Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

THE THERAPEUTIC EVIDENCE GRAPH UNIFIES MOLECULAR DATA, BIOLOGICAL MECHANISMS, DISEASE KNOWLEDGE, AND EXPERIMENTAL CLAIMS WITHOUT ERASING THEIR PROVENANCE

Ravi Patel[1]*, Maria Santos[2], Ananya Sharma[3]

1. ***Department of Knowledge Graphs and Evidence Integration, Faculty of Pharmacy, University of Mysore, Mysore, India.***
2. ***Department of Molecular Data Provenance and Biological Mechanisms, Faculty of Pharmacy, University of São Paulo, São Paulo, Brazil.***
3. ***Department of Therapeutic Evidence and Disease Knowledge Representation, Faculty of Pharmaceutical Sciences, University of Agricultural Sciences Bangalore, Bangalore, India.***

ARTICLE INFO

Received:
25 May 2024
Received in revised form:
28 September 2024
Accepted:
01 October 2024
Available online:
28 October 2024

Keywords: Therapeutic evidence graph, Pharmaceutical knowledge architecture, Biomedical knowledge graph, Evidence provenance, Molecular representation

ABSTRACT

Pharmaceutical knowledge is distributed across chemical databases, bioassay repositories, target resources, pathway models, disease ontologies, experimental studies, computational predictions, and curated therapeutic claims. Although knowledge graphs can connect these resources, conventional graph integration may obscure how an assertion was generated, which source supplied it, whether apparently supporting records are independent, and under which molecular, biological, experimental, or temporal conditions it applies. This article develops the Therapeutic Evidence Graph as an original pharmaceutical knowledge architecture for connecting molecular data, biological mechanisms, disease knowledge, studies, and experimental claims while retaining their evidential identities. The proposed architecture distinguishes therapeutic entities from computational representations, scientific assertions from their supporting evidence, and source agreement from evidential independence. Its central unit is a qualified therapeutic assertion accompanied by provenance, context, polarity, temporality, uncertainty, and applicability information. Contradictions remain represented as inspectable relationships rather than being removed through premature harmonization, while source lineage and version history permit conclusions to be examined under alternative evidence configurations. The architecture may support traceable reasoning for target discovery, drug repurposing, mechanism assessment, and evidence governance, but graph connectivity or predictive ranking cannot independently establish causality, pharmacological efficacy, developability, safety, clinical utility, or regulatory acceptability. The contribution is conceptual and requires implementation-specific evaluation of semantic fidelity, provenance completeness, evidence independence, temporal reconstruction, uncertainty calibration, interoperability, and decision usefulness. By treating provenance as part of therapeutic meaning rather than peripheral metadata, the Therapeutic Evidence Graph offers a bounded foundation for pharmaceutical evidence integration without converting heterogeneous observations and claims into an artificial consensus.

To Cite This Article: Patel R, Santos M, Sharma A. The Therapeutic Evidence Graph Unifies Molecular Data, Biological Mechanisms, Disease Knowledge, and Experimental Claims without Erasing Their Provenance. Pharmacophore. 2024;15(5):48-58. https://doi.org/10.51847/6V7MXb1SAc

Introduction

Pharmaceutical discovery depends on relationships that cross chemical, biological, experimental, and disease scales. A candidate molecule may be associated with a molecular representation, a measured bioactivity, a target, a pathway perturbation, a disease phenotype, an experimental study, and a proposed therapeutic use. Biomedical knowledge graphs provide a computational structure for organizing such heterogeneous entities and their typed relationships, supporting navigation, representation learning, link prediction, and other forms of graph-based analysis [1]. Their relevance to pharmaceutical science arises not simply from graph topology, but from the possibility of connecting evidence that would

Corresponding Author: Ravi Patel; Department of Knowledge Graphs and Evidence Integration, Faculty of Pharmacy, University of Mysore, Mysore, India. E-mail: ravi.patel@uom.ac.in.

otherwise remain separated by database structure, disciplinary vocabulary, identifier systems, and analytical scale. Nevertheless, connection alone does not determine whether two records are comparable, whether an edge represents an observation or prediction, or whether a proposed route through the graph has pharmacological meaning.

Knowledge graphs have accordingly been investigated across target identification, drug–target interaction analysis, drug repurposing, adverse-effect analysis, and molecular discovery. These applications demonstrate that graph structures may expose relations that are difficult to inspect when evidence remains divided among independent resources [2]. They do not, however, justify treating every connected path as a mechanistic explanation or every predicted edge as a therapeutically actionable relationship. A graph can organize an association between a molecule and disease while remaining unable to establish whether the association reflects direct target modulation, downstream pathway effects, shared literature, dataset leakage, or a computational inference that has not been experimentally examined. Pharmaceutical interpretation therefore requires an architecture that preserves the evidential status of each relation throughout integration and reasoning.

The practical value of a pharmaceutical knowledge graph is conditional on its construction choices, evidence coverage, semantic granularity, and intended application. Reviews of drug-discovery knowledge graphs distinguish technically interesting graph operations from scientifically useful tools and emphasize that target discovery and repurposing outputs remain dependent on the underlying information and the purpose for which the graph is queried [3]. This dependence becomes particularly important when the same graph contains manually curated mechanisms, high-throughput assay records, disease associations, text-mined statements, model predictions, and clinical observations. Without visible distinctions among those evidence types, graph density may be mistaken for scientific convergence, and computational proximity may be interpreted more strongly than the underlying sources permit.

Available drug-discovery datasets differ substantially in their primary entities, relationship types, scale, licensing arrangements, update cycles, identifiers, and suitability for knowledge-graph construction [4]. The unresolved problem is therefore not merely how to aggregate more pharmaceutical data, but how to unify relevant evidence without erasing the conditions under which it was produced. This article proposes the Therapeutic Evidence Graph as a provenance-preserving, temporally versioned, and context-qualified knowledge architecture. Its original contribution is the use of the qualified therapeutic assertion as the central graph unit: a typed subject–relationship–object claim connected to distinguishable evidence objects and accompanied by source lineage, methodological context, polarity, uncertainty, temporal status, and applicability boundaries. The architecture is intended to organize evidence and support inspectable reasoning; it is not presented as an empirically validated system, autonomous discovery platform, clinical decision instrument, or substitute for experimental and pharmacological evaluation.

Fragmentation of Pharmaceutical Evidence Across Data Sources and Scientific Scales

Fragmentation begins at the level of experimental measurement. Bioactivity resources contain molecular structures, targets, assays, activity values, units, endpoint descriptions, and study-derived annotations, but these elements are produced under assay-specific conditions and curation processes. ChEMBL illustrates how medicinal-chemistry information is extracted, standardized, mapped to controlled vocabularies, and supplemented by direct data deposition while retaining details about compounds, assays, targets, and activity measurements [5]. Such standardization makes records more searchable and reusable, yet it does not make all activity values experimentally equivalent. Measurements may differ by assay format, biological system, target construct, endpoint, concentration convention, unit conversion, or curation history. A therapeutic evidence architecture must consequently represent the assay and evidence context of an activity relation rather than retaining only a simplified molecule–target edge.

Fragmentation also arises because drug-centered resources combine scientific objects with different epistemic and practical meanings. DrugBank connects chemical structures and drug identities with targets, enzymes, carriers, transporters, pathways, pharmacological descriptions, interactions, classifications, and other drug-related information [6]. These relationships are valuable precisely because they are heterogeneous, but their integration into one resource does not mean that a curated target association, a metabolic relationship, an interaction warning, and a pathway annotation possess the same evidential basis. Moreover, a named drug may correspond to an active moiety, salt, stereochemical form, mixture, formulation, or regulated product. Collapsing these levels can produce false continuity between molecular prediction, pharmacological action, formulation-specific exposure, and therapeutic use. The proposed Therapeutic Evidence Graph therefore separates chemical identity, molecular form, drug concept, and product context whenever the supporting evidence requires those distinctions.

Disease evidence introduces a further layer of fragmentation. DisGeNET integrates and standardizes disease-associated genes and variants from curated repositories, literature-derived sources, animal-model information, and other evidence channels [7]. A disease–gene relation may therefore represent genetic association, variant evidence, curated knowledge, text-mined reporting, or a composite score built from multiple sources. These forms of support cannot be assumed to establish causality, target tractability, therapeutic direction, or relevance to a particular disease subtype. Disease labels themselves may encompass symptoms, traits, abnormal phenotypes, broad diagnostic categories, or molecularly defined conditions. A provenance-preserving graph must retain the originating evidence type, disease mapping, source version, and applicability context so that a repeated association across derivative resources is not misinterpreted as independent replication.

DrugCentral similarly demonstrates the value and complexity of connecting active pharmaceutical ingredients with indications, mechanisms of action, targets, pharmaceutical information, and regulatory or literature-derived records that may assist discovery and repositioning [8]. For the proposed architecture, fragmentation is not treated solely as a technical failure

to merge databases. It is a scientific condition produced by differences in what sources observe, curate, infer, normalize, and update. Integration must therefore preserve source-specific meaning while enabling cross-source traversal. The architecture should permit investigators to ask not only whether an association exists, but which evidence objects support it, whether they are independent, what contexts qualify them, whether contradictory records exist, and which graph state was current at the time of analysis. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop fragmentation of pharmaceutical evidence across data sources and scientific scales within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Fragmentation of pharmaceutical evidence across data sources and scientific scales

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Molecular identity	Are records referring to the same chemical entity, form, stereoisomer, salt, mixture, or product?	Chemical structures and drug records require standardization and explicit identity mappings [5].	Establishes the entity layer on which other evidence depends.	Different molecular forms may be merged under one drug name.	Shared names or structural similarity do not establish pharmaceutical interchangeability.
Bioassay evidence	Under which experimental conditions was a molecular activity observed?	Bioactivity records depend on assay design, target definition, endpoint, biological system, units, and curation [5].	Grounds molecule–target assertions in distinguishable experimental objects.	A normalized activity edge may conceal incomparable assays.	An observed assay effect does not alone establish in vivo action or therapeutic value.
Drug and pharmacological knowledge	Which type of relationship connects a drug to a target, enzyme, transporter, pathway, interaction, or indication?	Drug resources integrate multiple chemical, biological, and pharmacological evidence classes [6].	Prevents heterogeneous drug relations from being treated as semantically equivalent.	Database integration may be mistaken for uniform evidential strength.	A recorded drug relationship must be interpreted according to its relation type and source context.
Disease-genomics evidence	What kind of evidence supports a gene–disease or variant–disease association?	Disease knowledge platforms integrate curated, literature-derived, genetic, and other source classes [7].	Qualifies disease relationships by evidence type and disease concept.	Repeated derivative records may be counted as independent support.	Association does not independently establish causality, therapeutic direction, or target validity.
Indication and mechanism records	Is a relation an approved indication, curated mechanism, investigational use, literature claim, or repositioning hypothesis?	Drug-centered resources connect ingredients with indications, targets, mechanisms, and source-specific records [8].	Separates established uses from mechanistic or computational hypotheses.	A drug–disease link may be interpreted as evidence of efficacy.	Graph connectivity does not establish an appropriate indication or clinical recommendation.
Scientific scale	At what level was the evidence generated: molecular, cellular, tissue, organismal, disease, or therapeutic?	Pharmaceutical evidence crosses non-equivalent biological and experimental scales.	Provides applicability qualifiers for cross-scale reasoning.	Findings may be transferred between scales without justification.	Evidence at one scale supports another only when a validated translational relationship is available.
Source lineage	Do apparently distinct records originate from the same study, database, or derivation chain?	Integrated resources frequently reuse primary and secondary evidence.	Enables evidence-independence and source-dependence analysis.	Duplicate evidence can create artificial consensus.	Record multiplicity is not equivalent to independent replication.
Temporal state	Which source version, publication status, and graph release support the assertion?	Databases, ontologies, mechanisms, and classifications change over time.	Supports historical reconstruction and controlled updating.	New releases may overwrite prior interpretations or remove context.	Recency alone does not determine evidential quality.
Qualified therapeutic assertion	What exactly is claimed, under which conditions, and by what evidence?	The proposed architecture separates a typed assertion from the evidence objects supporting or challenging it.	Provides the central integrative unit of the Therapeutic Evidence Graph.	A bare edge may be treated as universal knowledge.	The qualified assertion is a proposed architectural construct requiring implementation and validation.

Representing Molecules, Targets, Pathways, Diseases, Studies, and Claims

A therapeutic graph must first distinguish a molecule from the representation through which a computational method processes it. Mol2vec, for example, learns vector representations of molecular substructures from their contextual occurrence, illustrating that a molecular embedding is a derived analytical object rather than the chemical entity itself [9]. Other molecular-learning approaches operate on calculated descriptors, fingerprints, strings, or molecular graphs, and comparative work shows that representation and model choice affect predictive behavior across property-prediction tasks [10]. The Therapeutic Evidence Graph should therefore maintain at least two linked objects: a versioned molecular-identity object and one or more representation objects describing the featurization method, parameters, software context, and model use. This distinction prevents a model-specific latent vector from being presented as intrinsic molecular knowledge and allows predictions based on incompatible representations to be examined separately.

Targets require similarly careful representation. A target node should specify whether it denotes a gene, protein, isoform, protein family, molecular complex, nucleic acid, or another biological entity, while target-state qualifiers may include organism, tissue, cell type, mutation, conformation, post-translational modification, or disease context. Molecule–target

assertions should then differentiate binding, inhibition, activation, degradation, stabilization, substrate relationships, expression effects, and predicted associations. Disease representation requires normalized identifiers, definitions, synonyms, and hierarchical relationships, as exemplified by the Human Disease Ontology's effort to provide consistent disease semantics across biomedical resources [11]. Even with ontology alignment, however, disease concepts may remain broader or narrower than the populations and biological states examined in an individual study. The graph must preserve both the normalized disease concept and the study-specific disease description rather than allowing normalization to erase clinically or biologically material differences.

Pathways and biological processes connect molecular interactions to higher-order mechanistic interpretations, but pathway resources vary in scope, granularity, boundaries, representation language, and curation policy. Pathway Commons integrates pathway and molecular-interaction information from multiple contributing resources and retains versioned releases, illustrating both the usefulness and the complexity of cross-resource pathway representation [12]. In the proposed architecture, a pathway should not be treated as a single universal container. It should remain linked to its source model, version, organism, cell context, constituent reactions, and evidence basis. A molecule's connection to a pathway may reflect direct target participation, downstream perturbation, enrichment analysis, curated membership, or computational inference. These relations should use different predicates and evidence qualifiers because pathway membership, pathway modulation, and therapeutic mechanism are not interchangeable claims.

Studies and claims form a final representational layer that is often flattened in biomedical graphs. The PubMed Knowledge Graph demonstrates that publications, authors, affiliations, funding, and biomedical entities can be represented as linked graph objects [13]. A pharmaceutical evidence architecture must go further by distinguishing the study from the assertions attributed to it and each assertion from the evidence reported within it. The proposed model therefore represents a study object, its experimental or analytical components, individual evidence objects, and qualified therapeutic assertions as connected but non-equivalent structures. A claim may support, contradict, refine, or remain incomparable with another claim depending on its molecular form, assay, target state, disease context, model system, direction of effect, and temporal status. This design permits molecules, targets, pathways, diseases, studies, and claims to be traversed within one architecture while ensuring that semantic normalization does not silently transform heterogeneous evidence into a universal or causally ordered account.

Provenance-Preserving Graph Architecture and Evidence-Layer Design

The Therapeutic Evidence Graph is proposed as a layered architecture rather than a flat network of entities and edges. Its first layer contains normalized therapeutic entities; its second contains qualified assertions; its third contains evidence objects; and its fourth contains provenance, uncertainty, temporal, and applicability qualifiers. This separation is necessary because link-prediction performance alone cannot demonstrate that graph semantics, source lineage, or evidential context have been preserved. Biomedical graph benchmarking can test whether models recover withheld relations under defined conditions, but benchmark success remains distinct from faithful pharmaceutical representation and therapeutic usefulness [14]. The architecture must therefore be evaluated both as a computational graph and as an evidence-preserving scientific record.

The central architectural unit is the qualified therapeutic assertion. Each assertion contains a subject entity, a typed and directional predicate, an object entity, and explicit qualifiers describing experimental context, polarity, temporal validity, and applicability. Pharmaceutical graphs such as PharmKG demonstrate the value of typed entities, relations, and attributes for biomedical data mining [15]. The proposed architecture extends this logic by preventing the assertion from absorbing its supporting evidence. A molecule–target assertion, for example, remains separate from the assay record, database statement, computational prediction, or publication claim that supports it. Multiple evidence objects may support, challenge, or qualify the same assertion while remaining individually inspectable.

Evidence access need not require complete physical centralization. BioThings Explorer shows that semantically annotated biomedical application programming interfaces can be combined into a federated query environment [16]. A Therapeutic Evidence Graph may therefore contain locally stored evidence, federated references, or hybrid representations, provided that every result retains a resolvable source identifier, source version, transformation history, access date, and extraction or curation route. Federation can reduce unnecessary duplication and preserve source-specific control, but it also introduces dependence on external availability, changing interfaces, and inconsistent annotation. The architecture should record whether an assertion was retrieved directly, transformed locally, inferred through a query chain, or copied from a derivative resource.

Multiscale integration is required because therapeutic reasoning crosses molecules, proteins, pathways, phenotypes, diseases, and interventions. PrimeKG illustrates how diverse biomedical entities can be connected within a precision-medicine graph while retaining typed relationships across biological scales [17]. PheKnowLator further demonstrates that ontology-grounded graph construction can be made configurable and reproducible across distinct representation choices [18]. The proposed contribution combines these principles with explicit evidence layering: the graph should reveal not only that two entities are connected, but which evidence class created the connection, which transformations were applied, and what scientific boundary constrains interpretation. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop provenance-preserving graph architecture and evidence-layer design within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Provenance-preserving graph architecture and evidence-layer design

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Therapeutic entity layer	Molecular structures, drug identifiers, genes, proteins, pathways, phenotypes, diseases, assays, studies	Maintains normalized but versioned scientific referents	Distinguishable entity objects with resolvable identifiers	Identity mappings may merge non-equivalent forms or split equivalent concepts	Expert-reviewed entity mapping, duplicate detection, and identity-resolution testing
Molecular representation layer	Fingerprints, descriptors, molecular graphs, embeddings, model-specific encodings	Records how a molecule was represented computationally	Representation objects linked to the underlying molecular identity	Representations may be task-specific, lossy, or model-dependent	Reproducibility of featurization and documentation of software, parameters, and version
Qualified assertion layer	Subject entity, typed predicate, object entity, direction, polarity, context	Expresses the exact scientific or therapeutic claim being made	Inspectable, context-qualified assertions	Predicates may be semantically ambiguous or broader than the source claim	Relation-level expert review, directionality checks, and context-completeness audit
Evidence-object layer	Assay results, curated statements, experimental observations, omics associations, computational predictions, literature claims	Preserves the record that supports, limits, or contradicts an assertion	Evidence objects linked without being merged into the assertion	Evidence may be incomplete, duplicated, dependent, or poorly reported	Source verification, evidence-type classification, and duplicate-lineage analysis
Provenance envelope	Source, version, record identifier, method, curator or algorithm, date, transformation history, license	Makes assertions and evidence traceable through their full lineage	Reconstructable origin and processing history	Provenance may be missing or partially recoverable	Independent lineage reconstruction and provenance-completeness audit
Qualification layer	Experimental system, tissue, organism, disease subtype, temporal interval, applicability boundary	Prevents context-specific evidence from appearing universal	Context-bounded assertions	Context fields may be absent or inconsistently encoded	Controlled-vocabulary mapping and expert assessment of applicability
Contradiction and uncertainty layer	Conflicting claims, measurement variability, model uncertainty, curation uncertainty	Preserves disagreement and differentiated uncertainty	Explicitly challenged or uncertain assertions	A single score may conceal incompatible uncertainty types	Calibration studies, contradiction benchmarks, and adjudicated test cases
Reasoning-trace layer	Graph queries, path searches, embeddings, ranking algorithms	Records the entities, assertions, evidence, and operations producing an output	Reproducible target, repurposing, or mechanism hypothesis trace	A faithful computational trace may still lack biological validity	Query replay, unsupported-step detection, and application-specific expert review
Governance and release layer	Ingestion policies, validation rules, versions, corrections, access controls	Maintains integrity across graph updates and releases	Versioned, auditable graph states	Governance decisions may be inconsistent or poorly documented	Release audit, correction tracking, deprecation testing, and role accountability

Figure 1 presents the provenance-preserving therapeutic evidence constellation, showing how the article's key components and boundaries are connected within provenance-preserving graph architecture and evidence-layer design.

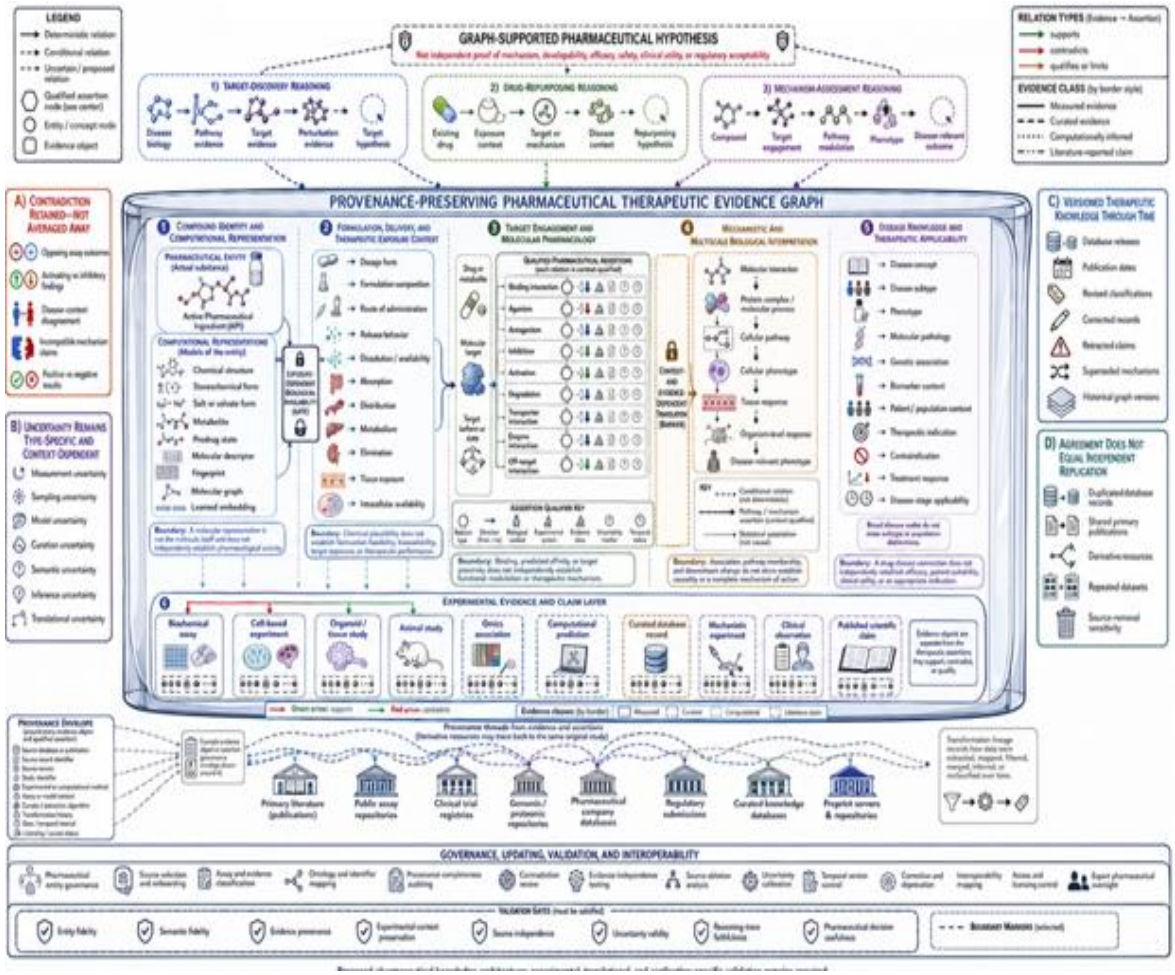

Figure 1. Provenance-Preserving Therapeutic Evidence Constellation.

The figure is an original conceptual synthesis that organizes the article's central contribution across fragmentation of pharmaceutical evidence across data sources and scientific scales, representing molecules, targets, pathways, diseases, studies, and claims, representing molecules, provenance-preserving graph architecture and evidence-layer design, encoding contradiction, uncertainty, temporality, and source dependence, and encoding contradiction. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Encoding Contradiction, Uncertainty, Temporality, and Source Dependence

Uncertainty must be represented as a family of distinguishable conditions rather than a single confidence score. Molecular property-prediction studies show that scalable uncertainty-estimation methods differ in calibration, behavior, and sensitivity to data conditions [19]. The proposed graph should therefore distinguish measurement uncertainty, sampling uncertainty, model uncertainty, curation uncertainty, semantic uncertainty, and inference uncertainty. An uncertainty annotation must identify what is uncertain, how the estimate was produced, and which task and dataset support its interpretation. A numerical confidence value without this context risks combining conceptually incompatible quantities and may create an appearance of precision that the evidence does not justify.

Evidential learning methods may estimate predictive evidence or uncertainty directly within molecular models [20]. Such outputs can be retained as evidence objects, but they should not become universal confidence attributes attached to molecular or therapeutic assertions. Model confidence is conditional on architecture, training data, representation, objective function, and evaluation setting. It does not quantify uncertainty in the underlying biology, assay quality, disease relevance, or translational applicability. The Therapeutic Evidence Graph should consequently store the predictive output, uncertainty method, training context, data version, and calibration evidence together, allowing later users to determine whether the estimate is suitable for a different query or application.

Source dependence must also remain visible because graph agreement may be produced by repeated use of related datasets. Ligand-based benchmarks can reward memorization when structurally similar compounds or closely related records are distributed across training and test sets [21]. MoleculeNet likewise demonstrates that molecular-learning conclusions depend on task definition, data split, representation, and evaluation metric [22]. A graph should therefore record dataset lineage,

overlap, derivation, and reuse. Source-ablation analyses can then test whether a target or repurposing ranking remains stable when a dominant database, assay family, publication group, or duplicated evidence chain is removed.

Contradiction and temporality should be encoded through explicit objects and validity intervals. Comparative drug-target prediction studies show that performance is conditioned by assay composition and evaluation design [23]. Apparent disagreement may therefore reflect different experimental systems, endpoint definitions, molecular forms, target constructs, or time-specific knowledge rather than a simple true–false conflict. The graph should retain incompatible findings, describe the dimensions on which they differ, and allow unresolved contradictions to remain unresolved. Corrections, retractions, database updates, revised classifications, and superseded mechanisms should create new graph states rather than silently overwriting historical evidence.

Reasoning Tasks for Target Discovery, Repurposing, and Mechanism Assessment

In target discovery, graph reasoning can organize disease associations, target biology, pathway context, molecular interactions, experimental perturbations, and existing pharmacological evidence. The objective is not to calculate one universal target-validity score but to expose the evidential route by which a candidate becomes prioritizable. Integrated biomedical networks have been used to rank drug-repurposing hypotheses through typed paths linking drugs, genes, diseases, and other biomedical entities [24]. Within the proposed architecture, an analogous path must remain connected to its supporting evidence objects, source lineage, and applicability qualifiers so that users can distinguish direct observations from inferred transitions.

Heterogeneous network integration may also support drug–target interaction prediction and computational repositioning by combining multiple similarity and association structures [25]. The Therapeutic Evidence Graph extends this approach by requiring that each input relation retain its evidential class and source dependence. A predicted drug–target edge may be scientifically useful as a hypothesis, but it should remain distinguishable from measured binding, functional modulation, and therapeutic mechanism. Repurposing routes should additionally preserve disease subtype, molecular form, exposure context, safety evidence, and temporal status because a plausible network connection does not establish that an existing drug can reach, modulate, or safely influence the relevant biological system.

Knowledge-graph embeddings can rank potential protein targets by learning relational patterns across a biomedical graph [26]. Such rankings may reveal candidates that are not obvious from isolated datasets, but their latent representations can obscure which sources or paths drove the result. The proposed architecture therefore requires a reasoning trace that records the graph version, model, training relations, excluded evidence, source concentrations, and evidence path supporting each candidate. Graph-based prioritization should be interpreted as hypothesis organization rather than mechanistic confirmation, and experimentally negative or contradictory evidence must remain visible even when it lowers model convenience or ranking coherence.

Mechanism assessment requires stricter evidence distinctions than target ranking. Graph-convolutional models have represented drug–drug–protein relations to predict polypharmacy side effects [27], demonstrating the value of typed relational structure for safety-oriented hypothesis generation. DrugMechDB further shows that mechanisms can be curated as multi-step paths connecting drugs to intermediate biological events and outcomes [28]. In the Therapeutic Evidence Graph, mechanism paths should distinguish direct biochemical evidence, cellular effects, pathway changes, phenotypes, and inferred transitions. A complete-looking path must not be labelled causal unless its intermediate relations and directionality are supported by appropriate evidence and remain valid in the relevant biological and therapeutic context.

Governance, Updating, Validation, and Interoperability Requirements

Governance must begin before data ingestion. Biopharmaceutical implementation of FAIR principles requires organizational stewardship, standards, incentives, infrastructure, and clearly assigned responsibilities [29]. For the Therapeutic Evidence Graph, governance should define source eligibility, ingestion frequency, identifier policies, evidence classifications, licensing constraints, conflict management, correction procedures, and access controls. Every transformation should be attributable to a curator, algorithm, or controlled workflow. Scientific responsibility cannot be delegated entirely to an automated extraction system because semantic correctness, evidence comparability, and applicability boundaries often require domain expertise.

Validation should be multidimensional rather than represented by one graph-quality score. FAIR design frameworks separate findability, accessibility, interoperability, and reusability into related but distinct properties that can be assessed through explicit metrics [30]. A therapeutic graph similarly requires separate tests for entity fidelity, semantic fidelity, provenance completeness, evidence independence, contradiction retention, temporal reconstruction, uncertainty representation, interoperability, and reasoning faithfulness. Passing one dimension cannot compensate automatically for failure in another. A graph may be technically accessible but scientifically misleading, semantically consistent but outdated, or computationally predictive while relying heavily on duplicated evidence.

Governance must also support recurring maturity assessment. Automated and community-governed FAIR evaluation can provide scalable, transparent tests for digital resources [31]. Comparable mechanisms may be adapted to graph releases, provided that automated checks are not treated as substitutes for scientific review. Each release should document source versions, additions, removals, corrected mappings, changed predicates, deprecated assertions, known limitations, and failed validation checks. Historical releases should remain reconstructable so that a therapeutic hypothesis can be interpreted against the evidence state that originally produced it.

Interoperability requires shared semantic structures and reproducible build processes. The Biolink Model provides common categories, predicates, and association patterns for biomedical and translational knowledge graphs [32]. KG-Hub demonstrates how versioned extraction, transformation, loading, reusable subgraphs, and standardized builds can support graph exchange and reproducibility [33]. The proposed architecture should map to such shared structures while retaining domain-specific qualifiers that cannot be represented safely through generic predicates. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop governance, updating, validation, and interoperability requirements within the article's central argument.

Table 3. Evaluation, implementation, and research priorities arising from The Therapeutic Evidence Graph Unifies Molecular Data, Biological Mechanisms, Disease Knowledge, and Experimental Claims without Erasing Their Provenance

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Entity fidelity	Whether molecules, targets, pathways, diseases, assays, and studies are mapped correctly	Expert-adjudicated mappings, duplicate testing, molecular-form checks	Non-equivalent entities are merged or equivalent concepts remain fragmented	Correct identity mapping does not validate associated claims	Determines whether graph queries refer to the intended scientific objects
Semantic fidelity	Whether predicates retain direction, biological meaning, and experimental context	Relation-level expert review and controlled-vocabulary testing	Binding, modulation, association, indication, and causality are conflated	Semantic consistency does not establish biological truth	Prevents inappropriate interpretation of graph paths
Provenance completeness	Whether every assertion and evidence object can be reconstructed from source to graph	Lineage audit, source replay, version reconstruction	Orphaned edges or undocumented transformations	Complete provenance does not guarantee source quality	Enables scientific audit and correction
Evidence independence	Whether apparently separate records derive from independent evidence	Citation lineage, database-derivation analysis, source overlap testing	One primary study is counted repeatedly through derivative resources	Independence may remain uncertain when lineage is missing	Prevents artificial inflation of evidential support
Contradiction retention	Whether incompatible findings remain represented and context-qualified	Adjudicated contradiction sets and conflict-recall testing	Disagreement is deleted, averaged, or labelled incorrectly	Preserving contradiction does not resolve scientific disagreement	Reveals where additional investigation is needed
Temporal reconstruction	Whether previous graph states and superseded assertions can be recovered	Release comparison, correction and deprecation testing	New data overwrite prior evidence without history	Newer evidence is not necessarily stronger evidence	Supports reproducibility of historical analyses
Uncertainty validity	Whether uncertainty annotations are calibrated and correctly typed	Calibration studies, coverage assessment, uncertainty-field audit	One confidence value combines incompatible uncertainty sources	Calibration may not transfer across tasks or datasets	Supports bounded rather than absolute interpretation
Source-dependence robustness	Whether outputs remain stable under source removal and deduplication	Source-ablation analysis and influence testing	One source or derivative evidence family determines the output	Stability does not prove correctness	Identifies fragile rankings and hidden evidence concentration
Reasoning faithfulness	Whether reported traces correspond to the data and computations producing outputs	Query replay, model logging, unsupported-step analysis	Post hoc paths are presented as the actual basis of prediction	Computational faithfulness is not equivalent to biological causality	Enables expert inspection of target and mechanism hypotheses
Interoperability	Whether exchange retains entity meaning, evidence context, and provenance	Cross-schema mapping and round-trip conversion	Identifiers transfer while qualifiers or lineage are lost	Technical compatibility does not guarantee semantic equivalence	Supports responsible reuse across platforms
Governance maturity	Whether updates, corrections, access, ownership, and deprecation are accountable	Release audit, policy compliance review, correction tracking	Errors persist without responsible ownership or visible correction	Process maturity cannot establish therapeutic validity	Determines whether the graph can be maintained as a trustworthy research resource
Application usefulness	Whether graph use improves a defined pharmaceutical workflow	Prospective expert studies and application-specific experimental follow-up	High graph performance does not improve decisions or evidence quality	Utility in one workflow does not establish universal applicability	Determines whether implementation is justified for target, repurposing, or mechanism work

Limitations and Boundary Conditions

The proposed Therapeutic Evidence Graph is a conceptual architecture and has not been empirically validated as an integrated system. Its effectiveness will depend on source quality, coverage, identifier resolution, ontology alignment, evidence extraction, update frequency, governance capacity, and the scientific competence of its users. Complete provenance may be impossible when databases do not expose derivation histories or when literature claims depend on inaccessible supplementary methods. The architecture can make missing provenance visible, but it cannot reconstruct information that was never reported or retained.

Graph representation also cannot remove the epistemic limits of the underlying evidence. Associations may remain non-causal, computational predictions may remain uncalibrated, assay effects may not translate to organisms, and mechanistic paths may contain unsupported intermediate steps. Source-aware graph reasoning may identify these limitations more clearly, but visibility should not be confused with resolution. Likewise, evidence quantity cannot substitute for independence or quality, and agreement among derivative databases cannot be interpreted as experimental replication.

The architecture is not intended to determine clinical treatment, regulatory acceptability, therapeutic superiority, formulation suitability, synthesizability, exposure, safety, or commercial viability. These judgments require evidence and expertise beyond graph topology. Human oversight remains necessary for source selection, relation interpretation, conflict assessment, application framing, and experimental follow-up. Any implementation should therefore define a narrow intended use, prohibited interpretations, required reviewers, and conditions under which graph outputs must not influence a pharmaceutical decision.

Research Agenda and Implementation Priorities

A first research priority is to establish a minimum provenance envelope for therapeutic assertions. This requires comparative evaluation across molecular, bioassay, pathway, disease, mechanism, and study resources to determine which source, version, method, transformation, temporal, and licensing fields are necessary for reproducibility. Parallel work should define a contradiction taxonomy capable of distinguishing direct opposition from differences caused by molecular form, assay design, organism, tissue, disease subtype, endpoint, or temporal context.

A second priority is to develop evaluation designs that move beyond conventional link prediction. Suitable benchmarks should include entity-resolution errors, semantic distortions, duplicated evidence, conflicting claims, incomplete provenance, temporal updates, source ablations, and deliberately misleading graph paths. Prospective studies should test whether the architecture improves expert target review, repurposing assessment, mechanism reconstruction, or evidence triage. Such studies should report when the graph changes a decision, when it merely reorganizes known information, and when it creates additional interpretive burden.

Implementation should proceed in stages. Initial prototypes should focus on a bounded therapeutic area, a limited source set, explicit assertion types, and reproducible versioned releases. Later stages may add federated access, automated extraction, uncertainty models, and cross-platform interoperability after semantic and provenance performance has been demonstrated. Reporting standards should require graph scope, source versions, ontology mappings, evidence classes, inference methods, uncertainty definitions, known omissions, correction policies, and prohibited uses. Expansion should remain conditional on evidence that each added layer preserves meaning rather than merely increasing graph size.

Conclusion

The Therapeutic Evidence Graph is proposed as a pharmaceutical knowledge architecture that connects molecular data, biological mechanisms, disease knowledge, studies, and experimental claims while preserving the provenance and evidential identity of each component. Its defining principle is that therapeutic meaning resides not only in which entities are connected, but also in how a relationship was observed or inferred, which source supplied it, when it applied, which uncertainties qualify it, and whether supporting records are independent. By separating entities, representations, assertions, evidence objects, contradictions, temporal states, and source lineages, the architecture may support more inspectable reasoning for target discovery, repurposing, mechanism assessment, and evidence governance. It cannot independently establish causality, efficacy, safety, clinical utility, or deployment readiness. Its scientific value will depend on rigorous semantic validation, provenance auditing, source-dependence testing, temporal reconstruction, interoperability assessment, application-specific evaluation, and sustained human oversight.

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. Zeng X, Tu X, Liu Y, Fu X, Su Y. Toward better drug discovery with knowledge graph. Curr Opin Struct Biol. 2022;72:114-26. doi:10.1016/j.sbi.2021.09.003
3. MacLean F. Knowledge graphs and their applications in drug discovery. Expert Opin Drug Discov. 2021;16(9):1057-69. doi:10.1080/17460441.2021.1910673

4. Bonner S, Barrett IP, Ye C, Swiers R, Engkvist O, Bender A, et al. A review of biomedical datasets relating to drug discovery: A knowledge graph perspective. Brief Bioinform. 2022;23(6). doi:10.1093/bib/bbac404
5. Mendez D, Gaulton A, Bento AP, Chambers J, De Veij M, Félix E, et al. ChEMBL: Towards direct deposition of bioassay data. Nucleic Acids Res. 2019;47(D1). doi:10.1093/nar/gky1075
6. 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
7. 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
8. Avram S, Bologa CG, Holmes J, Bocci G, Wilson TB, Nguyen DT, et al. DrugCentral 2021 supports drug discovery and repositioning. Nucleic Acids Res. 2021;49(D1). doi:10.1093/nar/gkaa997
9. Jaeger S, Fulle S, Turk S. Mol2vec: Unsupervised machine learning approach with chemical intuition. J Chem Inf Model. 2018;58(1):27-35. doi:10.1021/acs.jcim.7b00616
10. Yang K, Swanson K, Jin W, Coley CW, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370-88. doi:10.1021/acs.jcim.9b00237
11. Schriml LM, Munro JB, Schor M, Olley D, McCracken C, Felix V, et al. The Human Disease Ontology 2022 update. Nucleic Acids Res. 2022;50(D1). doi:10.1093/nar/gkab1063
12. Rodchenkov I, Babur O, Luna A, Aksoy BA, Wong JV, Fong D, et al. Pathway Commons 2019 update: Integration, analysis and exploration of pathway data. Nucleic Acids Res. 2020;48(D1). doi:10.1093/nar/gkz946
13. Xu J, Kim S, Song M, Jeong M, Kim D, Kang J, et al. Building a PubMed knowledge graph. Sci Data. 2020;7(1):205. doi:10.1038/s41597-020-0543-2
14. Breit A, Ott S, Agibetov A, Samwald M. OpenBioLink: A benchmarking framework for large-scale biomedical link prediction. Bioinformatics. 2020;36(13):4097-8. doi:10.1093/bioinformatics/btaa274
15. Zheng S, Rao J, Song Y, Zhang J, Xiao X, Fang EF, et al. PharmKG: A dedicated knowledge graph benchmark for bomedical data mining. Brief Bioinform. 2021;22(4). doi:10.1093/bib/bbaa344
16. Callaghan J, Xu CH, Xin J, Alvarado Cano M, Riutta A, Zhou E, et al. BioThings Explorer: A query engine for a federated knowledge graph of biomedical APIs. Bioinformatics. 2023;39(9). doi:10.1093/bioinformatics/btad570
17. Chandak P, Huang K, Zitnik M. Building a knowledge graph to enable precision medicine. Sci Data. 2023;10(1):67. doi:10.1038/s41597-023-01960-3
18. 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
19. Scalia G, Grambow CA, Pernici B, Li YP, Green WH. Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction. J Chem Inf Model. 2020;60(6):2697-717. doi:10.1021/acs.jcim.9b00975
20. Soleimany AP, Amini A, Goldman S, Rus D, Bhatia SN, Coley CW. Evidential deep learning for guided molecular property prediction and discovery. ACS Cent Sci. 2021;7(8):1356-67. doi:10.1021/acscentsci.1c00546
21. Wallach I, Heifets A. Most ligand-based classification benchmarks reward memorization rather than generalization. J Chem Inf Model. 2018;58(5):916-32. doi:10.1021/acs.jcim.7b00403
22. Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: A benchmark for molecular machine learning. Chem Sci. 2018;9(2):513-30. doi:10.1039/c7sc02664a
23. Mayr A, Klambauer G, Unterthiner T, Steijaert M, Wegner JK, Ceulemans H, et al. Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. Chem Sci. 2018;9(24):5441-51. doi:10.1039/c8sc00148k
24. 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
25. Luo Y, Zhao X, Zhou J, Yang J, Zhang Y, Kuang W, et al. A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information. Nat Commun. 2017;8(1):573. doi:10.1038/s41467-017-00680-8
26. Mohamed SK, Nováček V, Nounu A. Discovering protein drug targets using knowledge graph embeddings. Bioinformatics. 2020;36(2):603-10. doi:10.1093/bioinformatics/btz600
27. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics. 2018;34(13). doi:10.1093/bioinformatics/bty294
28. Gonzalez-Cavazos AC, Tanska A, Mayers M, Carvalho-Silva D, Sridharan B, Rewers PA, et al. DrugMechDB: A curated database of drug mechanisms. Sci Data. 2023;10:632. doi:10.1038/s41597-023-02534-z
29. Wise J, Grebe de Barron A, Splendiani A, Balali-Mood B, Vasant D, Little E, et al. Implementation and relevance of FAIR data principles in biopharmaceutical R&D. Drug Discov Today. 2019;24(4):933-8. doi:10.1016/j.drudis.2019.01.008
30. Wilkinson MD, Sansone SA, Schultes E, Doorn P, da Silva Santos LOB, Dumontier M. A design framework and exemplar metrics for FAIRness. Sci Data. 2018;5:180118. doi:10.1038/sdata.2018.118
31. Wilkinson MD, Dumontier M, Sansone SA, da Silva Santos LOB, Prieto M, Batista D, et al. Evaluating FAIR maturity through a scalable, automated, community-governed framework. Sci Data. 2019;6:174. doi:10.1038/s41597-019-0184-5

32. Unni DR, Moxon SAT, Bada M, Brush M, Bruskiewich R, Clemons PA, et al. Biolink Model: A universal schema for knowledge graphs in clinical, biomedical, and translational science. Clin Transl Sci. 2022;15(8):1848-55. doi:10.1111/cts.13302
33. Caufield JH, Putman T, Schaper K, Unni DR, Hegde H, Callahan TJ, et al. KG-Hub—building and exchanging biological knowledge graphs. Bioinformatics. 2023;39(7). doi:10.1093/bioinformatics/btad418