed, dispensed, administered, recommended, intended, or actually taken. Community prescribing and dispensing data can broaden reconciliation evidence, but electronic integration does not by itself establish current use, prescriber intention, or patient understanding [6].

A provenance disagreement arises when the origin, author, transformation, verification status, or copying history of an assertion is unclear. An omission disagreement occurs when the absence of a medication is potentially meaningful relative to the expected coverage of a source. Absence from a dispensing record, for example, is different from documented discontinuation, while absence from a patient report may reflect recall limitations, nonadherence, or misunderstanding. Patient participation can expose discrepancies that institutional records do not resolve, particularly when actual use diverges from documented instructions [7]. Patient-generated assertions should therefore remain visible with their own provenance and uncertainty rather than being automatically subordinated or accepted.

An uncertainty disagreement exists when evidence cannot responsibly distinguish among competing medication states. This category prevents premature resolution merely because one assertion is easier to process, more recent in the interface, or repeated across copied records. Repetition is not necessarily independent corroboration.

Finally, disagreement is embedded in workflow. Medication reconciliation depends on role allocation, information availability, communication, implementation support, and local practice rather than on software matching alone [8]. A conflict may therefore result from delayed documentation, fragmented care, ambiguous authority, inaccessible external information, or unresolved communication between professionals and patients. The proposed taxonomy organizes these mechanisms but does not determine the clinical importance of an individual discrepancy.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture.

Table 1. Definitions, components, relationships, and intended uses for the article When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Medication assertion	Records are often treated as interchangeable facts despite different origins and purposes.	Medication identity, dose, route, frequency, status, source, author, and time.	**Proposed:** represent every entry as a source-bounded assertion. The need follows from observed medication-list inconsistency [1].	Preserves the distinction between recorded information and verified medication state.	Extraction fidelity and expert review across source types.	Incorrect extraction or loss of source context.	An assertion is not equivalent to clinical truth.
Semantic disagreement	Equivalent or distinct therapies may be confused because of coding, terminology, formulation, or unit variation.	Names, codes, strengths, dosage forms, routes, units, and mappings.	**Proposed:** normalize terminology while retaining the original representation.	Reduces false conflict without erasing clinically meaningful differences.	Terminology-validation studies and expert adjudication.	False merging or false separation.	Normalization does not authorize a medication change.
Temporal disagreement	Old and current assertions may appear simultaneously valid.	Event time, documentation time, order time, dispensing time, administration time, and inferred validity interval.	**Proposed:** compare assertions only after constructing temporal relationships.	Distinguishes change over time from true contemporaneous conflict.	Temporally annotated records and longitudinal expert review.	Stale therapy accepted as current or valid therapy rejected as obsolete.	Recency alone does not establish correctness.
Medication-intent disagreement	Prescribed, dispensed, administered, intended, and taken therapy may differ.	Orders, dispensing events, administration events, notes, and patient reports.	**Proposed:** encode medication-process function for each assertion. Electronic source integration provides relevant but incomplete evidence [6].	Makes the reason for disagreement explicit.	Cross-source validation and confirmation of intended use.	One process stage is mistaken for another.	Dispensing does not prove ingestion; prescribing does not prove continuation.
Provenance envelope	Copied, transformed, or unattributed entries can appear independently confirmed.	Source system, author, extraction method, copied-from relationship, verification status, and transformation history.	**Proposed:** attach immutable provenance metadata to each assertion.	Supports traceability and detection of dependent evidence.	Provenance completeness and reconstruction testing.	Repeated copies mistaken for independent corroboration.	More appearances do not necessarily mean stronger evidence.
Omission disagreement	A missing medication may represent true absence, incomplete coverage, or failed documentation.	Expected source coverage, access, reporting behavior, and care setting.	**Proposed:** interpret absence relative to what the source can reasonably contain.	Prevents absence from being converted automatically into discontinuation.	Source-coverage studies and subgroup analyses.	False discontinuation or missed omission.	Missingness is not evidence of non-use without contextual support.
Patient or caregiver assertion	Actual use and patient understanding may be absent from institutional records.	Patient report, caregiver report, medication containers, routines, and barriers.	**Proposed:** retain patient evidence as a distinct source requiring respectful verification. Patient engagement can identify unresolved discrepancies [7].	Improves visibility of use outside institutional documentation.	Usability, concordance, equity, and communication studies.	Patient evidence ignored, misunderstood, or accepted without verification.	Patient reports are evidentiary inputs, not automatically definitive answers.
Discrepancy class	A surface difference does not reveal its underlying mechanism.	Semantic, temporal, intent, provenance, omission, and uncertainty features.	**Proposed extension:** assign one or more disagreement dimensions using established discrepancy classification as a foundation [5].	Supports more specific reasoning and evaluation.	Interrater reliability, construct validity, and clinical-relevance testing.	Overlapping or unstable classes.	The proposed extension has not been validated.

Conflict set	Competing assertions may be silently overwritten during list consolidation.	All assertions concerning a medication concept and relevant interval.	**Proposed:** preserve competing assertions, support, dissent, and unresolved questions as one object.	Improves auditability and prevents hidden conflict deletion.	Reconstruction studies and comparison with existing reconciliation processes.	Conflict proliferation or inappropriate consolidation.	Preserving conflict does not imply that all assertions are equally credible.
Provisional reconciled medication state	Conventional outputs may appear final even when uncertainty remains.	Supported assertions, adjudication decisions, unresolved conflict, authority, and follow-up.	**Proposed:** produce a time-bounded state with explicit confidence and unresolved items.	Communicates what is accepted, contested, deferred, or escalated.	Prospective comprehension and workflow evaluation.	Provisional output interpreted as certified truth.	The state is decision support and documentation, not autonomous prescribing authority.

Proposed Multi-Agent Reconciliation Architecture

Medication-reconciliation automation encompasses separable tasks, while contemporary healthcare-agent systems provide concepts for bounded roles, tool use, communication, and role-specialized deliberation [9–11]. The proposed architecture applies these concepts without assuming that autonomous agents should determine a patient's medication regimen. Its central design principle is bounded epistemic responsibility: each agent is accountable for a defined information-processing function, but no individual agent is authorized to convert its output directly into a clinical medication decision.

The architecture begins with source-specific agents. Separate agents interface with prescribing records, dispensing or claims data, medication-administration records, clinical documentation, active medication lists, and patient or caregiver reports. Each agent extracts assertions using rules and models appropriate to its source. It also records missingness, expected coverage, access limitations, and transformation history. Source specialization is intended to localize error: a failure can be attributed to a particular source interface or extraction process rather than disappearing inside one undifferentiated model.

A shared medication-identity normalization service maps terminology, strengths, dosage forms, routes, and units while retaining original source expressions. A provenance and temporal-envelope service records who or what generated an assertion, when the underlying event occurred, when it was documented, and the interval during which it may have been valid. A discrepancy-classification agent then constructs conflict sets according to semantic, temporal, intent, provenance, omission, and uncertainty dimensions.

An evidence-suitability agent characterizes how useful each assertion is for the question under consideration. It does not assign permanent authority to any source. A pharmacy dispensing record may be highly informative about supply but insufficient to establish ingestion, while a patient account may describe actual use but leave dose or timing uncertain. Suitability is therefore assertion-specific and decision-specific.

A reconciliation orchestrator integrates these structured outputs and generates candidate medication states. It is constrained from silently deleting conflicting assertions or treating repeated copied entries as independent confirmation. A separate challenge agent searches for unsupported assumptions, missing sources, inconsistent timelines, and conclusions that depend on correlated evidence. Traceable clinical-agent research supports the importance of linking outputs to source evidence and intermediate reasoning [12]. In this architecture, traceability includes the original assertion, transformations, temporal interpretation, evidence assessment, agent dissent, and authority used for any resolution.

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

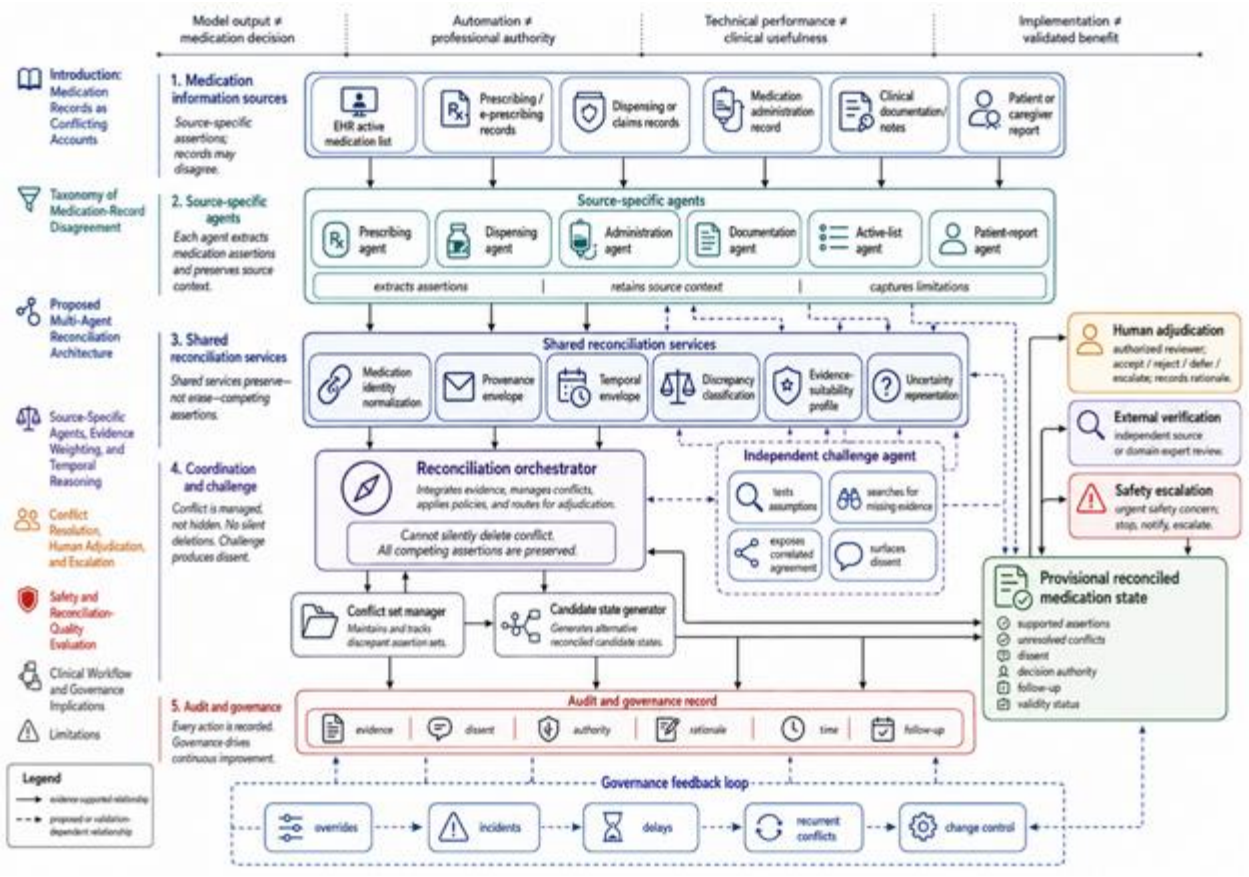

Figure 1. Multi-Agent Medication-Record Reconciliation Architecture

The architecture produces a provisional reconciled medication state, not an asserted ground truth. The state identifies supported assertions, rejected interpretations, unresolved conflicts, dissenting evidence, decision authority, validity period, and required follow-up. Conflicts outside configured authority or evidentiary sufficiency are transferred to human adjudication, external verification, or escalation. Every operation is written to an audit and change log.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article When the Medication Record Disagrees with Itself: A Multi-Agent Reconciliation Architecture.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Source-specific agents	Extract source-bounded medication assertions.	Source record to structured assertion plus limitations.	Supplies normalization, temporal, and conflict services.	Fidelity to the source and correct representation of missingness.	Omitted, altered, or fabricated assertion.	Approve source access, validate adapters, and review source-specific errors.	Source extraction performance does not establish reconciliation accuracy.
Identity normalization	Align medication concepts, formulations, strengths, routes, and units.	Original expression to normalized concept while preserving source text.	Precedes conflict classification but cannot resolve clinical intent.	Correct equivalence and separation decisions.	False merge or false distinction.	Maintain terminology services and review clinically consequential mappings.	Terminology alignment is not medication authorization.

Provenance and temporal envelope	Preserve origin, transformation, event time, documentation time, and validity.	Assertion metadata to an auditable longitudinal representation.	Informs suitability, conflict classification, and challenge.	Completeness and reconstructability of provenance and time.	Copied records treated as independent or stale assertions treated as current.	Define provenance requirements and resolve unavailable metadata.	Complete metadata cannot guarantee factual correctness.
Discrepancy-classification agent	Characterize the dimensions of disagreement.	Normalized assertions to one or more conflict classes.	Structures orchestrator and adjudicator review.	Agreement with expert classification and stability across settings.	Incorrect or overly narrow conflict class.	Establish classification governance and clinical-review procedures.	Classification does not determine clinical significance.
Evidence-suitability agent	Characterize relevance and limitations for a specific reconciliation question.	Source and context descriptors to an assertion-specific profile.	Informs orchestration without imposing permanent source authority.	Calibration, consistency, and context appropriateness.	Universal source ranking or hidden weighting bias.	Approve weighting dimensions and monitor subgroup effects.	A suitability profile is not a probability that an assertion is true.
Reconciliation orchestrator	Construct candidate medication states without erasing conflict.	Conflict sets and suitability profiles to provisional candidates.	Receives challenge-agent critique and authority constraints.	Traceability, consistency, and appropriate preservation of unresolved evidence.	Premature closure or correlated evidence mistaken for corroboration.	Configure permissible operations and review recurring errors.	Candidate generation is not autonomous medication decision-making.
Independent challenge agent	Test assumptions, search for missing evidence, and expose dissent.	Candidate state and evidence trace to challenge findings.	Returns critique to the orchestrator or escalates uncertainty.	Added detection of unsupported resolution without excessive noise.	Shared model error or ceremonial challenge with no independent value.	Ensure functional independence and review challenge failures.	Agent disagreement is informative but not inherently correct.
Human adjudication interface	Support accountable acceptance, rejection, deferral, or escalation.	Competing assertions, evidence, dissent, and consequence to an authorized reviewer.	Updates the provisional state and audit log.	Comprehension, reproducibility, workload, and resistance to automation bias.	Rubber-stamp acceptance or excessive cognitive burden.	Assign authority, training, time, and escalation support.	Human presence alone does not guarantee safety.
Escalation controller	Route unresolved or consequential conflict to appropriate action.	Uncertainty, consequence, authority, and time sensitivity to verification or escalation.	Coordinates pharmacy, prescriber, patient, caregiver, and safety processes.	Appropriate routing and timely completion.	Under-escalation, over-escalation, or unowned follow-up.	Define escalation ownership and monitor completion.	Escalation thresholds require local validation.
Audit and governance loop	Preserve decisions and support monitoring, change control, and incident review.	Source evidence and actions to immutable history and governance feedback.	Spans every architecture layer.	Reconstruction of decisions, overrides, changes, and incidents.	Untraceable action or uncontrolled system change.	Maintain audit retention, incident review, and rollback authority.	Auditability supports accountability but does not establish clinical benefit.

The architecture is intentionally modular. Alternative implementations could combine or separate functions differently, and conventional software may perform some components more appropriately than generative or agentic models. The conceptual claim is not that every reconciliation task requires an intelligent agent. Rather, any automated component should have a bounded role, visible evidentiary basis, explicit authority, and testable failure conditions.

Source-Specific Agents, Evidence Weighting, and Temporal Reasoning

Source-specific reasoning depends first on interoperability. Medication identity, route, formulation, status, and timing cannot be compared reliably when systems exchange syntactically available but semantically incompatible data [13]. Interoperability in the proposed architecture therefore includes terminology alignment, source identifiers, medication-process meaning, time representation, and preservation of original values. It is a prerequisite for comparison, not proof that the exchanged information is complete or correct.

Each source agent should produce both medication assertions and a source-quality profile. Electronic health record quality is multidimensional, involving completeness, correctness, concordance, plausibility, and currency rather than one universal quality score [14]. The architecture extends this principle by requiring agents to state what a source is expected to contain, what is unavailable, and which transformations were applied. A medication-administration record may be strong evidence that a dose was documented

as administered in a facility but weak evidence of later outpatient continuation. A dispensing history may indicate supply but not ingestion. A copied active-list entry may be recent in the interface but old in clinical meaning.

Temporal reasoning is central because medication disagreement is often disagreement about time. Clinical text and structured records can contain event order, duration, and temporal relations that require explicit extraction and interpretation [15]. The proposed temporal envelope therefore separates event time, documentation time, system-entry time, and assertion validity interval. It also records whether time was directly stated, inferred from context, or unresolved. A discharge instruction entered after a prescription was discontinued should not be interpreted solely according to database insertion order. Similarly, a medication mentioned historically in a note should not be treated as current merely because the note is recent.

Evidence weighting should be renamed evidence-suitability assessment to avoid implying that one numerical score can determine medication truth. Data suitability is conditional on the purpose for which information is being used [16]. The proposed assessment therefore asks whether an assertion is relevant to the specific reconciliation question, sufficiently direct, temporally applicable, independently supported, and produced by a source expected to observe the medication-process event at issue.

Suitability dimensions include source identity, directness of observation, process function, temporal proximity, internal consistency, cross-source corroboration, transformation history, completeness, plausibility, expected coverage, and patient confirmation. These dimensions should not be assumed to compensate freely for one another. Strong semantic matching cannot repair absent provenance, and repeated copied entries should not offset a direct contemporaneous contradiction. Potential medication consequence must also influence the required evidentiary standard: uncertainty that might produce duplicate anticoagulation, omitted insulin, or inappropriate continuation warrants a different response from a minor historical-list ambiguity. This principle identifies the need for escalation; it does not supply a validated clinical threshold.

The orchestrator should retain the evidence-suitability profile alongside each candidate state. Where evidence is insufficient, the appropriate output is not forced agreement but a documented uncertainty state with a verification action. Candidate states may be compared, but the architecture should not present the highest-scoring option as automatically correct. Weighting models could reproduce documentation bias, access inequalities, or local workflow assumptions, particularly when some patients have fewer connected records or less opportunity to correct errors.

Validation must therefore evaluate the source agents separately and jointly. Required tests include extraction fidelity, terminology equivalence, provenance completeness, temporal ordering, conflict detection, calibration of uncertainty, and sensitivity to missing or dependent evidence. Testing should span institutions, record systems, care transitions, medication classes, and source-availability patterns. Until these relationships are empirically calibrated, the proposed weighting dimensions remain structured prompts for reasoning rather than validated probabilities, clinical rules, or autonomous decision thresholds.

Conflict Resolution, Human Adjudication, and Escalation

Conflict resolution should not mean selecting the most frequent, recent, or computationally confident assertion. The proposed architecture first determines whether disagreement can be explained by semantic normalization or temporal sequence. Remaining conflicts enter an adjudication packet containing competing assertions, provenance, temporal context, evidence suitability, uncertainty, potential consequences, and agent dissent. Human–computer collaboration may improve some clinical tasks, but its value depends on task structure and interaction design [17]. Human review must therefore be treated as an accountable decision process rather than an automatic safety mechanism.

Reviewers can accept erroneous computational recommendations, particularly when advice is presented confidently or appears institutionally endorsed [18]. The interface should consequently preserve independent evidence, expose disagreement, and require the adjudicator to record whether an assertion was accepted, rejected, deferred, or referred for verification.

Escalation should reflect potential consequence, unresolved uncertainty, time sensitivity, and professional authority. Organizational complexity, workflow conditions, and local responsibilities affect whether technologies are adopted and sustained [19]. Alert burden, interruption, and inappropriate overrides must also be considered [20]. Proposed resolution states are supported agreement, semantically normalized difference, temporally explained difference, clinically adjudicated difference, deferred verification, urgent safety escalation, and persistent uncertainty. These states and thresholds require prospective validation and must not be interpreted as clinical rules.

Safety and Reconciliation-Quality Evaluation

Evaluation should progress from responsible problem and data formulation through early live clinical assessment to later comparative evaluation capable of supporting stronger clinical claims [21–23]. Technical testing should examine extraction fidelity, normalization, provenance completeness, temporal ordering, discrepancy classification, uncertainty calibration, and preservation of unresolved evidence. Human-factors testing should evaluate comprehension, workload, adjudication reproducibility, automation bias, and responses to agent dissent.

Workflow evaluation should assess interruptions, duplicate work, unresolved tasks, communication failures, and escalation completion. Safety evaluation should distinguish missed consequential discrepancies from system-induced changes, delays, and false escalations. High technical performance cannot exclude unintended clinical consequences [24]. Equity analyses should examine performance and escalation across patients with unequal documentation continuity, pharmacy access, portal participation, language support, and connected data sources. Claims of reduced medication harm would require comparative prospective outcome evidence; no earlier validation stage could substitute for that evidence.

Clinical Workflow and Governance Implications

Implementation requires explicit accountability for source access, agent configuration, adjudication, overrides, follow-up, and resulting medication actions. Ethical implementation also requires attention to privacy, consent, bias, and responsibility [25]. Governance should extend across authorization, monitoring, incident review, change control, restriction, rollback, and retirement [26].

Explanation displays may support inspection, but they cannot certify that an agent's output is correct or safe [27]. Users should instead receive structured information about intended use, source coverage, known limitations, prohibited interpretations, and required escalation conditions [28].

Figure 2 presents the original conceptual synthesis for source conflict, temporal reasoning, human adjudication, and escalation, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

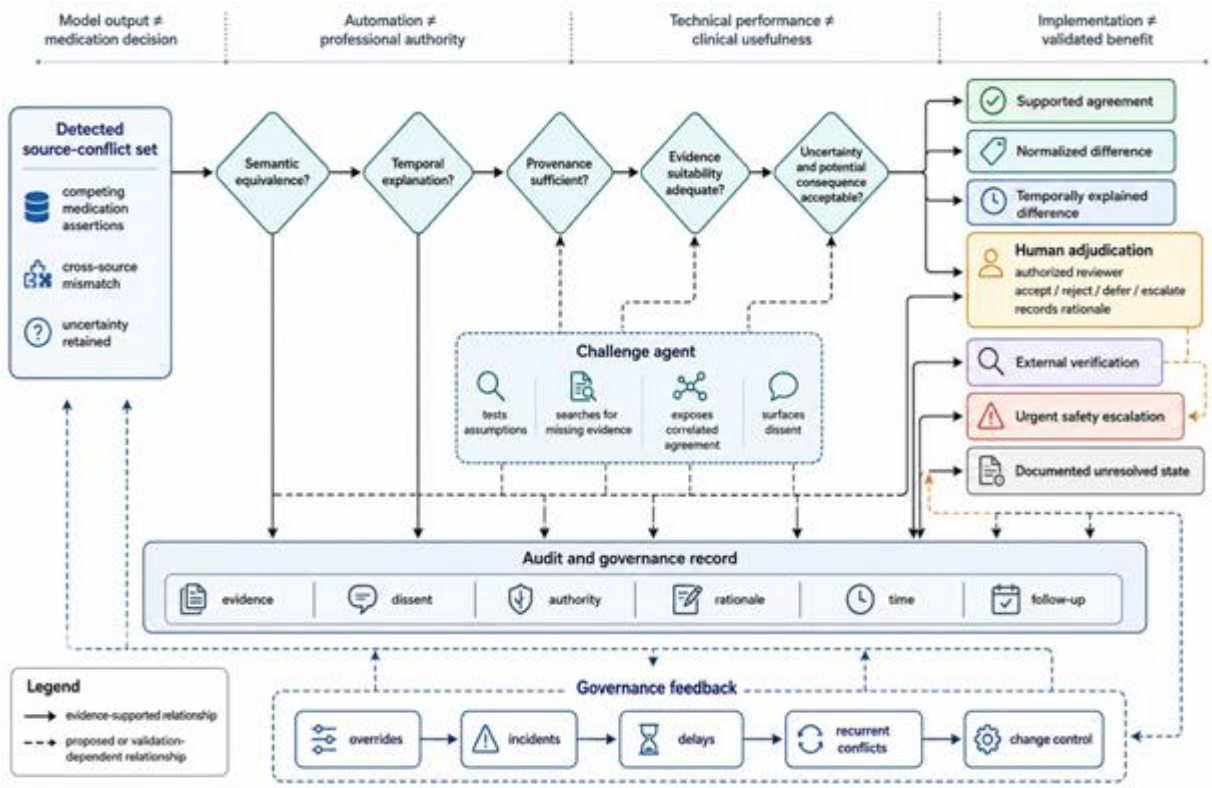

Figure 2. Source Conflict, Temporal Reasoning, Human Adjudication, and Escalation

Limitations

Reconciliation agents may reproduce inequity through inappropriate proxies, unequal access to recorded care, nonrandom missingness, and documentation practices embedded in electronic health data [29, 30]. The architecture also assumes that source roles, temporal representations, evidence-suitability dimensions, and escalation responsibilities can be configured consistently, although these may vary across institutions and jurisdictions. External validation, workflow integration, prospective comparison, and postdeployment monitoring remain necessary before clinical-impact or transportability claims can be made [31]. The proposal should therefore be interpreted as a research architecture, not as evidence of safety, effectiveness, or deployment readiness.

CONCLUSION

Medication-record reconciliation is not merely the consolidation of lists. It is a governed interpretation of conflicting assertions generated by different people, systems, purposes, and times. The proposed architecture preserves source provenance, temporal context, uncertainty, and dissent through bounded source agents, shared reasoning services,

orchestration, independent challenge, human adjudication, escalation, and audit.

Its principal contribution is a structured way to distinguish record agreement from evidentiary sufficiency and provisional reconciliation from verified medication truth. Its value remains conditional on technical, human-factors, workflow, safety, equity, outcome, and governance validation. It does not replace pharmacists, prescribers, patients, or existing reconciliation responsibilities, and it does not establish that agentic automation is preferable to simpler systems. The architecture is intended to make disagreement visible, traceable, and accountable while defining the evidence required before stronger clinical claims could responsibly be considered.

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

References

1. Walsh KE, Marsolo KA, Davis C, Todd T, Martineau B, Arbaugh C, et al. Accuracy of the medication list in the electronic health record—implications for care, research, and improvement. J Am Med Inform Assoc. 2018;25(7):909-12. doi:10.1093/jamia/ocy027
2. Bell SK, Delbanco T, Elmore JG, Fitzgerald PS, Fossa A, Harcourt K, et al. Frequency and types of patient-reported errors in electronic health record ambulatory care notes. JAMA Netw Open. 2020;3(6). doi:10.1001/jamanetworkopen.2020.5867
3. Tamblyn R, Abrahamowicz M, Buckeridge DL, Bustillo M, Forster AJ, Girard N, et al. Effect of an electronic medication reconciliation intervention on adverse drug events: A cluster randomized trial. JAMA Netw Open. 2019;2(9). doi:10.1001/jamanetworkopen.2019.10756
4. 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
5. Almanasreh E, Moles R, Chen TF. The medication discrepancy taxonomy (MedTax): The development and validation of a classification system for medication discrepancies identified through medication reconciliation. Res Social Adm Pharm. 2020;16(2):142-8. doi:10.1016/j.sapharm.2019.04.005
6. Tamblyn R, Winslade N, Lee TC, Motulsky A, Meguerditchian AN, Bustillo M, et al. Improving patient safety and efficiency of medication reconciliation through the development and adoption of a computer-assisted tool with automated electronic integration of population-based community drug data: The RightRx project. J Am Med Inform Assoc. 2018;25(5):482-95. doi:10.1093/jamia/ocx107
7. Prey JE, Polubriaginof F, Grossman LV, Masterson Creber R, Tsapepas D, Perotte R, et al. Engaging hospital patients in the medication reconciliation process using tablet computers. J Am Med Inform Assoc. 2018;25(11):1460-9. doi:10.1093/jamia/ocy115
8. Schnipper JL, Mixon AS, Stein J, Wetterneck TB, Kaboli PJ, Mueller SK, 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
9. Tabja Bortesi JP, Becerra MP, Ranisau J, Wen B, Nadesan P, Devereaux PJ, et al. AI-based automation for medication reconciliation: Scoping review. J Med Internet Res. 2026;28. doi:10.2196/86760
10. Xu G, Li X, Chen Y, Duan Y, Wu S, Yu H, et al. A comprehensive survey of AI agents in healthcare. J Biomed Inform. 2026;179:105045. doi:10.1016/j.jbi.2026.105045
11. Chen X, Yi H, You M, Liu W, Wang L, Li H, et al. Enhancing diagnostic capability with multi-agents conversational large language models. npj Digit Med. 2025;8:159. doi:10.1038/s41746-025-01550-0
12. Zhao W, Wu C, Fan Y, Qiu P, Zhang X, Sun Y, et al. An agentic system for rare disease diagnosis with traceable reasoning. Nature. 2026;651(8106):775-84. doi:10.1038/s41586-025-10097-9
13. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. npj Digit Med. 2019;2:79. doi:10.1038/s41746-019-0158-1
14. Lewis AE, Weiskopf N, Abrams ZB, Foraker R, Lai AM, Payne PRO, et al. Electronic health record data quality assessment and tools: A systematic review. J Am Med Inform Assoc. 2023;30(10):1730-40. doi:10.1093/jamia/ocad120
15. Alfattni G, Peek N, Nenadic G. Extraction of temporal relations from clinical free text: A systematic review of current approaches. J Biomed Inform. 2020;108:103488. doi:10.1016/j.jbi.2020.103488
16. Shang N, Weng C, Hripcsak G. A conceptual framework for evaluating data suitability for observational studies. J Am Med Inform Assoc. 2018;25(3):248-58. doi:10.1093/jamia/ocx095
17. Tschandl P, Rinner C, Apalla Z, Argenziano G, Codella N, Halpern A, et al. Human–computer collaboration for skin cancer recognition. Nat Med. 2020;26(8):1229-34. doi:10.1038/s41591-020-0942-0
18. 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;4:31. doi:10.1038/s41746-021-00385-9
19. 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
20. 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
21. 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
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. 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. Lancet Digit Health. 2020;2(10). doi:10.1016/S2589-7500(20)30218-1
24. 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
25. Char DS, Shah NH, Magnus D. Implementing machine learning in health care—addressing ethical challenges. N Engl J Med. 2018;378(11):981-3. doi:10.1056/NEJMp1714229
26. 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
27. 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
28. 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
29. Obermeyer Z, Powers BW, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. doi:10.1126/science.aax2342
30. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. JAMA Intern Med. 2018;178(11):1544-7. doi:10.1001/jamainternmed.2018.3763
31. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17:195. doi:10.1186/s12916-019-1426-2