Pharmacophore

ISSN-2229-5402

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

ONE PHARMACEUTICAL REASONING SPACE FOR MOLECULES, TARGETS, PATHWAYS, PHENOTYPES, DISEASES, AND THERAPEUTIC EVIDENCE

Yifan Wang[1]*, Kristina Hansen[2], Moussa Diarra[3], Henrik Lund[4]

1. ***Department of Unified Pharmaceutical Reasoning, Faculty of Pharmacy, Nanjing Agricultural University, Nanjing, China.***
2. ***Department of Molecular–Target–Pathway Integration, Faculty of Pharmaceutical Sciences, Aarhus University, Aarhus, Denmark.***
3. ***Department of Disease–Phenotype–Therapeutics Reasoning, Faculty of Pharmacy, University of Bamako, Bamako, Mali.***
4. ***Department of Therapeutic Evidence and Decision Space, Faculty of Pharmacy, Lund University, Lund, Sweden.***

ARTICLE INFO

Received:
28 October 2024
Received in revised form:
15 December 2024
Accepted:
16 December 2024
Available online:
28 December 2024

Keywords: Pharmaceutical knowledge representation, Biomedical knowledge graphs, Semantic interoperability, Evidence provenance, Context-qualified assertions, Drug discovery

ABSTRACT

Pharmaceutical knowledge is distributed across chemical, biological, preclinical, clinical, safety, formulation, and regulatory-adjacent data environments that encode different objects, relations, contexts, and evidentiary standards. Although contemporary computational methods can connect parts of this landscape, predictive performance on an isolated task does not establish that the underlying representation is scientifically coherent, evidentially traceable, or suitable for pharmaceutical reasoning. This article proposes a unified pharmaceutical reasoning space as an original knowledge-representation architecture for organizing molecules and drug products, targets, pathways, phenotypes, diseases, and therapeutic evidence within one context-sensitive semantic substrate. The approach is architectural rather than empirical: it synthesizes representational requirements from pharmaceutical data integration, biomedical knowledge graphs, curated molecular resources, pathway knowledgebases, phenotype ontologies, and evidence-aware reasoning. The central contribution is the context-qualified, evidence-bearing assertion, defined as a typed relation between canonical entities that retains provenance, evidence class, biological and experimental context, temporal status, uncertainty, and contradiction without collapsing these dimensions into a single score. The architecture separates source ingestion, semantic representation, reasoning services, validation, and decision use so that graph connectivity, embedding proximity, or ranking performance cannot silently become claims of mechanism, causality, clinical utility, or regulatory acceptability. Its principal limitations are dependence on source quality, incomplete ontological alignment, unresolved contextual gaps, and the absence of empirical validation. The proposed reasoning space may nevertheless provide a disciplined basis for integrating heterogeneous pharmaceutical evidence, comparing competing explanations, and designing future evaluation programs across discovery, safety, and repurposing.

To Cite This Article: Wang Y, Hansen K, Diarra M, Lund H. One Pharmaceutical Reasoning Space for Molecules, Targets, Pathways, Phenotypes, Diseases, and Therapeutic Evidence. Pharmacophore. 2024;15(6):118-26. https://doi.org/10.51847/rWq6JO8bqU

Introduction

Computational pharmaceutical science operates across a chain of questions that cannot be reduced to one data modality or one predictive endpoint. Molecular structure, target engagement, pathway modulation, cellular phenotype, disease biology, formulation, exposure, safety, and clinical observation each contribute different forms of evidence, often at different scales and under different assumptions. Machine-learning applications now span much of this chain, yet their usefulness remains conditional on data quality, validation strategy, and the intended pharmaceutical decision; performance within a benchmark does not by itself establish that a model has captured a biologically or pharmaceutically meaningful relationship [1]. The representational problem is therefore prior to, and broader than, optimization: the system must state what is being connected, under which conditions, on the basis of what evidence, and with what limits on interpretation.

Corresponding Author: Yifan Wang; Department of Unified Pharmaceutical Reasoning, Faculty of Pharmacy, Nanjing Agricultural University, Nanjing, China. E-mail: yifan.wang@njau.edu.cn.

Drug repurposing makes this problem especially visible. A computationally attractive drug–disease link may be derived from shared targets, transcriptional reversal, phenotypic similarity, clinical co-occurrence, or network proximity, but these signals are not scientifically interchangeable. Repurposing progresses only when hypotheses are qualified through mechanistic plausibility, experimental support, clinical evidence, feasibility, safety, and development context, rather than being accepted because one ranking score is high [2]. A representation intended to support pharmaceutical reasoning must therefore preserve the route by which a claim was produced and prevent evidence from one domain from being silently promoted into a stronger claim in another.

Biomedical knowledge graphs offer a practical basis for connecting compounds, proteins, genes, pathways, diseases, phenotypes, and publications through typed relationships. They can support retrieval, path-based exploration, link prediction, and hypothesis generation, but their outputs inherit the limitations of source coverage, graph construction, relation semantics, and downstream validation [3]. The unresolved gap is not simply the absence of a larger graph. Existing systems frequently place heterogeneous records into a common computational container while leaving critical distinctions—such as association versus mechanism, experimental context versus universal truth, and prediction versus confirmation—under-specified.

This article proposes the Unified Pharmaceutical Reasoning Space as an original knowledge-representation architecture. Its purpose is to organize canonical pharmaceutical entities, governed relations, contextual qualifiers, evidence objects, provenance, temporal state, uncertainty, and competing assertions within one reasoning substrate, while keeping representation, inference, validation, and decision authority distinct. The proposal builds on the complementary need for cross-stage evidence integration, translational triangulation, and graph-based organization [1–3]. It does not claim empirical validation, universal ontology coverage, clinical utility, regulatory acceptance, or deployment readiness. Its intended contribution is a testable conceptual specification for asking whether a pharmaceutical claim is semantically well formed, evidentially traceable, contextually bounded, and appropriate for a defined reasoning task.

Fragmented Representations across Pharmaceutical Data Systems

Fragmentation in pharmaceutical data systems is not merely the physical separation of databases. Resources differ in the entities they include, the identifiers they privilege, the granularity of their concepts, the semantics assigned to relations, the provenance they retain, their licensing conditions, and the frequency and transparency of updates. Comparative analyses of drug-discovery datasets show that overlapping resources may represent nominally similar drugs, genes, diseases, pathways, or interactions in materially different ways, creating integration problems that affect reproducibility as well as coverage [4]. A unified space must therefore preserve source-specific meaning and version history rather than treating ingestion as lossless consolidation.

Fragmentation also originates in how assertions are created. Expert curation, structured database import, text mining, experimental pipelines, and model-derived inference produce records with different error structures and evidentiary standing. Biomedical knowledge-graph construction choices determine which entities become nodes, which statements become edges, how uncertainty is represented, and what downstream learning methods can legitimately recover [5]. An extracted sentence, a curated interaction, and an embedding-predicted link may all be computationally traversable, but they should not be represented as epistemically equivalent. The proposed architecture consequently separates source acquisition and semantic normalization from the later reasoning services that rank or infer relationships.

Clinical and omics evidence demonstrates why context cannot be treated as optional metadata. Interpreting proteomic observations may require explicit links among projects, participants, biospecimens, assays, measured proteins, tissues, diseases, drugs, analytical workflows, and publications [6]. Removing these links can transform a condition-specific observation into an apparently general biological statement. Across pharmaceutical resources, fragmentation therefore arises from at least three non-substitutable dimensions: heterogeneous source coverage, heterogeneous construction routes, and heterogeneous experimental or clinical context [4–6]. The proposed reasoning space addresses these dimensions by retaining source records and representing their assertions through a shared semantic form without erasing the circumstances under which the assertions were generated.

Disease-associated evidence adds a further problem: even when sources describe the same broad relation, they may rely on different disease vocabularies, gene identifiers, curation practices, literature-extraction methods, and evidence scores. Integrated disease-genomics resources show the importance of standardized identifiers, controlled vocabularies, provenance, and evidence qualification for keeping gene–disease and variant–disease assertions interpretable [7]. The architectural implication is that integration should produce a network of traceable, context-qualified assertions rather than a flat collection of de-duplicated edges. Data availability is not equivalent to data suitability, and semantic alignment is not evidence that two records carry the same scientific meaning. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop fragmented representations across pharmaceutical data systems within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Fragmented representations across pharmaceutical data systems

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Source coverage and scope	Which pharmaceutical entities and relations	Databases are designed for different scientific purposes and therefore provide	Defines fragmentation as a scientific coverage problem	Apparent completeness may reflect source density	A large integrated graph cannot be assumed to represent the full

	are represented, omitted, or disproportionately detailed?	unequal coverage of molecules, targets, pathways, diseases, phenotypes, products, and evidence.	rather than only a technical integration problem.	rather than biological or therapeutic completeness.	pharmaceutical evidence landscape.
Identity and canonicalization	When do records from different sources denote the same entity, and when must they remain distinct?	Molecules, salts, stereoisomers, active moieties, drug products, genes, proteins, disease concepts, and phenotype terms may use overlapping or non-equivalent identifiers.	Establishes canonical entities while preserving mappings to source-specific records.	Incorrect merging can fabricate relations; excessive separation can conceal valid cross-source evidence.	Canonicalization is a governed mapping decision, not proof of scientific equivalence.
Relation semantics	What precisely does each connection assert?	Relations such as binds, inhibits, associates with, treats, expresses, participates in, and is-part-of carry different directionality and evidentiary implications.	Requires typed predicates with explicit domain, range, direction, and permissible qualifiers.	Generic edges can silently upgrade association into mechanism, causality, or therapeutic effect.	Graph connectivity alone does not determine the scientific meaning or strength of a relation.
Experimental and clinical context	Under what species, tissue, cell, assay, dose, route, disease state, population, and time conditions was the assertion observed?	Pharmaceutical observations are conditional on biological and study context.	Motivates a context envelope attached to every decision-relevant assertion.	Missing qualifiers can make a local observation appear universally applicable.	An insufficiently contextualized assertion may be searchable but is not necessarily suitable for reasoning or decision support.
Evidence type and provenance	What evidence supports the assertion, and how was it curated, extracted, or inferred?	Curated databases, literature statements, assays, clinical observations, and computational predictions have different evidentiary status.	Makes evidence and provenance first-class representational objects.	A common graph format may create false equivalence among fundamentally different evidence classes.	Provenance identifies origin and processing history; it does not guarantee correctness.
Version, update, and licensing state	Which release, access conditions, and update cycle govern the record?	Source content and availability change over time, and integration may be constrained by licensing or release practices.	Supports reproducibility, auditability, and controlled graph reconstruction.	Unversioned integration can produce irreproducible paths or retain superseded claims.	A reasoning result is interpretable only relative to the source versions and permissions from which it was derived.
Data suitability and representational fitness	Is the available information adequate for the intended reasoning task?	Coverage, context completeness, evidence quality, missingness, and semantic compatibility determine task suitability.	Prevents the presence of data from being treated as sufficient justification for its use.	Dense but poorly qualified data can dominate inference and create false confidence.	Availability, interoperability, and computational accessibility do not alone establish pharmaceutical suitability.

Core Entities, Relations, Contexts, and Evidence Types

The proposed reasoning space begins with canonical entities, but canonicalization must not flatten pharmaceutically important distinctions. A “drug” may refer to a molecular structure, active moiety, stereochemical form, salt, metabolite, investigational substance, approved product, dose form, or combination product. Drug-centered resources connect structures with targets, enzymes, transporters, mechanisms, interactions, products, and pharmaceutical information, demonstrating why a drug cannot be modeled adequately as one undifferentiated node [8]. The architecture therefore distinguishes molecule-level identity from product- and use-level identity while retaining explicit mappings among them.

Relations require equally disciplined semantics. Protein networks may combine experimentally determined interactions, curated pathway associations, co-expression, genomic context, text-derived evidence, and computational predictions. These channels can support functional discovery, but an association score is not equivalent to a direct physical interaction, causal influence, or therapeutically actionable target relationship [9]. In the proposed architecture, every predicate has a defined meaning, direction, domain, range, and permitted qualifiers. “Binds,” “inhibits,” “is associated with,” “participates in,” and “treats” must remain separate relation types even when they connect the same broad classes of entities.

Pathways should not be reduced to labels appended to genes or targets. Curated pathway resources represent biological processes as nested events and reaction-level structures involving participants, locations, inputs, outputs, and supporting literature [10]. This event-based view enables a target to be situated within a biological sequence rather than merely assigned to a pathway category. Drugs, protein relations, and pathways consequently require different entity and relation semantics, which a single generic node–edge convention cannot preserve without explicit qualification [8–10]. The proposed atomic knowledge unit is therefore a context-qualified, evidence-bearing assertion: a canonical subject and object connected by a governed predicate, accompanied by the context, evidence type, provenance, and assertion state required to interpret the claim.

Phenotypes extend the architecture from molecular and pathway representation to observable biological states. Their meaning depends on ontology alignment, species, genotype, developmental stage, disease context, observation method, and supporting evidence. Cross-species phenotype platforms show that normalized phenotypic descriptions can connect model-organism observations to human disease hypotheses, while also requiring careful control of species and evidence context [11]. Within the proposed reasoning space, phenotype similarity may support comparison or prioritization, but it does not establish that

mechanisms, therapeutic effects, or safety findings will translate across species or populations. Evidence types are therefore attached to assertions as explicit qualifiers, and reasoning services must return not only a ranked relationship but also the traceable evidentiary path and the contexts under which that relationship can be considered.

Architecture of the Unified Pharmaceutical Reasoning Space

The proposed Unified Pharmaceutical Reasoning Space is organized as a layered semantic environment rather than a single undifferentiated graph. Its first layer preserves source records, identifiers, versions, and extraction routes; its second resolves those records to canonical pharmaceutical entities; and its third represents typed, context-qualified, evidence-bearing assertions. Systematic integration of drugs, diseases, genes, pathways, adverse effects, and other biomedical concepts has shown that heterogeneous evidence can be connected for hypothesis prioritization when relation types and provenance remain recoverable [12]. The present architecture extends that principle by making traceability and interpretive boundaries mandatory properties of every reasoning path.

The canonical semantic substrate contains distinct entity classes for molecular structures, active moieties, drug products, targets, pathways, phenotypes, diseases, biological samples, interventions, study populations, and evidence objects. Relations are governed through explicit schemas rather than inferred solely from node proximity. Large biomedical knowledge resources demonstrate the value of standardized entity categories, relation vocabularies, source attribution, and reusable graph structures for precision-medicine reasoning [13]. Nevertheless, standardization does not imply that every source assertion is equivalent. The architecture therefore retains the original source representation alongside its canonical mapping, allowing users to inspect where semantic harmonization required interpretation.

Above the semantic substrate sits a reasoning-services layer that supports retrieval, path analysis, evidence reconciliation, link prediction, and hypothesis prioritization. Knowledge-graph embeddings can identify latent structural patterns and prioritize potential protein targets, but the resulting scores reflect graph topology, training assumptions, negative-sampling strategies, and source composition [14]. In the proposed reasoning space, an embedding-generated link is stored as a predicted assertion rather than merged with experimentally established evidence. Its output must retain the model version, training graph, inference date, uncertainty representation, and supporting or conflicting paths.

The architecture is completed by a governance and decision-use boundary. Target-oriented knowledge graphs illustrate how compounds, genes, diseases, pathways, expression evidence, and other biomedical relations can be assembled for target discovery [15]. The proposed contribution adds a requirement that any discovery, safety, or repurposing output remain linked to its evidence path and intended context of use. Human reviewers must be able to distinguish observed, curated, extracted, predicted, mechanistically supported, experimentally confirmed, and clinically evaluated assertions. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop architecture of the unified pharmaceutical reasoning space within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Architecture of the unified pharmaceutical reasoning space

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Source-preservation layer	Database records, publications, assays, omics studies, clinical observations, safety resources, formulation and exposure records	Retain original identifiers, source semantics, release information, licensing state, and extraction history	Versioned source objects that can be reconstructed and audited	Source incompleteness, outdated records, access restrictions, and undocumented processing	Reproducible ingestion testing, source-version verification, and record-level trace audits
Identity-resolution layer	Source-specific molecule, product, gene, protein, pathway, phenotype, and disease identifiers	Map records to canonical entities while preserving distinctions relevant to pharmaceutical interpretation	Canonical entities with explicit mappings to source records	Incorrect merging, unresolved synonyms, stereochemical ambiguity, and granularity mismatch	Expert-reviewed mapping samples, competency tests, and error analysis by entity class
Typed-relation layer	Curated predicates, extracted relations, experimentally reported interactions, and model-derived associations	Specify relation meaning, directionality, domain, range, and permitted qualifiers	Semantically governed assertions rather than generic graph edges	Predicate ambiguity and inconsistent relation granularity across sources	Relation-conformance checks and domain-specific expert review
Context layer	Species, tissue, cell type, genotype, disease state, assay, dose, route, population, formulation, exposure, and time	Define the conditions under which an assertion can be interpreted	Context-qualified assertions with explicit applicability limits	Missing or incompatible contextual metadata	Context-completeness assessment and task-specific suitability review
Evidence and provenance layer	Curated database evidence, publications, experiments, clinical observations, text extraction, and computational predictions	Record how each assertion was generated and what evidence class supports it	Traceable evidence objects linked to claims and sources	Provenance does not itself establish accuracy; evidence classes may remain heterogeneous	Evidence-type validation, source-path reconstruction, and curator agreement studies
Scientific-state layer	Hierarchical relations, causal claims, temporal states, contradiction sets,	Prevent association, prediction, mechanism, causality, and	Explicit claim-state and epistemic qualifiers	Some distinctions require expert interpretation and cannot	Domain-specific adjudication protocols and longitudinal update testing

	uncertainty, and version histories	confirmation from being treated as equivalent	be resolved by schema alone		
Reasoning-services layer	Governed graph assertions, evidence paths, embeddings, rules, and task definitions	Support discovery, safety, repurposing, evidence reconciliation, and gap identification	Ranked, provenance-linked hypotheses or evidence summaries	Model bias, graph incompleteness, shortcut learning, and task mismatch	Benchmark, retrospective, prospective, and context-of-use evaluation
Human-governance layer	Reasoning outputs, provenance trails, conflicts, uncertainties, and intended decisions	Control interpretation, escalation, approval, and communication of outputs	Bounded recommendations for expert consideration	Review burden, automation bias, and inconsistent institutional practices	Usability studies, decision-impact evaluation, auditability testing, and governance review

Figure 1 presents the unified pharmaceutical reasoning galaxy, showing how the article's key components and boundaries are connected within architecture of the unified pharmaceutical reasoning space.

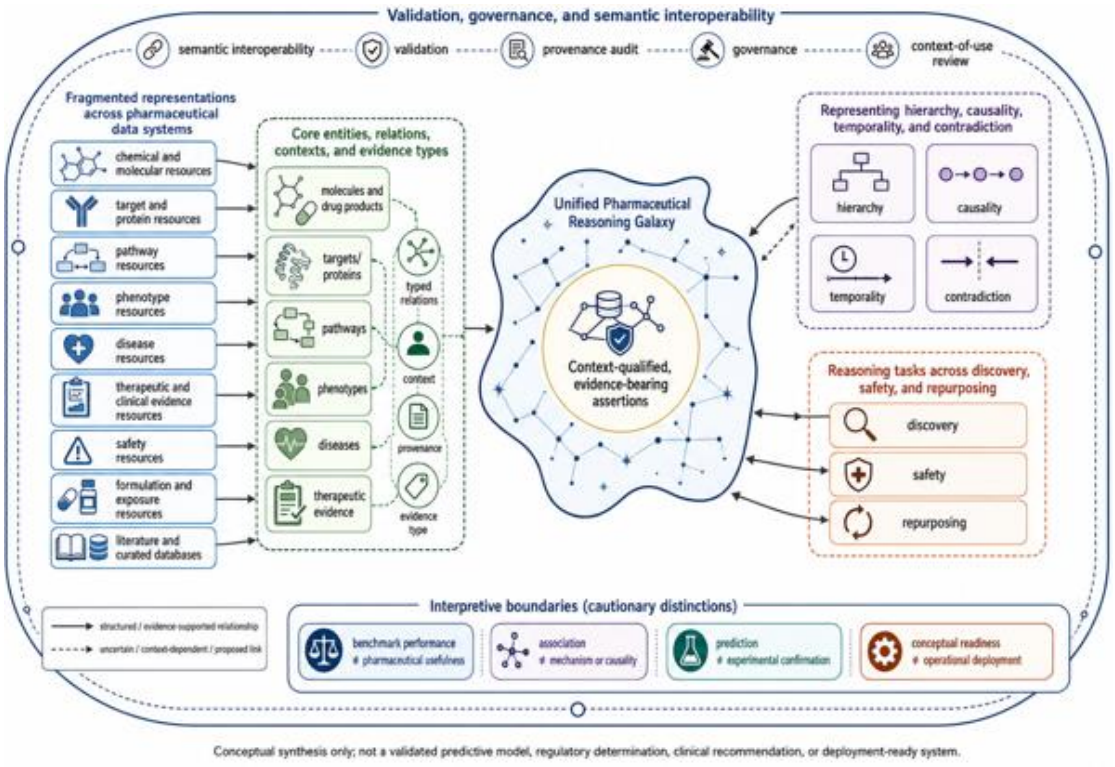

Figure 1. Unified Pharmaceutical Reasoning Galaxy.

The figure is an original conceptual synthesis that organizes the article's central contribution across fragmented representations across pharmaceutical data systems, core entities, relations, contexts, and evidence types, core entities, evidence types, architecture of the unified pharmaceutical reasoning space, representing hierarchy, causality, temporality, and 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.

Representing Hierarchy, Causality, Temporality, and Contradiction

Hierarchy determines how reasoning moves between broad and specific concepts. Disease ontologies differ in scope, classification principles, parent–child structures, and mappings to clinical or biological vocabularies [16]. Treating these structures as interchangeable may create false equivalence between a general disease category, a molecular subtype, a clinical manifestation, and a phenotype. The proposed architecture therefore stores hierarchical relations as typed assertions with ontology provenance. Inference through a hierarchy is permitted only when the relation semantics justify inheritance, and the system must identify whether a result depends on exact identity, descendant expansion, or cross-ontology mapping.

Causal status must be represented separately from association and prediction. Predictive models may identify variables or graph patterns associated with an outcome without establishing that changing those variables would change the outcome. Causal and counterfactual reasoning requires explicit assumptions about interventions, confounding, selection, time ordering, and transportability [17]. The reasoning space therefore assigns distinct states to statistical association, predicted relation, mechanistic hypothesis, experimentally supported mechanism, and intervention-linked causal claim. No graph path or embedding score may automatically promote an assertion from one state to another.

Temporality affects both biological meaning and evidentiary validity. Disease progression can follow population- and sequence-dependent trajectories in which one condition precedes, follows, or modifies the interpretation of another [18]. Pharmaceutical assertions similarly change with treatment stage, exposure duration, resistance, disease progression, product reformulation, new safety findings, and database updates. Each assertion should consequently carry observation time, validity interval, source release, and supersession status where available. A current reasoning result must remain distinguishable from a historically valid claim or an assertion derived from an obsolete source version.

Contradiction is modeled as a retained relationship among assertions rather than as an error to be removed automatically. Evidence ontologies provide controlled terms for describing experimental, computational, author-statement, and curator-inference evidence [19]. These annotations allow apparently conflicting statements to be compared according to their evidence types and contexts. Opposing findings from different tissues, species, doses, formulations, populations, or time points may represent context dependence rather than direct falsification. The architecture therefore creates contradiction sets that preserve all relevant assertions, expose contextual differences, and defer truth adjudication when the available evidence cannot support a justified resolution.

Reasoning Tasks across Discovery, Safety, and Repurposing

Safety reasoning requires representations that can distinguish individual-drug effects from combination-specific outcomes. Graph-convolutional modeling of polypharmacy illustrates that different drug pairs may be associated with different adverse-effect types and that relation-specific structure matters [20]. Within the proposed architecture, a safety query should therefore include the relevant drug forms, co-medications, patient or experimental context, exposure conditions, adverse-event definition, evidence source, and temporal window. A predicted interaction is a signal for investigation, not confirmation that an adverse effect will occur in a patient or population.

Repurposing reasoning can draw on multilayer relationships among compounds, targets, pathways, diseases, phenotypes, and existing therapeutic evidence. Knowledge-graph learning can extend association-based reasoning across multiple biomedical layers and identify candidate drug–disease relationships [21]. The proposed reasoning space adds evidence-path comparison: candidates can be ranked not only by model output but also by whether their paths converge through independent evidence types, depend on one highly connected entity, conflict with safety information, or lack disease-context specificity. Such ranking remains hypothesis-generating and does not establish therapeutic efficacy.

Human-centered repurposing also requires that outputs correspond to clinically meaningful questions and expose the evidence required for expert interpretation. Clinician-centered modeling demonstrates the importance of evaluating repurposing systems in relation to realistic therapeutic indications and user-facing evidence [22]. The proposed architecture therefore separates model-level ranking from decision-level readiness. A candidate may have high computational priority while remaining unsuitable because of contraindications, formulation constraints, route requirements, exposure limitations, weak mechanistic support, or absence of relevant clinical evidence.

Discovery reasoning can likewise support target prioritization by combining genetic, expression, pathway, disease, tractability, and pharmacological information. Machine-learning approaches within integrated target platforms show how heterogeneous evidence features can assist identification of target–disease associations [23]. In the proposed reasoning space, target discovery produces a structured hypothesis package rather than a single score. The package should identify supporting and opposing evidence, context gaps, causal status, known compounds, pathway position, safety liabilities, and the experimental observations needed to advance or reject the hypothesis.

Validation, Governance, and Semantic Interoperability

Semantic interoperability requires more than exchanging graph files. Shared schemas can align biomedical entity categories and predicates so that independently constructed knowledge graphs can be queried and combined more consistently [24]. The proposed reasoning space adopts this principle while requiring explicit mappings from local source semantics to the shared model. Interoperability testing should assess whether entities, predicates, qualifiers, and evidence objects retain their intended meaning after translation, not merely whether the transferred records satisfy a technical syntax.

Computational validation must also separate graph-learning performance from pharmaceutical validity. Benchmarking frameworks enable standardized comparison of biomedical link-prediction methods through shared datasets, tasks, and evaluation procedures [25]. Such benchmarks are valuable for reproducibility but cannot establish that a predicted link is mechanistically correct, experimentally reproducible, clinically useful, or safe. Evaluation should therefore include semantic correctness, leakage control, temporal holdouts, calibration, evidence-path fidelity, context sensitivity, contradiction handling, and task-specific utility in addition to conventional ranking measures.

Governance begins with stewardship of the underlying data and metadata. FAIR-oriented evaluation provides a basis for assessing whether digital resources are findable, accessible, interoperable, and reusable through transparent indicators [26]. For a pharmaceutical reasoning space, these properties should be extended to versioned identifiers, machine-readable provenance, evidence-type documentation, licensing conditions, update responsibility, correction mechanisms, and reproducible transformations. FAIRness improves the conditions for responsible reuse but does not guarantee scientific quality, lack of bias, or suitability for a particular pharmaceutical decision.

Operational reproducibility depends on repeatable ingestion, normalization, canonicalization, semantic standardization, and graph construction. Translational knowledge-graph systems demonstrate how heterogeneous biomedical sources can be

processed into a semantically standardized resource through documented build procedures [27]. The proposed architecture additionally requires change logs, mapping tests, conflict reports, and output-level audit trails. Governance should assign responsibility for ontology changes, source retirement, model updates, contradiction escalation, and expert approval. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop validation, governance, and semantic interoperability within the article's central argument.

Table 3. Evaluation, implementation, and research priorities arising from One Pharmaceutical Reasoning Space for Molecules, Targets, Pathways, Phenotypes, Diseases, and Therapeutic

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Entity identity and canonicalization	Whether source records are mapped to the correct molecular, product, target, pathway, phenotype, or disease entities	Expert-reviewed mapping samples, identity competency questions, stereochemical and product-form test cases	Incorrect merging, unresolved ambiguity, or loss of pharmaceutically relevant distinctions	High mapping accuracy in one entity class may not generalize to others	Determines whether downstream evidence is attached to the correct object
Relation semantics	Whether predicates retain their intended direction, domain, range, and scientific meaning	Schema-conformance tests, curator review, relation-specific error analysis	Generic or misdirected edges that conflate binding, association, treatment, mechanism, or causality	Formal validity does not establish factual correctness	Controls the claims that reasoning paths are permitted to support
Context completeness	Whether decision-relevant assertions retain species, tissue, assay, dose, route, formulation, population, disease-state, and time qualifiers	Context audits, missingness analysis, task-based competency questions	Outputs depend on assertions lacking critical applicability conditions	Some historical or literature-derived evidence may not report sufficient context	Determines whether evidence is suitable for a defined pharmaceutical question
Evidence and provenance fidelity	Whether each assertion can be traced to its source, evidence class, extraction method, and version	Source reconstruction, provenance-path testing, curator agreement analysis	Broken source links, undocumented transformations, or evidence-class mislabeling	Traceability does not guarantee reliability or reproducibility of the original evidence	Supports audit, correction, and evidence-strength assessment
Hierarchy, causality, time, and contradiction	Whether the system preserves distinct claim states and avoids invalid inheritance or causal promotion	Ontology tests, causal-state audits, temporal holdouts, contradiction case studies	Association presented as mechanism, obsolete assertions treated as current, or conflicts silently removed	Schema alone cannot adjudicate every scientific disagreement	Protects users from interpreting computational organization as scientific proof
Reasoning reproducibility	Whether equivalent queries, graph versions, models, and parameters produce reconstructable outputs	Version-controlled benchmark runs, provenance-linked query logs, deterministic build checks	Irreproducible rankings or outputs with unrecorded dependencies	Reproducibility does not establish usefulness	Enables comparison, debugging, and accountable model updating
Predictive and retrieval performance	Whether the system retrieves or ranks relevant relationships under controlled tasks	Link-prediction benchmarks, information-retrieval assessment, temporal validation, calibrated uncertainty analysis	Leakage, shortcut learning, unstable ranking, or poor calibration	Benchmark performance is conditional on task design and graph composition	Supports model selection but not automatic pharmaceutical decision making
Pharmaceutical task utility	Whether outputs improve target review, safety assessment, repurposing analysis, or evidence-gap detection	Expert comparison studies, retrospective case analysis, prospective silent evaluation, workflow studies	Outputs are accurate in benchmark terms but irrelevant, unactionable, or misleading in practice	Utility may differ by institution, disease area, evidence availability, and user expertise	Determines whether the architecture may support a defined context of use
FAIRness and accessibility	Whether data, metadata, schemas, and provenance can be discovered, accessed, exchanged, and reused appropriately	FAIR-oriented indicators, metadata inspection, licensing review, interoperability exercises	Unresolvable identifiers, inaccessible evidence, undocumented restrictions, or non-reusable mappings	FAIR properties do not establish unbiased or scientifically adequate content	Influences reproducibility, collaboration, and long-term maintenance
Governance and human oversight	Whether responsibilities, review thresholds, update controls, and escalation procedures are explicit	Governance audit, role mapping, incident review, usability assessment, decision-impact evaluation	Unowned data changes, automation bias, absent conflict escalation, or undocumented overrides	Governance quality depends on institutional capacity and cannot be inferred from software design	Defines whether outputs remain advisory, auditable, and appropriately bounded

Limitations and Boundary Conditions

The Unified Pharmaceutical Reasoning Space is a proposed architecture, not a complete ontology or validated implementation. Its representational quality would depend on the coverage, curation practices, update cycles, and licensing constraints of its contributing sources. Biomedical datasets differ substantially in scope and semantics [4]. Disease ontologies also embody different classification choices [16]. Canonicalization may therefore remain contested, particularly for molecular forms, drug products, disease subtypes, phenotypic manifestations, and relations that lack direct cross-source equivalence.

The architecture cannot convert observational association into causality or compensate automatically for missing experimental context. Causal interpretation requires assumptions and study designs beyond graph structure [17]. Temporal ordering can strengthen interpretation, but longitudinal patterns may still reflect confounding, healthcare-system processes, or population-specific factors [18]. Contradictions may be organized and exposed, yet semantic representation alone cannot determine which finding is true when studies differ in quality, population, assay, dose, formulation, or endpoint.

The reasoning space also does not establish synthesizability, developability, safety, efficacy, clinical utility, or regulatory acceptability. A ranked target, interaction, or repurposing candidate remains conditional on the evidence used to generate it. Clinician-centered evaluation may reveal whether outputs correspond to meaningful user questions [22], while graph benchmarks can test reproducible computational behavior [25]. Neither form of evaluation alone justifies operational deployment. Human review, experimental confirmation, prospective assessment, and institution-specific governance remain necessary.

Research Agenda and Implementation Priorities

The first research priority is representational validation. Studies should test whether pharmaceutical experts can map source records consistently to canonical entities, distinguish predicates reliably, and identify which contextual qualifiers are necessary for specific reasoning tasks. Shared schemas provide a starting point for cross-resource alignment [24], but domain-specific extensions will be needed for molecular form, drug product, formulation, route, exposure, assay, and therapeutic-use context. Reproducible graph-building procedures should be evaluated through independently reconstructed releases, mapping-error audits, and change-impact analyses [27].

The second priority is multidimensional reasoning evaluation. Future studies should compare unqualified graph inference with context-aware, provenance-aware, temporally constrained, and contradiction-preserving alternatives. Standard link-prediction benchmarks can test controlled computational properties [25], but prospective silent evaluations and expert-review studies are needed to determine whether evidence paths improve pharmaceutical interpretation. Clinician-centered repurposing work provides one model for aligning evaluation with intended users [22]. Similar designs are needed for target selection, safety analysis, formulation-linked reasoning, and evidence-gap detection.

The third priority is staged implementation. Early systems should begin with bounded scientific questions, transparent source sets, high-value entity classes, and explicit exclusion rules rather than attempting universal coverage. Evidence-type annotation can support controlled comparison of claims [19], while FAIR-oriented assessments can guide stewardship and reuse [26]. Progression from prototype to decision-support evaluation should require predefined semantic, provenance, calibration, utility, and governance criteria. Operational influence should expand only when evidence demonstrates suitability for the exact context of use.

Conclusion

The Unified Pharmaceutical Reasoning Space is proposed as a context-sensitive knowledge-representation architecture connecting molecules and drug products, targets, pathways, phenotypes, diseases, and therapeutic evidence without collapsing their distinct meanings or evidentiary status. Its central unit is the context-qualified, evidence-bearing assertion, embedded within explicit layers for identity, relation semantics, provenance, hierarchy, causal status, temporality, contradiction, uncertainty, and governance. The architecture permits discovery, safety, and repurposing reasoning while keeping graph connectivity, model performance, mechanistic interpretation, experimental confirmation, clinical usefulness, and regulatory acceptability separate. Its value must be established through semantic, computational, scientific, workflow, and governance evaluation rather than one performance measure. The contribution is therefore a bounded and testable architectural proposal for organizing pharmaceutical knowledge, not a validated model or deployment-ready decision system.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
2. Pushpakom S, Iorio F, Eyers PA, Escott KJ, Hopper S, Wells A, et al. Drug repurposing: Progress, challenges and recommendations. Nat Rev Drug Discov. 2019;18(1):41-58. doi:10.1038/nrd.2018.168
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. 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
6. Santos A, Colaço AR, Nielsen AB, Niu L, Strauss M, Geyer PE, et al. A knowledge graph to interpret clinical proteomics data. Nat Biotechnol. 2022;40(5):692-702. doi:10.1038/s41587-021-01145-6
7. Piñero J, Bravo À, Queralt-Rosinach N, Gutiérrez-Sacristán A, Deu-Pons J, Centeno E, et al. DisGeNET: A comprehensive platform integrating information on human disease-associated genes and variants. Nucleic Acids Res. 2017;45(D1). doi:10.1093/nar/gkw943
8. 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
9. Szklarczyk D, Gable AL, Lyon D, Junge A, Wyder S, Huerta-Cepas J, et al. STRING v11: Protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. Nucleic Acids Res. 2019;47(D1). doi:10.1093/nar/gky1131
10. Jassal B, Matthews L, Viteri G, Gong C, Lorente P, Fabregat A, et al. The Reactome pathway knowledgebase. Nucleic Acids Res. 2020;48(D1). doi:10.1093/nar/gkz1031
11. Mungall CJ, McMurry JA, Köhler S, Balhoff JP, Borromeo C, Brush M, et al. The Monarch Initiative: An integrative data and analytic platform connecting phenotypes to genotypes across species. Nucleic Acids Res. 2017;45(D1). doi:10.1093/nar/gkw1128
12. 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
13. 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
14. 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
15. Zhou C, Cai CP, Huang XT, Wu S, Yu JL, Wu JW, et al. TarKG: A comprehensive biomedical knowledge graph for target discovery. Bioinformatics. 2024;40(10). doi:10.1093/bioinformatics/btae598
16. Haendel MA, McMurry JA, Relevo R, Mungall CJ, Robinson PN, Chute CG. A census of disease ontologies. Annu Rev Biomed Data Sci. 2018;1:305-31. doi:10.1146/annurev-biodatasci-080917-013459
17. Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. Nat Mach Intell. 2020;2(7):369-75. doi:10.1038/s42256-020-0197-y
18. Siggaard T, Reguant R, Jørgensen IF, Haue AD, Lademann M, Aguayo-Orozco A, et al. Disease trajectory browser for exploring temporal, population-wide disease progression patterns in 7.2 million Danish patients. Nat Commun. 2020;11(1):4952. doi:10.1038/s41467-020-18682-4
19. Nadendla S, Jackson R, Munro J, Quaglia F, Mészáros B, Olley D, et al. ECO: The Evidence and Conclusion Ontology, an update for 2022. Nucleic Acids Res. 2022;50(D1). doi:10.1093/nar/gkab1025
20. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics. 2018;34(13). doi:10.1093/bioinformatics/bty294
21. Bang D, Lim S, Lee S, Kim S. Biomedical knowledge graph learning for drug repurposing by extending guilt-by-association to multiple layers. Nat Commun. 2023;14(1):3570. doi:10.1038/s41467-023-39301-y
22. Huang K, Chandak P, Wang Q, Havaldar S, Vaid A, Leskovec J, et al. A foundation model for clinician-centered drug repurposing. Nat Med. 2024;30(12):3601-13. doi:10.1038/s41591-024-03233-x
23. Han Y, Klinger K, Rajpal DK, Zhu C, Teeple E. Empowering the discovery of novel target-disease associations via machine learning approaches in the Open Targets platform. BMC Bioinformatics. 2022;23(1):232. doi:10.1186/s12859-022-04753-4
24. Unni DR, Moxon SAT, Bada M, Brush M, Bruskiewich R, Clemons P, 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
25. 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
26. Wilkinson MD, Sansone SA, Schultes E, Doorn P, Bonino da Silva Santos LO, Dumontier M. A design framework and exemplar metrics for FAIRness. Sci Data. 2018;5:180118. doi:10.1038/sdata.2018.118
27. Wood EC, Glen AK, Kvarfordt LG, Womack F, Acevedo L, Yoon TS, et al. RTX-KG2: A system for building a semantically standardized knowledge graph for translational biomedicine. BMC Bioinformatics. 2022;23(1):400. doi:10.1186/s12859-022-04932-3