Original Article

Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph

Mateo Alvarez[1]*, Sofia Herrera[2], Diego Cruz[1], Andres Castro[3]

[1]Department of Pharmacy Knowledge Graphs and Reasoning, Faculty of Pharmacy, National University of Colombia, Bogota, Colombia. [2]Department of Medication Advice Representation, Faculty of Pharmacy, University of Chile, Santiago, Chile. [3]Department of AI Provenance and Pharmacy Knowledge, Faculty of Pharmacy, University of Antioquia, Medellin, Colombia.

Abstract

Medication advice is often represented digitally as an isolated recommendation, alert, or relation between a medicine and an action. Such representations may omit why the advice exists, when it applies, which patient or setting qualifies it, and which source or evidence supports it. This article proposes an assertion-centred pharmacy knowledge-graph architecture for representing medication advice as a context-bearing, temporally bounded, provenance-linked knowledge object. The central construct, the Medication Advice Assertion, connects medication identity, action, rationale, intended clinical goal, applicability conditions, temporal scope, source version, evidence item, evidence-strength assessment, contradiction state, responsible agent, and intended-use boundary. The architecture separates extracted associations from normative advice, evidence strength from model confidence, provenance from truth, and technical performance from clinical usefulness. It also proposes qualified relations among assertions, including support, refinement, exception, contradiction, and supersession, so that disagreement is preserved rather than silently overwritten. Evaluation is organized as a sequence of semantic-conformance, identity-resolution, provenance, temporal, applicability, contradiction, technical, workflow, safety, equity, and governance checks. Clinical use would additionally require pharmacist and stakeholder assessment, external validation, human-factors testing, monitoring, and controlled maintenance. The proposed architecture does not establish the correctness of any medication recommendation, clinical benefit, regulatory acceptance, or deployment readiness. Its original contribution is a representational and governance framework for making the why, when, and source of medication advice independently inspectable while preserving professional authority and uncertainty.

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

Introduction

Biomedical knowledge graphs can integrate heterogeneous entities and relations across scientific and clinical sources, but graph construction alone does not determine whether a relation is suitable for medication advice [1]. A relation such as “drug treats condition” may summarize curated evidence, a statistical association, an extracted statement, or an inferred hypothesis. When the intended use is pharmacy practice, these forms are not interchangeable. Advice requires an explicit action, a reason for that action, a defined recipient or population, an applicable context, a temporal scope, and a traceable source.

Graphs learned from electronic health records illustrate this distinction. Relations recovered from clinical documentation may reflect recorded co-occurrence, local practice patterns, or extraction error rather than universally applicable advice [2]. A pharmacy knowledge graph must therefore preserve the epistemic status of each assertion and avoid converting observation into recommendation through graph structure

Address for correspondence: Mateo Alvarez, Department of Pharmacy Knowledge Graphs and Reasoning, Faculty of Pharmacy, National University of Colombia, Bogota, Colombia.
mateo.alvarez@unal.edu.co
Received: 21 January 2026; **Accepted:** 05 April 2026

How to cite this article: Alvarez M, Herrera S, Cruz D, Castro A. Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph. Arch Pharm Pract. 2026;17(2):13-20. https://doi.org/10.51847/v5Z0LHHaFi

alone. Alternative explanations—including documentation bias, incomplete capture, or site-specific workflow—must remain visible.

Medication knowledge is also more granular than a single drug node. Curated resources distinguish ingredients, products, strengths, formulations, targets, indications, interactions, and multiple source records [3]. Advice may differ across those levels. "Avoid," "adjust," "monitor," or "counsel" may apply to an ingredient class, one formulation, a dose range, a route, or a particular combination. Collapsing these distinctions can create apparently coherent but clinically misleading edges.

Clinical decision-support literature further shows that usefulness and safety depend on workflow integration, usability, maintenance, and the consequences of acting on an output [4]. The problem is therefore not merely how to connect medication data. It is how to represent advice so that users can inspect why it exists, when it applies, where it came from, how strongly it is supported, whether it conflicts with other advice, and whether it is eligible for a particular use. This article develops a proposed, non-empirical architecture for that purpose. It does not claim to validate a graph, recommend medication actions, or transfer professional authority to an automated system.

Knowledge-Representation Requirements

A medication-advice graph requires more than nodes and edges. Knowledge-graph design encompasses schema, entity identity, relation semantics, contextual qualification, validation constraints, and query behavior [5]. For pharmacy use, these requirements imply that medication, patient, action, rationale, source, evidence, and temporal concepts must have explicit types and allowable relationships. The graph should support queries that retrieve not only an action but also its conditions, provenance, and limits.

Representation, acquisition, completion, reasoning, and temporal modeling are distinct technical functions [6]. A relation imported from a curated source should therefore remain distinguishable from one extracted from text or predicted by a model. Likewise, a historical assertion should not become current merely because it remains in the graph. The proposed architecture treats provenance and time as properties of an assertion rather than optional annotations added after graph construction.

Provenance should disclose what was produced, from which input, through which process, and by which agent [7]. In medication advice, this includes the source document and version, the extraction or curation method, the responsible organization or curator, and the graph build in which the assertion appeared. Provenance improves traceability, but it does not establish that the source is correct or applicable.

Computable biomedical knowledge also requires maintenance, distribution, versioning, and stewardship across its lifecycle [8]. Consequently, the architecture must support update review, supersession, withdrawal, rollback, and retirement. These controls are proposed requirements for representational adequacy; they do not demonstrate medication-safety benefit or clinical readiness.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph.

Table 1. Definitions, components, relationships, and intended uses for the article Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph.

Component or Scientific Dimension	Problem Addressed	Inputs or Determinants	Proposed Mechanism or Relationship	Expected Contribution	Evidence Required	Failure or Uncertainty Risk	Boundary Statement
Graph Semantic Core	Ambiguous entities and predicates	Typed entities, relation definitions, constraints	Evidence-derived schema principles [5] linked to proposed pharmacy-specific constraints	Queryable meaning and conformance	Ontology review and competency testing	Semantically valid but clinically wrong assertions	Schema conformance does not establish clinical validity
Medication Identity	Conflation of ingredient, product, strength, route, and formulation	Curated medication identifiers and mappings [3]	Proposed identity-resolution layer	Advice attached to the correct medication level	Terminology mapping and pharmacist review	Duplicate, merged, or mis-scoped medication nodes	Identity resolution does not prove advice correctness
Advice Assertion	Context-free medication–action edges	Medication, action, rationale, goal, recipient	**Proposed synthesis:** reified Medication Advice Assertion	Inspectable advice object	Structural and domain-expert validation	Missing rationale or undefined action scope	The assertion is a representation, not a recommendation
Applicability Context	Advice generalized beyond its conditions	Population, patient factors, setting, trigger, exception	**Proposed synthesis:** applicability gate attached to each assertion	Conditional retrieval	Scenario-based validation	Unsafe overgeneralization or false exclusion	Formal conditions cannot replace clinical assessment
Temporal Scope	Obsolete advice treated as current	Valid time, event anchor, transaction time	Temporal qualification informed by knowledge-graph temporal requirements [6]	Time-aware retrieval and supersession	Temporal consistency testing	Incorrect or absent validity intervals	A current timestamp does not certify current evidence

Provenance	Source and transformation history lost	Source version, curator, extraction method, graph build	Assertion-level provenance based on established provenance functions [7]	Traceability and auditability	Source-to-assertion reconstruction	Traceable misinformation or incomplete lineage	Provenance indicates origin, not truth
Lifecycle Stewardship	Silent change and unmanaged obsolescence	Review status, update trigger, version, owner	Proposed review, rollback, withdrawal, and retirement states informed by computable-knowledge lifecycle needs [8]	Controlled maintenance	Governance-process and change-log audit	Unreviewed updates or orphaned assertions	Change control does not establish safe deployment
Intended-Use Boundary	Technical outputs used beyond evidence	Use category, professional role, review status	**Proposed synthesis:** use-eligibility relation	Separation of informational, research, and decision-support uses	Human-factors and organizational evaluation	Automation bias or authority displacement	Professional accountability remains external to the graph

Proposed Pharmacy Knowledge-Graph Schema

Heterogeneous biomedical graphs demonstrate that multiple entity and relation types can be integrated without reducing all knowledge to one generic association [9]. The proposed schema applies this principle to medication advice but changes the organizing unit. Instead of a direct edge from a medication to an action, it introduces a reified Medication Advice Assertion that can carry rationale, applicability, time, provenance, and governance attributes.

Source-specific evidence and curation information can coexist within an integrated knowledge platform [10]. Accordingly, each assertion is linked to a source record and source version rather than merely naming a journal, database, or guideline in free text. The graph also records whether the assertion was curated, extracted, transformed, or inferred. These distinctions permit source-aware retrieval and later reconstruction, but they do not rank sources automatically or establish that one source should prevail.

Medication relations must also remain semantically specific. Polypharmacy modeling shows why a generic "interacts with" edge is insufficient when different combinations produce different adverse-effect relations [11]. The proposed schema therefore separates action type, expected consequence, interaction mechanism, and medication-related problem. "Avoid," "adjust," "monitor," "counsel," and "refer" are represented as different action classes whose applicability may depend on dose, route, co-medication, laboratory state, or care setting.

Graph-predicted medication relations should remain distinguishable from externally checked or curated assertions. Previous work has compared graph-predicted adverse drug reactions with electronic health record evidence, illustrating both the value and the limits of external checking [12]. In the proposed architecture, inference status is never converted automatically into advice status. A predicted relation may generate a review candidate, but decision-use eligibility requires separate content, workflow, safety, and governance evaluation.

The schema has four connected layers. The medication and clinical-context layer represents ingredients, products, doses, routes, formulations, indications, patient factors, co-medications, laboratory states, and settings. The advice-assertion layer represents action, rationale, intended goal, recipient, eligibility, exceptions, and temporal scope. The evidence and provenance layer represents source versions, evidence items, extraction or curation processes, agents, and evidence-strength assessments. The governance layer represents contradiction sets, review states, supersession, withdrawal, retirement, and intended use. These layers are an original proposed synthesis and require empirical validation.

Figure 1 presents the original conceptual synthesis for pharmacy knowledge-graph schema for medication advice, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

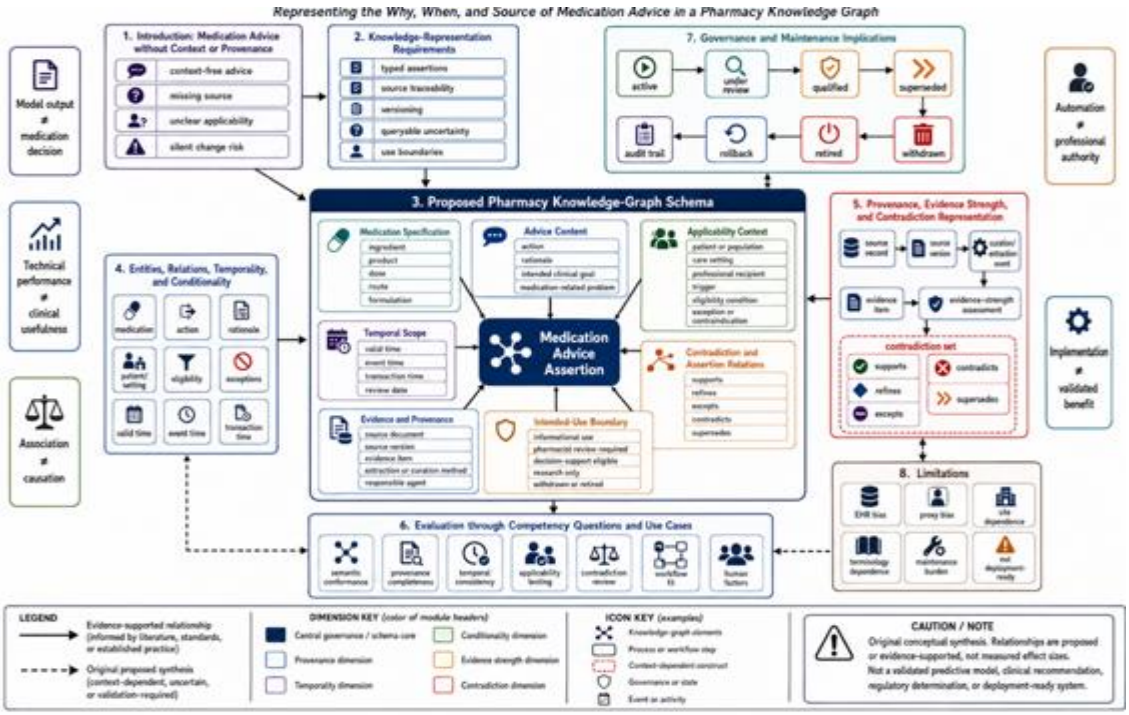

Figure 1. Pharmacy Knowledge-Graph Schema for Medication Advice

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Representing the Why, When, and Source of Medication Advice in a Pharmacy Knowledge Graph.

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 Ingestion	Preserve source identity and version	Source record → normalized evidence item	Feeds provenance and assertion creation	Reconstructable import based on provenance principles [7]	Lost version or undocumented transformation	Data steward and curator	Import success does not establish source quality
Entity Resolution	Map medication and clinical concepts	Raw identifiers → typed graph entities	Supports all assertion links	Correct ingredient, product, dose-form, and route resolution [3]	Entity collision or incorrect equivalence	Terminology specialist and pharmacist	Mapping accuracy does not validate advice
Heterogeneous Integration	Retain relation and entity differences	Multiple sources → typed integrated graph	Supplies context for assertions	Semantic preservation across source types [9]	Semantic collapse or duplicated meaning	Knowledge engineer and domain reviewer	Integration does not imply causal coherence
Assertion Construction	Bind action, rationale, context, time, and source	Evidence item → proposed assertion	Connects semantic, temporal, and governance layers	Required-field completeness under the proposed schema	Context-free or rationale-free assertion	Pharmacist curator and system owner	Complete fields do not establish correctness
Relation Typing	Distinguish advice and medication effects	Medication pair or state → typed relation	Constrains actions and consequences	Effect-specific relation accuracy informed by polypharmacy modeling [11]	Generic "interaction" relation hides consequence	Clinical-pharmacology reviewer	Relation specificity is not clinical benefit
Source And Evidence Qualification	Keep source attributes visible	Source metadata → provenance and evidence objects	Supports review and contradiction analysis	Source-specific metadata retained after integration [10]	Prestige substituted for evidence assessment	Evidence reviewer and governance body	Source authority is not certainty
Inference Review	Separate prediction from accepted advice	Model output → review candidate	May generate, but cannot authorize, assertions	External checking and error analysis informed by adverse-	Predicted edge displayed as recommendation	Model developer, pharmacist, and safety lead	Technical validation is

				reaction graph validation [12]			insufficient for decision use
Lifecycle Control	Govern change, supersession, and retirement	Updated evidence → reviewed assertion version	Updates provenance, contradiction, and use status	Version traceability and auditable state transition [8]	Silent change or irreversible overwrite	Named content owner and organization	Controlled maintenance does not equal deployment readiness

Entities, Relations, Temporality, and Conditionality

Ontology-backed biomedical data science distinguishes explicit domain knowledge, formal semantics, observed data, and inductive model outputs [13]. The proposed schema preserves that distinction. Core entities include medication ingredient, clinical product, dose, route, formulation, indication, medication-related problem, patient or population, care setting, professional recipient, source document, evidence item, and responsible agent. The Medication Advice Assertion is a proposed intermediary entity that binds these objects without treating them as equivalent kinds of evidence.

Relations are typed according to their function. An assertion may concern a medication, recommend an action, address a medication-related problem, pursue a clinical goal, apply to a population, require a condition, exclude an exception, derive from a source, and be reviewed by an agent. Relations between assertions are separately represented as supports, refines, excepts, contradicts, or supersedes. These categories are proposed; their reliability and inter-rater interpretability must be evaluated.

Clinical temporal information is heterogeneous and difficult to extract reliably from free text [14]. The schema therefore separates three temporal dimensions. Valid time states when the advice is intended to apply. Event time anchors it to a clinical occurrence, such as initiation, dose change, laboratory result, or transition of care. Transaction time records when the assertion entered or changed within the graph. This separation supports historical reconstruction and prevents a newly imported source from appearing to have always been valid. It does not guarantee that extracted dates or event relations are correct.

Computable guideline recommendations can be decomposed into recommendation, action, population, applicability, and evidence-related artifacts [15]. The proposed graph extends that logic to pharmacy advice by treating applicability as a structured gate. Conditions may involve indication, age, organ function, dose, route, duration, co-medication, laboratory state, pregnancy, previous response, care setting, or professional role. Exceptions and contraindications are represented explicitly rather than buried in prose.

Medication-interaction knowledge may depend on mechanism, interacting substances, exposure conditions, and evidence paths that do not fully agree [16]. Consequently, conditionality is not a binary property attached to the medication alone. It is evaluated at the assertion level and may differ among formulations, doses, routes, populations, and time periods. A query may therefore return several qualified assertions rather than one decontextualized answer. Formalizing these distinctions can make uncertainty inspectable, but it cannot establish that source information is complete, that a condition has been measured accurately, or that an assertion is clinically appropriate for an individual patient.

Provenance, Evidence Strength, and Contradiction Representation

Reproducible knowledge-graph construction requires explicit representation choices, versioned artifacts, and reconstructable transformation pipelines [17]. Each proposed Medication Advice Assertion should therefore retain its source record, source version, extraction or curation event, responsible agent, and graph release. A later source update should create a new assertion version rather than silently replacing the earlier state.

Standards, repositories, and policy metadata can be connected through governed, machine-readable descriptions [18]. The proposed architecture uses those links to identify applicable terminologies, source repositories, access conditions, and stewardship responsibilities. Such metadata can support reuse and auditability, but neither standardization nor FAIRness establishes clinical quality.

Evidence strength must remain separate from effect size, source prestige, citation frequency, or model confidence. Evidence-certainty methods demonstrate the need to assess confidence in evidence independently from quantitative estimates [19]. The graph would therefore store an evidence-strength assessment as a provenance-linked object, with its method, date, assessor, and scope visible.

Contradictory drug evidence should not be erased through source integration because disagreement may reflect population, intervention, outcome, or contextual differences [20]. The proposed contradiction set first tests whether assertions overlap in medication identity, action, population, condition, outcome, and valid time. Non-overlapping assertions may be refinements, exceptions, or supersessions rather than direct contradictions. Temporal ambiguity, evidence certainty, and contradiction are distinct but interacting uncertainty sources and should remain separately queryable [14, 19, 20]. This representation exposes uncertainty; it does not resolve disagreement or determine the correct medication action.

Evaluation through Competency Questions and Use Cases

Evaluation should begin with competency questions: Can the graph identify the source and version of an assertion? Can it retrieve why advice was issued, when it applies, which exceptions qualify it, and which assertions conflict? Technical testing should then examine identity resolution, relation constraints, temporal consistency, provenance completeness, and contradiction classification.

Aggregate link-prediction scores can obscure effects of relation frequency, graph structure, and evaluation design [21]. Model comparison also requires standardized implementations, datasets, hyperparameter control, and transparent reporting [22]. Accordingly, technical performance should be reported by relation type and error class rather than treated as a global measure of clinical quality.

Early clinical evaluation must additionally examine workflow, human factors, user behavior, safety, and implementation context [23]. Proposed use cases should span heterogeneous biomedical integration, adverse-reaction checking, and context-sensitive interaction reasoning [9, 12, 16]. Passing these tests would demonstrate only specified technical or workflow properties. Decision-use eligibility would still require pharmacist review, external validation, safety analysis, subgroup assessment, monitoring, and evidence that the graph supports rather than displaces professional judgment.

Governance and Maintenance Implications

Responsible healthcare machine learning requires governance across development, deployment, monitoring, and stakeholder responsibility [24]. For the proposed graph, named owners are needed for source ingestion, terminology mapping, assertion curation, evidence assessment, contradiction review, software maintenance, clinical oversight, and retirement decisions.

Transparency, replicability, ethics, effectiveness, and patient benefit should be evaluated as separate questions [25]. An inspectable source trail may improve accountability without proving that advice is accurate, useful, equitable, or safe. Governance should therefore distinguish content review from model review and organizational authorization from technical availability.

Trustworthy healthcare artificial intelligence further requires attention to fairness, universality, traceability, usability, robustness, and explainability [26]. **Figure 2** presents the original conceptual synthesis for provenance, temporality, conditionality, evidence strength, and contradiction, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

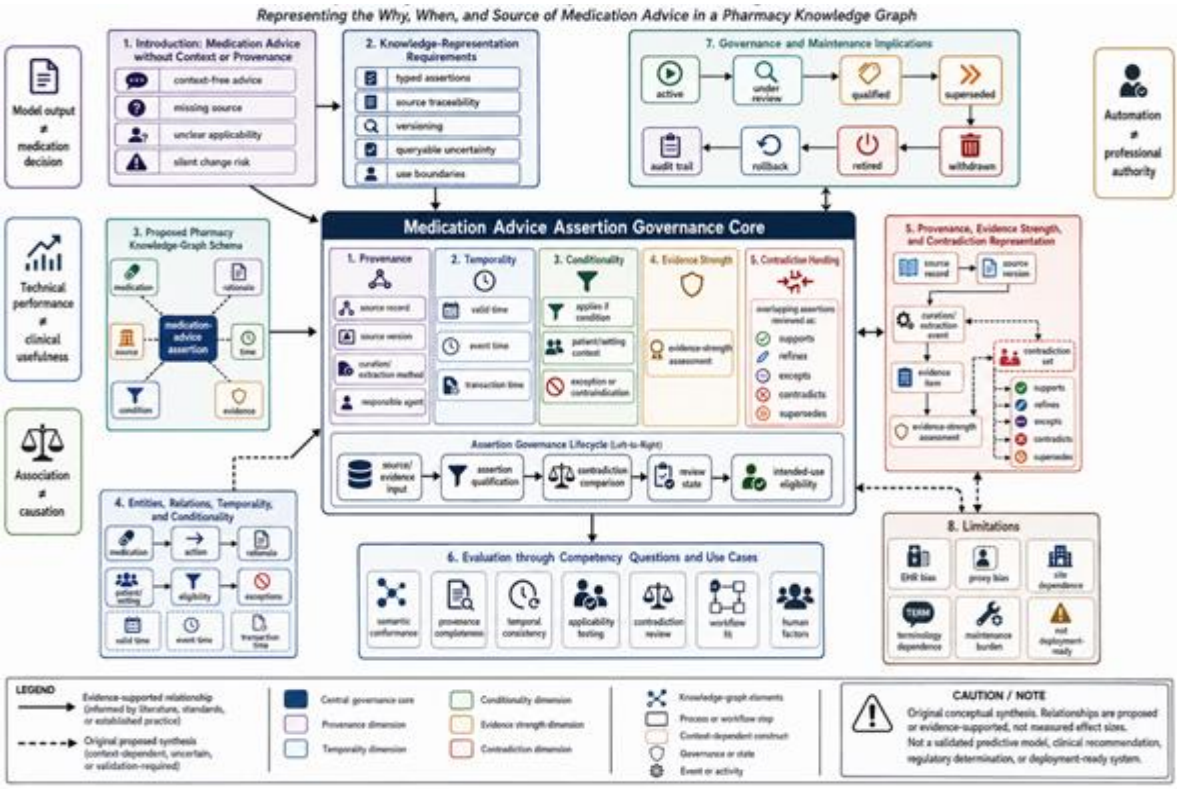

Figure 2. Provenance, Temporality, Conditionality, Evidence Strength, and Contradiction

Maintenance should connect assertion-level provenance, computable-knowledge stewardship, and versioned graph reconstruction [7, 8, 17]. Proposed state transitions include active, under review, qualified, superseded, withdrawn, and retired. Every transition should preserve the prior assertion, reason for change, responsible reviewer, and effective date. These controls require organizational and empirical validation and do not confer regulatory acceptance or safe deployment.

Limitations

Electronic-health-record-derived evidence may contain missingness, misclassification, documentation bias, and selection effects [27]. The graph cannot reconstruct knowledge that was never recorded or correct systematic distortion through representation alone. Apparently neutral proxies may also generate inequitable decisions or coverage even when protected characteristics are absent [28]. Equity evaluation must therefore examine source coverage, subgroup representation, and the consequences of missing or conflicting advice.

Technical quality and conceptual plausibility do not establish clinical impact, transferability, workflow benefit, or implementation success [29]. Additional limitations include dependence on terminology mappings, source licensing, curator agreement, contradiction-classification reliability, local professional roles, and continuing maintenance capacity.

Conclusion

This article proposes an assertion-centred pharmacy knowledge-graph architecture for representing medication advice together with its rationale, applicability, temporal scope, source, evidence strength, contradiction state, and intended-use boundary. Its central contribution is to replace context-free medication–action edges with inspectable Medication Advice Assertions linked to versioned evidence and governance processes.

The architecture may support clearer retrieval, review, updating, and comparison of medication advice, but representation does not establish truth, clinical appropriateness, or patient benefit. Evidence strength is not model confidence; provenance is not validation; technical performance is not clinical usefulness; and automation does not acquire professional authority. Future work must evaluate semantic conformance, pharmacist interpretation, temporal and contradiction accuracy, workflow effects, safety, equity, transportability, governance, and maintenance in defined settings. Until such evidence exists, the framework should be treated as an original conceptual proposal rather than a clinical recommendation, regulatory standard, or deployment-ready system.

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

References

1. Nicholson DN, Greene CS. Constructing knowledge graphs and their biomedical applications. Comput Struct Biotechnol J. 2020;18:1414-28. doi:10.1016/j.csbj.2020.05.017
2. Rotmensch M, Halpern Y, Tlimat A, Horng S, Sontag D. Learning a health knowledge graph from electronic medical records. Sci Rep. 2017;7(1):5994. doi:10.1038/s41598-017-05778-z
3. Wishart DS, Feunang YD, Guo AC, Lo EJ, Marcu A, Grant JR, et al. DrugBank 5.0: A major update to the DrugBank database for 2018. Nucleic Acids Res. 2018;46(D1). doi:10.1093/nar/gkx1037
4. 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
5. Hogan A, Blomqvist E, Cochez M, d'Amato C, de Melo G, Gutierrez C, et al. Knowledge graphs. ACM Comput Surv. 2021;54(4):71. doi:10.1145/3447772
6. Ji S, Pan S, Cambria E, Marttinen P, Yu PS. A survey on knowledge graphs: Representation, acquisition, and applications. IEEE Trans Neural Netw Learn Syst. 2022;33(2):494-514. doi:10.1109/TNNLS.2021.3070843
7. Herschel M, Diestelkämper R, Ben Lahmar H. A survey on provenance: What for? What form? What from? VLDB J. 2017;26(6):881-906. doi:10.1007/s00778-017-0486-1
8. Wong D, Peek N. Does not compute: Challenges and solutions in managing computable biomedical knowledge. BMJ Health Care Inform. 2020;27(2). doi:10.1136/bmjhci-2019-100123
9. Himmelstein DS, Lizee A, Hessler C, Brueggeman L, Chen SL, Hadley D, et al. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. Elife. 2017;6. doi:10.7554/eLife.26726
10. Piñero J, Ramírez-Anguita JM, Saüch-Pitarch J, Ronzano F, Centeno E, Sanz F, et al. The DisGeNET knowledge platform for disease genomics: 2019 update. Nucleic Acids Res. 2020;48(D1). doi:10.1093/nar/gkz1021
11. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics. 2018;34(13). doi:10.1093/bioinformatics/bty294
12. Bean DM, Wu H, Iqbal E, Dzahini O, Ibrahim ZM, Broadbent M, et al. Knowledge graph prediction of unknown adverse drug reactions and validation in electronic health records. Sci Rep. 2017;7(1):16416. doi:10.1038/s41598-017-16674-x
13. Callahan TJ, Tripodi IJ, Pielke-Lombardo H, Hunter LE. Knowledge-based biomedical data science. Annu Rev Biomed Data Sci. 2020;3:23-41. doi:10.1146/annurev-biodatasci-010820-091627
14. 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
15. Lichtner G, Alper BS, Jurth C, Spies C, Boeker M, Meerpohl JJ, et al. Representation of evidence-based clinical practice guideline recommendations on FHIR. J Biomed Inform. 2023;139:104305. doi:10.1016/j.jbi.2023.104305
16. Taneja SB, Callahan TJ, Paine MF, Kane-Gill SL, Kilicoglu H, Joachimiak MP, et al. Developing a knowledge graph for pharmacokinetic natural product-drug interactions. J Biomed Inform. 2023;140:104341. doi:10.1016/j.jbi.2023.104341
17. Callahan TJ, Tripodi IJ, Stefanski AL, Cappelletti L, Taneja SB, Wyrwa JM, et al. An open source knowledge graph ecosystem for the life sciences. Sci Data. 2024;11(1):363. doi:10.1038/s41597-024-03171-w
18. Sansone SA, McQuilton P, Rocca-Serra P, Gonzalez-Beltran A, Izzo M, Lister AL, et al. FAIRsharing as a community approach to standards, repositories and policies. Nat Biotechnol. 2019;37(4):358-67. doi:10.1038/s41587-019-0080-8
19. Brignardello-Petersen R, Bonner A, Alexander PE, Siemieniuk RA, Furukawa TA, Rochwerg B, et al. Advances in the GRADE approach to rate the certainty in estimates from a network meta-analysis. J Clin Epidemiol. 2018;93:36-44. doi:10.1016/j.jclinepi.2017.10.005

20. Sosa DN, Altman RB. Contexts and contradictions: A roadmap for computational drug repurposing with knowledge inference. Brief Bioinform. 2022;23(4). doi:10.1093/bib/bbac268
21. Rossi A, Firmani D, Matinata A, Merialdo P, Barbosa D. Knowledge graph embedding for link prediction: A comparative analysis. ACM Trans Knowl Discov Data. 2021;15(2):14. doi:10.1145/3424672
22. Ali M, Berrendorf M, Hoyt CT, Vermue L, Galkin M, Sharifzadeh S, et al. Bringing light into the dark: A large-scale evaluation of knowledge graph embedding models under a unified framework. IEEE Trans Pattern Anal Mach Intell. 2022;44(12):8825-45. doi:10.1109/TPAMI.2021.3124805
23. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
24. 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. Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368. doi:10.1136/bmj.l6927
26. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
27. 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
28. 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
29. 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