Traditional Medicine Knowledge Graphs for Drug Discovery: Cultural Context, Evidence Grading, Mechanistic Links, and Benefit-Sharing Logic

Daniel Fischer[1]*, Laura Meier[1], Stefan Koch[2]

[1]Department of AI for Anti-diabetic Flavonoid Glycosides, Faculty of Pharmacy, University of Freiburg, Freiburg, Germany

[2]Department of Computational Pharmacology for Insulin Sensitizers, Faculty of Pharmacy, Technical University of Munich, Munich, Germany.

ABSTRACT

Traditional medicine knowledge has long informed natural-product research, yet its representation in computational drug discovery raises scientific, cultural, and governance challenges. Knowledge graphs may help organize heterogeneous information concerning traditional uses, botanical identity, preparation context, phytochemical composition, bioactivity, molecular targets, pathways, safety, and literature provenance. However, graph construction can also detach knowledge from its cultural setting, merge incompatible evidence types, expose sensitive information, overstate mechanistic certainty, or support extractive research practices. This article proposes a responsible knowledge framework for traditional medicine knowledge graphs in drug discovery. The framework places cultural context preservation, source provenance, sensitive knowledge triage, evidence grading, citation traceability, uncertainty annotation, and conflicting-evidence management before mechanistic interpretation or candidate prioritization. It further distinguishes reported traditional use, published ethnopharmacological evidence, experimental bioactivity, computational inference, mechanistic hypotheses, safety evidence, and validated therapeutic claims. Benefit-sharing logic, access governance, ethical sourcing, community-interest recognition, expert and stakeholder review, and feedback updating are incorporated as explicit framework dimensions rather than treated as external considerations. The proposed architecture is intended to support respectful, transparent, and evidence-aware hypothesis generation. It does not establish therapeutic efficacy, confirm biological mechanism, certify ethical or legal compliance, represent community approval, or demonstrate commercialization readiness. Its principal contribution is an integrated conceptual structure for connecting traditional medicine knowledge with modern drug-discovery evidence while preserving epistemic boundaries, cultural meaning, governance responsibilities, and validation requirements.

Key Words: *Traditional medicine, Knowledge graphs, Drug discovery, Cultural context, Evidence grading, Mechanistic links*

eIJPPR 2026; 16(3):1-19

HOW TO CITE THIS ARTICLE: Fischer D, Meier L, Koch S. Traditional Medicine Knowledge Graphs for Drug Discovery: Cultural Context, Evidence Grading, Mechanistic Links, and Benefit-Sharing Logic. Int J Pharm Phytopharmacol Res. 2026;16(3):1-19. https://doi.org/10.51847/2kl4U4xqG3

INTRODUCTION

Traditional medicine knowledge has contributed substantially to the identification of natural products, molecular scaffolds, and biologically active substances that have influenced modern pharmacotherapy. Its relevance to drug discovery arises not only from specific materials but also from accumulated observations concerning preparation, administration, therapeutic purpose, ecological availability, and perceived biological effects. Nevertheless, historical use and natural-product origin must not be interpreted as proof of contemporary therapeutic efficacy. Reviews of natural-product-derived medicines demonstrate that natural sources remain

Corresponding author: Daniel Fischer
Address: Department of AI for Anti-diabetic Flavonoid Glycosides, Faculty of Pharmacy, University of Freiburg, Freiburg, Germany.
E-mail: ✉ daniel.fischer@pharmazie.uni-freiburg.de
Received: 19 February 2026; **Revised:** 28 April 2026; **Accepted:** 30 April 2026

This is an **open access** journal, and articles are distributed under the terms of the Creative Commons Attribution-Non Commercial-ShareAlike 4.0 License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.

important to pharmaceutical innovation, while also showing that successful translation depends on extensive chemical, pharmacological, toxicological, and clinical investigation [1]. Traditional medicine knowledge should therefore be understood as a potential source of research questions and discovery hypotheses rather than as an independent validation system. This distinction is especially important when information is converted into computational formats, because structured representations may make associations appear more precise, standardized, or conclusive than the underlying evidence permits.

The scientific interpretation of traditional medicine materials is further complicated by variability in biological identity, preparation methods, chemical composition, extraction procedures, experimental models, assay design, and reporting quality. A traditional preparation cannot be treated as interchangeable with an isolated constituent, a commercially sourced extract, or a computationally predicted compound set unless the relationship between those materials is documented. Best-practice recommendations in phytopharmacological research emphasize botanical authentication, reproducible preparation, adequate chemical characterization, suitable controls, transparent methods, and cautious interpretation of pharmacological and toxicological results [2]. These requirements are directly relevant to computational knowledge representation. A graph that connects a traditional use to a compound, target, pathway, or safety outcome without preserving material identity and evidentiary context may generate misleading mechanistic narratives. Responsible graph development must consequently begin with the quality and meaning of the source evidence rather than with the technical convenience of creating connections.

Knowledge graphs offer a potentially valuable structure for organizing heterogeneous biomedical information as typed entities, relations, evidence objects, and provenance records. They can connect literature, compounds, biological targets, diseases, pathways, phenotypes, assays, and other knowledge domains while supporting semantic search, integration, and hypothesis generation. Biomedical knowledge-graph research nevertheless shows that graph quality depends on source selection, entity normalization, relation definition, ontology design, provenance, completeness, bias, and evaluation [3]. These dependencies become more demanding when traditional medicine knowledge is included. Traditional terms may not map exactly to biomedical concepts, preparation contexts may be incompletely described, multiple sources may conflict, and some knowledge may remain sensitive or governed by communities. A technically coherent graph can therefore remain scientifically weak, culturally reductive, or ethically inappropriate if its edges conceal uncertainty, its terminology erases local meaning, or its access model assumes that all published information is freely reusable.

The ethical challenges are not resolved merely by adding general principles such as transparency, fairness, privacy, or accountability to a graph-development protocol. Analyses of the global AI ethics landscape show broad agreement around several principles but substantial variation in interpretation, implementation, responsibility, and enforcement [4]. For traditional medicine knowledge graphs, this implementation gap is particularly consequential because the represented information may involve cultural identity, community-held interests, biological resources, restricted knowledge, commercial value, and unequal research relationships. A responsible architecture must therefore convert broad ethical intentions into specific practices for source screening, context preservation, sensitivity assessment, access control, evidence grading, provenance, expert review, stakeholder engagement, validation planning, and feedback. The objective of this article is to propose an original responsible knowledge framework that integrates these practices with ontology design, mechanistic interpretation, safety evidence, benefit-sharing logic, and explicit claim boundaries for traditional medicine knowledge graphs used in drug-discovery research.

Traditional medicine knowledge in drug discovery

Ethnopharmacological knowledge can guide natural-product discovery by identifying organisms, preparations, therapeutic purposes, administration practices, and observations considered relevant within particular medical and cultural settings. Its scientific value lies partly in its capacity to direct attention toward materials and hypotheses that might otherwise be overlooked. Ethnopharmacology has therefore contributed to bioprospecting and natural-product research by connecting traditional use reports with botanical, chemical, and pharmacological investigation [5]. However, a traditional-use report is not a simple compound–disease association. It may refer to a complex preparation, a locally defined illness category, a preventive or supportive practice, a particular route of administration, or a context that has no exact biomedical equivalent. Responsible knowledge representation must retain these distinctions. The corresponding graph record should preserve the reported use, source context, identity information, preparation details, and evidentiary status rather than translating the record directly into an efficacy relation or validated therapeutic indication.

Traditional medicine knowledge graphs also require careful differentiation among the material and evidentiary entities involved in drug discovery. Botanical identity

should be linked to accepted taxonomy, reported local nomenclature, plant part, specimen or source information where available, and the preparation to which a use report applies. Preparation context should distinguish whole materials, mixtures, processed products, extracts, fractions, and isolated constituents. Indication mapping should indicate whether a traditional term has an exact, approximate, broader, narrower, or unresolved relationship to a biomedical concept. Structured traditional-medicine resources demonstrate the technical feasibility of connecting medicines, ingredients, compounds, targets, diseases, experimental findings, and bibliographic sources [6]. Such structures can improve retrieval and comparison, but their value depends on maintaining the distinction between database assertions, literature-derived associations, experimental results, and computationally generated links. The presence of a connection in a structured resource does not determine its evidentiary strength or translational meaning.

The integration of phytochemical, bioactivity, safety, mechanistic, and literature evidence can extend traditional medicine knowledge graphs beyond descriptive cataloguing. Multimodal natural-product informatics may combine chemical structures, spectral or compositional data, assay results, target information, pathway annotations, omics evidence, textual evidence, and known safety signals within a common computational environment. Knowledge graphs can support this integration by representing both the biological entities and the evidence objects that justify their relationships. Recent analysis of artificial intelligence in natural-product science highlights the potential of multimodal data and knowledge graphs to support evidence integration and discovery-oriented reasoning [7]. Yet multimodality also increases the risk of evidentiary collapse. A measured constituent, a literature-reported constituent, a predicted constituent, and a compound inferred from a related species should not share the same relation status. Similarly, in vitro activity, target prediction, pathway enrichment, animal evidence, clinical observation, and traditional use must remain separately typed and qualified.

Responsible use of traditional medicine knowledge in computational discovery also requires attention to inclusivity, cultural respect, data governance, and unequal power. Artificial intelligence may increase the scale at which ethnopharmacological literature is extracted, translated, linked, and ranked, but scale does not remove the obligations associated with the origin and meaning of the information. Responsible AI scholarship in ethnopharmacology emphasizes the need to avoid extractive data practices, represent diverse knowledge traditions appropriately, address bias, and protect culturally sensitive information [8]. A knowledge graph should consequently support hypothesis generation without implying that all accessible knowledge is ethically unrestricted, commercially available, or suitable for automated reuse. Traditional use alone cannot establish mechanism, safety, efficacy, candidate validation, or commercialization readiness. These conclusions require separate evidence and governance processes, including material verification, experimental testing, toxicity assessment, expert interpretation, engagement with relevant stakeholders, and review of intended downstream use.

Cultural context and evidence grading

Cultural context is not an optional descriptive field added after graph construction; it is part of the meaning of the knowledge being represented. Ethnopharmacological field research standards emphasize the importance of documenting the research setting, knowledge source, participant or practitioner context, collection approach, terminology, biological material, preparation, use, and reporting process [9]. These elements determine what an assertion actually means and how far it may be generalized. A graph that retains only a plant name and a biomedical disease label may erase distinctions among local illness concepts, ritual or supportive uses, preventive practices, multi-ingredient preparations, and therapeutic claims. Cultural-context preservation therefore requires a source-linked record that maintains the original terminology and reported use while allowing qualified mappings to biomedical vocabularies. Where a mapping is approximate or contested, the graph should preserve the original concept and annotate uncertainty rather than enforce a false equivalence.

Local nomenclature, preparation details, use conditions, and source provenance are equally important for evidence grading. Recommended standards for ethnopharmacological field studies call for clear documentation of local names, taxonomic identification, participant selection, collection methods, preparation, administration, and use context [10]. These requirements should be translated into minimum metadata rules for graph ingestion. Evidence assessment should also recognize that ethnopharmacological research proceeds through multiple stages, including question formulation, field investigation, botanical identification, chemical analysis, pharmacological testing, toxicological evaluation, and interpretation [11]. Each stage produces a different kind of evidence. A recorded use may support the assertion that a practice was reported in a particular context; it does not establish that a material is pharmacologically active. A bioassay may support activity under defined experimental conditions; it does not establish clinical benefit. Evidence grades should therefore

describe what a source supports rather than assign a single universal score to all graph relations.

Published evidence must be evaluated by type, directness, material correspondence, methodological quality, and relevance to the specific graph claim. Methodological literature in ethnopharmacology distinguishes field evidence, botanical and chemical characterization, experimental pharmacology, toxicology, and clinical research as related but non-interchangeable domains [12]. For extracts and complex preparations, adequate chemical characterization is especially important because poorly defined materials prevent reproducible interpretation and weaken links between traditional preparations, measured constituents, and biological outcomes [13]. A responsible graph should therefore attach evidence objects to claims and retain details concerning study design, material identity, preparation, analytical method, assay system, controls, exposure, endpoint, limitations, and source citation where reported. Citation traceability should allow users to determine whether a graph relation is directly reported, curator-interpreted, text-mined, computationally predicted, or inferred from another relation. Conflicting evidence should be preserved as parallel, evidence-qualified assertions rather than silently resolved through majority counting or selective inclusion.

Community-held knowledge sensitivity and restricted knowledge boundaries require a governance process separate from conventional literature quality assessment. Publication status alone should not be treated as proof that knowledge is culturally unrestricted, ethically suitable for unrestricted mining, or available for commercial exploitation. Sensitive records may require exclusion, aggregation, reduced granularity, controlled access, or deferral pending appropriate review. Expert review should include relevant botanical, pharmacological, toxicological, informatics, ethnopharmacological, linguistic, cultural, and governance expertise according to the record and intended use. Best-practice standards can strengthen transparency and reproducibility, but they do not independently prove therapeutic validity or ethical acceptability [14]. For that reason, the graph should contain both a no therapeutic-claim boundary and a no ethics-compliance boundary. The first prevents traditional-use, bioactivity, or computational evidence from being presented as clinical efficacy. The second prevents provenance fields, access labels, citations, or review records from being presented as ethical certification, legal sufficiency, community approval, or authorization for commercialization.

Table 1 summarises cultural context and evidence grading requirements for responsible traditional medicine knowledge graphs in drug discovery.

Table 1. Cultural Context and Evidence Grading Requirements for Traditional Medicine Knowledge Graphs in Drug Discovery: Traditional Use Context, Local Nomenclature, Preparation Context, Indication Mapping, Source Provenance, Published Evidence, Community-Held Knowledge Sensitivity, Restricted Knowledge Boundaries, Citation Traceability, Uncertainty, and Claim Limits

Knowledge component	Core responsibility question	Required context or evidence	Evidence-grading function	Potential output	Risk of misrepresentation	Governance or validation need	Framework role
Traditional medicine knowledge	What kind of knowledge is represented, and within which knowledge system or research context?	Knowledge type, source context, terminology, reported purpose, provenance, and scope.	Distinguishes situated knowledge from generalized biomedical interpretation.	Contextualized traditional medicine knowledge record.	Treating diverse traditions as a single standardized body of data.	Domain and cultural-context review.	Primary knowledge layer.
Traditional use context	What use was reported, by whom, under which circumstances, and for what purpose?	Reported use, setting, preparation, administration, timing, population, and source.	Grades completeness and proximity to the original account.	Source-linked traditional-use assertion.	Converting a use report into an efficacy or treatment claim.	Ethnopharmacological review.	Use-evidence layer.
Local or traditional nomenclature	What does the reported term denote, and is its translation exact or approximate?	Original term, language, transliteration,	Assigns confidence to terminology	Preserved local term with qualified	Treating approximate translations as exact biomedical equivalents.	Linguistic, cultural, botanical, or clinical terminology	Terminology-preservation layer.

		translation, synonyms, geographic context, and mapping rationale.	and semantic mappings.	cross-reference.		review as appropriate.	
Preparation context	Which material or formulation does the evidence concern?	Biological material, part, ingredients where reported, processing, extraction, route, and preparation conditions.	Determines whether evidence applies to a traditional preparation, extract, fraction, or isolated constituent.	Preparation-specific evidence record.	Attributing the activity of an isolated compound directly to a complex preparation.	Chemical characterization and material correspondence assessment.	Material-definition layer.
Use context	Is the reported practice therapeutic, preventive, supportive, symptomatic, diagnostic, or otherwise context-specific?	Intended purpose, practitioner role, timing, setting, co-practices, and reported population.	Separates distinct categories of use and limits generalization.	Qualified use relation.	Collapsing multiple practices into a generic treatment edge.	Ethnopharmacological and cultural-context review.	Relation-qualification layer.
Indication mapping	Does a traditional illness or symptom term correspond to a biomedical concept?	Original wording, symptom description, local classification, biomedical candidate term, and mapping confidence.	Classifies mappings as exact, partial, broader, narrower, approximate, or unresolved.	Qualified indication crosswalk.	False equivalence between traditional and biomedical diagnoses.	Joint cultural-domain and biomedical review.	Semantic mapping layer.
Knowledge source provenance	Can the origin and transformation of every assertion be reconstructed?	Source publication or record, contributor or curator where appropriate, extraction method, transformation history, date, and version.	Distinguishes primary, secondary, curated, transformed, and inferred evidence.	Provenance chain linked to the assertion.	Presenting transformed or inferred information as directly reported knowledge.	Mandatory provenance audit.	Traceability backbone.
Published evidence	What does the cited study directly support?	Study design, material identity, methods, results, limitations,	Grades relevance, directness, methodological strength, and corresponde	Literature-supported evidence object.	Citation laundering or citing a source that does not support the represented claim.	Claim-to-source verification and updating.	Published-evidence layer.

		publication status, and persistent identifier.	nce to the graph claim.				
Community-held knowledge sensitivity	Could representation expose knowledge that remains governed, sensitive, restricted, or harmful if disclosed?	Sensitivity indicators, expressed expectations, disclosure status, potential harms, and permitted-use information where available.	Determines whether knowledge should be open, controlled, summarized, deferred, or excluded.	Sensitivity classification or non-ingestion decision.	Assuming that publication or accessibility eliminates community interests.	Appropriate stakeholder and governance review.	Sensitive-knowledge gate.
Restricted knowledge boundary	Should the information be represented, and at what level of detail?	Restriction status, sacred or confidential character, ecological or commercial sensitivity, harm assessment, and approved scope.	Applies exclusion, minimization, abstraction, or controlled access.	Withheld, coarse-grained, or access-controlled representation.	Revealing protected knowledge or enabling reconstruction through graph relations.	Access control, security review, and audit.	Hard governance boundary.
Ethnopharmacological record	Does the source meet minimum documentation requirements?	Field methods, identity evidence, local terminology, use, preparation, source context, and reporting quality.	Grades record completeness and reliability.	Quality-coded ethnopharmacological assertion.	Treating incomplete and well-documented records as equivalent.	Best-practice checklist and expert review.	Core evidence object.
Citation traceability	Can each claim be traced to the exact supporting source?	Persistent identifier, relevant passage, dataset record, extraction note, and curation history.	Confirms correspondence between claim and evidence.	Auditable claim–source connection.	Decorative citation without claim-level support.	Citation audit and correction process.	Claim-audit layer.
Uncertainty annotation	Which aspects are unknown, ambiguous, weakly supported,	Missing information, confidence, mapping uncertainty,	Preserves uncertainty instead of converting it	Explicit uncertainty status and rationale.	False precision and overconfident graph interpretation.	Calibration and expert review.	Uncertainty-management layer.

	disputed, or computationally inferred?	extraction uncertainty, and model uncertainty.	into binary truth.				
Conflicting evidence	Do sources disagree about identity, use, composition, activity, safety, or mechanism?	Competing claims, study designs, materials, dates, methods, and evidence grades.	Retains parallel claims and compares their evidentiary bases.	Conflict set with unresolved or reviewed status.	Silently selecting a preferred source or averaging away disagreement.	Expert adjudication or explicit unresolved-conflict label.	Conflict-management layer.
No therapeutic-claim boundary	Could a user mistake traditional use, bioactivity, or computational evidence for proof of efficacy?	Claim type, evidence class, validation status, intended audience, and downstream use.	Blocks unsupported efficacy, treatment, and clinical-utility predicates.	Hypothesis-only or evidence-limited statement.	Therapeutic overclaiming or clinical misinformation.	Pharmacological, toxicological, and clinical validation before stronger claims.	Mandatory scientific boundary.
No ethics-compliance boundary	Could governance metadata be misread as proof of ethical, legal, or community authorization?	Access status, stakeholder review, permissions where documented, unresolved issues, scope, and review date.	Separates recorded governance activity from certification or approval.	Governance-status record with explicit limitations.	Ethics washing, false legal assurance, or implied community approval.	Continuing human governance and context-specific review.	Mandatory governance boundary.

Mechanistic links and knowledge graph design

A responsible traditional medicine knowledge graph requires an explicit ontology and schema that define what kinds of entities, evidence objects, relations, and qualifiers the graph is permitted to contain. Reviews of traditional Chinese medicine knowledge graphs show that graph development commonly involves source acquisition, entity recognition, relation extraction, knowledge fusion, schema construction, graph completion, application, and quality evaluation [15]. For responsible drug-discovery use, this technical workflow should be expanded to distinguish traditional medicine systems, ethnopharmacological records, biological materials, preparations, compounds, assays, targets, pathways, safety findings, literature sources, knowledge holders, access conditions, and governance decisions. Ontology classes should not force culturally specific terms into biomedical categories when correspondence is incomplete. Community-developed ontologies demonstrate that traditional-drug properties and domain concepts can be represented through formal classes and relations while retaining domain terminology [16]. Such semantic preservation is essential because interoperability should enable qualified comparison without erasing differences in meaning.

Node and edge definitions determine the scientific interpretation of a knowledge graph. A node may represent a taxon, specimen, preparation, compound, assay, disease concept, pathway, publication, community-interest record, or evidence object, but these entities must not be treated as interchangeable. Likewise, an edge should specify whether a relation was directly reported, experimentally observed, manually curated, text-mined, predicted, disputed, or inferred. Representation-learning approaches have been used to construct and query traditional medicine knowledge graphs, including relations among prescriptions, herbs, symptoms, and related entities [17]. However, embedding proximity or link prediction cannot independently establish biological mechanism or therapeutic relevance. Evidence-specific relation modeling is particularly important for natural product–drug

interactions, where enzyme, transporter, exposure, and pharmacokinetic evidence may arise from different experimental systems and levels of certainty [18]. Predicted or indirect relations should therefore remain visibly separated from verified observations and should never overwrite the source assertions from which they were derived.

Mechanistic interpretation requires a layered evidentiary path rather than a visually persuasive sequence of connected nodes. A preparation may be linked to reported constituents, measured constituents, isolated compounds, or computationally predicted compounds, each of which has a different evidentiary status. Bioactivity relations should specify the tested material, assay system, concentration, exposure duration, controls, endpoint, and source. Target links should distinguish direct binding, functional modulation, genetic evidence, pathway inference, docking, text-mined association, and model prediction. Knowledge graphs have been applied to drug discovery tasks such as target identification, repurposing, candidate prioritization, and mechanism-oriented reasoning, but their outputs remain discovery hypotheses that require experimental validation [19]. Graph reliability also depends on the coverage, identifiers, licenses, provenance, relation semantics, curation practices, and biases of the contributing datasets [20]. Consequently, a mechanistic path should be interpreted as an evidence-organizing structure, not as proof that every intermediate relation is correct or causal.

Network pharmacology, omics integration, ADMET prediction, and toxicity evidence can enrich mechanistic interpretation when they are represented with appropriate context and uncertainty. Network-derived associations may reflect database popularity, incomplete interactomes, parameter choices, or circular use of the same evidence across input and validation stages. Omics results are conditional on the sample, platform, preprocessing, statistical model, biological contrast, and preparation tested. ADMET and toxicity evidence should distinguish experimental observations, clinical signals, curated reports, and computational predictions, while absence of a graph relation must not be interpreted as evidence of safety. Studies of knowledge-graph completion in traditional Chinese medicine illustrate the need to evaluate predicted links separately from asserted knowledge and to assess graph quality explicitly [21]. The responsible design proposed here therefore includes a causal boundary: graph connectivity, statistical association, enrichment, or prediction may generate a mechanistic hypothesis, but causality requires suitable perturbational, temporal, or intervention evidence. It also includes a drug-discovery hypothesis boundary that prevents ranked compounds, targets, pathways, or preparations from being described as validated candidates before analytical, experimental, toxicological, and translational review.

Table 2 maps mechanistic link and knowledge graph design requirements for responsible drug discovery interpretation.

Table 2. Mechanistic Links and Knowledge Graph Design Requirements for Traditional Medicine Drug Discovery: Ontology Design, Graph Schema, Node and Edge Definitions, Semantic Interoperability, Botanical Identity, Phytochemical Evidence, Bioactivity Evidence, Target Evidence, Pathway Evidence, Safety Evidence, Mechanistic Hypotheses, and Causal Boundaries

Graph design component	Core design question	Required input	Knowledge graph function	Potential output	Risk of overinterpretation	Validation requirement	Decision boundary
Ontology design	Which concepts must remain distinct, and which mappings are justified?	Traditional medicine terminology, ethnopharmacological concepts, biomedical ontologies, evidence classes, governance concepts, and competency questions.	Defines classes, properties, constraints, and qualified mappings.	Versioned responsible-knowledge ontology.	Presenting ontology categories as universal or culturally neutral.	Logical testing, domain review, cultural-context review, and version governance.	Ontological representation does not establish scientific or cultural truth.
Knowledge graph schema	Which entity, evidence, and governance types may	Research purpose, source domains, access rules, claim boundaries, and	Specifies the permitted structure and	Modular graph schema covering knowledge, evidence,	Omitting difficult contextual information to simplify the graph.	Use-case testing and multidisciplinary review.	No ingestion outside the documented schema and purpose.

	be represented ?	validation requirements.	mandatory metadata.	mechanism, safety, and governance.			
Node definition	What real-world, conceptual, documentary, or governance object does each node denote?	Identity criteria, identifiers, labels, provenance, temporal scope, and version.	Creates typed entities for materials, concepts, sources, evidence, and governance records.	Provenance-linked typed node.	Merging taxa, samples, preparations, compounds, and source assertions.	Identity-resolution audit.	Unresolved entities remain separate.
Edge definition	What exact assertion does the relation encode?	Subject, predicate, object, source, evidence type, qualifiers, direction, and uncertainty.	Encodes an auditable claim rather than an unqualified connection.	Evidence-qualified edge.	A direct-looking edge may be interpreted as causal or validated.	Predicate review and source verification.	No high-impact edge without evidence linkage.
Relation type	Was the relation reported, observed, curated, predicted, inferred, or disputed?	Evidence origin, extraction method, transformation history, and review status.	Separates epistemic status through typed relations.	Asserted, observed, predicted, inferred, or contested relation.	Mixing model-generated links with experimentally reported findings.	Automated consistency checks and manual sampling.	Inferred relations cannot overwrite source assertions.
Semantic interoperability	Can different vocabularies be connected without erasing local meaning?	Standard identifiers, original terminology, synonyms, crosswalks, and mapping confidence.	Supports exchange and integration through qualified mappings.	Interoperable concepts with preserved original labels.	Flattening culturally specific categories into biomedical terminology.	Reversible mapping review and confidence assessment.	Uncertain mappings remain explicitly uncertain.
Botanical identity	Which taxon, specimen, plant part, and material are represented ?	Scientific name, authority, voucher or source data where reported, plant part, collection information, and taxonomic version.	Anchors preparations and evidence to a defined biological identity.	Taxonomically qualified material node.	Taxonomic synonymy or misidentification contaminating downstream relations.	Botanical and taxonomic validation.	No species-level inference without adequate identity support.
Phytochemical evidence	Was a constituent measured, isolated, annotated, reported, or predicted?	Analytical method, sample identity, compound identifier, abundance or concentration where reported, and source.	Connects materials to constituents with evidence qualifiers.	Measured, isolated, annotated, reported, or predicted constituent relation.	Treating predicted or literature-associated constituents as measured in the preparation.	Analytical-method and material-correspondence review.	Constituent status must remain evidence-specific.
Bioactivity evidence	Under which conditions	Tested material, assay model, dose or	Represents context-dependent	Assay-qualified	Translating in vitro or model-	Assay-quality and biological-	Bioactivity alone cannot

	was biological activity observed?	concentration, duration, controls, endpoint, statistics, and source.	biological observations.	bioactivity relation.	system activity into therapeutic efficacy.	relevance review.	establish clinical utility.
Target evidence	What kind of evidence supports the material–target or compound–target relation?	Binding assay, functional assay, genetics, expression, prediction, docking, literature extraction, or pathway inference.	Encodes target relations by method and confidence.	Evidence-specific target relation.	Treating docking, text mining, or prediction as confirmed target engagement.	Orthogonal biochemical or cellular validation.	Predicted target links remain hypothesis-level.
Pathway evidence	How was the pathway relation established?	Curated pathway source, enrichment analysis, expression data, perturbation study, or functional experiment.	Connects molecular evidence to biological processes.	Method-qualified pathway relation.	Reading enrichment as causal activation or inhibition.	Independent functional validation.	No causal pathway claim from enrichment alone.
Network pharmacology evidence	Are the network inputs, algorithms, and assumptions transparent?	Compound set, target sources, network database, parameters, null models, and provenance.	Generates network modules, centrality measures, and ranked mechanistic hypotheses.	Network-supported hypothesis.	Database popularity bias, circularity, and inflated multi-target claims.	Sensitivity analysis and experimental follow-up.	Network position is not mechanistic validation.
Omics evidence where supported	Which sample and analytical workflow generated the molecular signal?	Omics type, biological model or cohort, preparation, sampling context, platform, preprocessing, and statistical contrast.	Connects context-specific molecular measurements to mechanisms.	Sample-bound omics evidence object.	Generalizing a model-specific signature to other materials or populations.	Replication and biological validation.	Omics relations remain conditional on the measured context.
ADMET evidence	Is the disposition or toxicity property observed or predicted?	Endpoint, species or system, dose, exposure, method, model applicability domain, and uncertainty.	Represents absorption, distribution, metabolism, excretion, and toxicity evidence.	Experimental, curated, clinical, or predicted ADMET relation.	Presenting computational predictions as established human pharmacokinetics.	Fit-for-purpose experimental confirmation.	Predicted ADMET cannot establish safety or suitability.
Toxicity evidence	What adverse outcome occurred, under	Material, dose, duration, species or population, route, comparator,	Makes adverse effects and safety	Exposure-qualified toxicity relation.	Ignoring formulation, dose, interaction, or	Toxicological and clinical review as appropriate.	Absence of reported toxicity does not establish safety.

	which exposure conditions?	outcome, and source.	signals visible.		population differences.		
Mechanistic link	What evidentiary chain connects preparation, compound, target, pathway, phenotype, and outcome?	Evidence-qualified intermediate relations and provenance.	Constructs transparent mechanistic paths.	Mechanistic evidence path.	Assuming that graph connectivity proves mechanism.	Validation of critical links, directionality, and alternatives.	A connected path is not mechanism confirmation.
Mechanistic hypothesis	Which graph-supported explanation is testable and falsifiable?	Evidence paths, uncertainty, competing explanations, biological plausibility, and safety context.	Generates explicit research hypotheses.	Ranked hypothesis with rationale and limitations.	Ranking may conceal weak, circular, or incomplete evidence.	Experimental design and competing-hypothesis review.	Hypothesis status persists until tested.
Causal boundary	Which relations are associational, intervention-supported, or causally unresolved?	Study design, temporal ordering, perturbation evidence, confounding assessment, and replication.	Labels causal status and blocks unsupported causal predicates.	Associational, intervention-supported, or unresolved causal relation.	Association-to-causation inference.	Perturbational, longitudinal, or causal-inference evidence.	No causal predicate without adequate design.
Drug discovery hypothesis boundary	What may be prioritized without being called a validated candidate?	Evidence grade, uncertainty, safety flags, novelty, feasibility, governance status, and intended use.	Produces research-priority records with explicit limitations.	Prioritized compound, preparation, target, pathway, or interaction hypothesis.	Candidate ranking interpreted as efficacy or development readiness.	Analytical, experimental, safety, translational, and governance review.	No validated-candidate or commercialization-readiness label.

Benefit-sharing logic

Benefit-sharing logic should be treated as a core design requirement rather than as a statement added after knowledge extraction. Traditional medicine and natural-product research may involve knowledge, materials, institutions, and communities located within unequal scientific and commercial relationships. Comparative case analyses of access and benefit-sharing in ethnopharmacology demonstrate that workable arrangements depend on the research context, source-country environment, institutional relationships, intended use, and practical opportunities for collaboration [22]. A knowledge graph cannot determine or satisfy these obligations by itself. It can, however, make relevant provenance, contribution, use restrictions, collaboration records, and decision points visible. The graph should therefore contain governance objects that identify the intended research purpose, source and material provenance, permitted-use boundaries, review status, access tier, stakeholder interests, and triggers for renewed review when the project purpose or downstream use changes.

Inclusive provenance is necessary because conventional biomedical metadata often record publications, databases, and technical transformations while omitting the people, communities, and relationships from which knowledge or biological materials originated. Scholarship on benefit sharing argues that provenance metadata should recognize contributors, communities, origins, and associated interests rather than treating provenance as a purely technical lineage [23]. Such recognition should not be reduced to a decorative attribution field. Where appropriate and authorized, the system should preserve preferred forms of attribution, contributor roles, source-country partnerships,

and governance contacts while avoiding claims that the graph speaks for a community. Local Contexts approaches further illustrate how notices and labels may communicate community expectations and provenance-related responsibilities to researchers [24]. These mechanisms may support awareness and responsible handling, but they do not replace dialogue, consent, authorization, negotiated arrangements, or context-specific governance.

Ethical sourcing and biodiversity responsibility are also relevant to graph design because medicinal-plant discovery depends on accurate biological identity, geographic and ecological context, collection practices, and sustainable access to biological resources. Modern biodiversity science can improve taxonomy, phylogenetic understanding, distribution knowledge, and conservation awareness in medicinal-plant research [25]. The responsible graph should therefore preserve sourcing information at a level appropriate to the research purpose while protecting sensitive ecological locations and other information that could increase harm. Traditional knowledge respect similarly requires recognition that medicinal-plant knowledge may remain embedded in Indigenous or local knowledge systems, social relationships, and collective interests. Reviews of Indigenous knowledge in medicinal-plant research emphasize both its potential contribution to therapeutic investigation and the challenges of equitable collaboration, recognition, protection, and translation [26]. Source-country participation and community-interest recognition should consequently be evaluated as substantive relationships rather than inferred from a publication address, geographic label, or citation.

Operationally, the proposed framework uses data access tiering, stakeholder review, audit trails, and commercialization boundaries to prevent research access from being silently converted into unrestricted downstream use. Records may be designated open, registered, controlled, restricted, summarized, deferred, or excluded according to sensitivity, permitted purpose, and governance review. Audit logs should record access, edits, exports, inferred relations, and material changes in intended use. A commercialization boundary should trigger renewed review when research hypotheses, data products, intellectual property strategies, licensing discussions, or development plans extend beyond the purpose under which the knowledge was initially represented. Feedback to knowledge holders may be appropriate where relevant relationships and agreed communication channels exist, but feedback procedures must not expose confidential identities or imply universal community representation. The graph should contain both a no legal-compliance claim boundary and a no ethical-certification claim boundary: technical controls, provenance records, attribution, notices, stakeholder review, or benefit-sharing fields do not establish legal sufficiency, ethical certification, community approval, or permission for commercialization.

Proposed responsible knowledge framework

The proposed framework begins with research question definition and knowledge source screening. Before ingestion, the research team should define the scientific objective, intended users, permitted outputs, anticipated downstream uses, relevant cultural and governance risks, and the types of evidence needed to address the question. Eligible sources should then be screened for provenance, methodological relevance, access status, material identity, sensitivity, and fitness for the proposed use. Biomedical knowledge-graph resources demonstrate that heterogeneous entities and relation types can be integrated through a documented multimodal schema [27]. In the present framework, however, technical integrability is not sufficient for inclusion. A source may be scientifically relevant yet unsuitable for unrestricted representation because its context is incomplete, its knowledge status is uncertain, its permitted use is unclear, or its content is sensitive. Source screening therefore produces an explicit inclusion, controlled-inclusion, deferral, minimization, or exclusion decision.

The second stage combines cultural context preservation, sensitive knowledge triage, evidence grading, and ontology design. Original terminology, local or traditional nomenclature, use context, preparation context, source provenance, and uncertainty should be retained before mappings to standardized vocabularies are created. Sensitive knowledge triage determines whether a record should be open, controlled, restricted, summarized, or excluded. Evidence grading then distinguishes traditional-use documentation, literature interpretation, botanical identification, chemical characterization, bioactivity, target evidence, pathway evidence, omics evidence, safety findings, computational prediction, and clinical evidence. The ontology and graph schema should represent these evidence classes explicitly while preserving source-specific limitations. Semantically standardized biomedical graph systems show the value of canonical identifiers, ontology-aligned predicates, reproducible integration pipelines, and documented source attribution [28]. The proposed framework extends these practices by requiring governance metadata, sensitivity status, access tier, claim boundaries, and culturally qualified mappings as first-class graph components.

The third stage integrates mechanistic and safety evidence without collapsing observational, experimental, and computational findings. Preparation–compound, compound–target, target–pathway, pathway–phenotype,

and material–safety relations should be represented as evidence-qualified paths. Omics evidence may add biological context when the sample, platform, processing workflow, and experimental contrast remain visible. Knowledge graphs designed for clinical proteomics illustrate how experimentally generated molecular measurements can be connected to curated biomedical knowledge to support interpretation [29]. The responsible framework requires the same separation between sample-derived findings and established knowledge. It also requires lifecycle provenance for every transformation, including extraction, normalization, entity resolution, mapping, inference, human review, and version change. Reviews of biomedical provenance show that useful provenance may include the origin of data, responsible agents, workflow steps, transformations, timestamps, and version histories [30]. Provenance is therefore treated as an active accountability mechanism rather than as a static citation field.

The fourth stage combines benefit-sharing logic mapping, access governance review, expert and stakeholder review, drug-discovery hypothesis generation, and validation planning. Graph algorithms may identify connected evidence paths or prioritize compounds, preparations, targets, pathways, and interactions, but these outputs should remain explicitly labeled as hypotheses. Systematic biomedical knowledge integration has shown that graph-based features can prioritize drug-repurposing candidates, while the resulting rankings still require external and experimental testing [31]. Before a hypothesis progresses, qualified reviewers should examine botanical identity, chemical correspondence, assay relevance, mechanistic plausibility, safety gaps, cultural context, access conditions, and benefit-sharing considerations. Responsible machine-learning guidance in health care similarly emphasizes multidisciplinary design, validation, monitoring, and harm-aware oversight across the system lifecycle [32]. In the proposed framework, validation planning defines the analytical, experimental, toxicological, or translational evidence needed to test each hypothesis and specifies in advance what outcomes would support, weaken, refute, or leave the hypothesis unresolved.

The final stage is feedback updating under continuing governance. New publications, experimental results, corrections, retractions, taxonomic changes, safety signals, stakeholder decisions, access changes, and failed hypotheses should update the graph without erasing the historical record. Each revision should be versioned and linked to the actor, reason, evidence, and affected claims. Ethical principles should be translated into concrete responsibilities, documentation tools, review gates, audit procedures, and escalation pathways rather than being presented as abstract assurances [33]. The framework therefore retains two non-negotiable boundaries throughout its lifecycle. The no therapeutic-claim boundary prevents reported use, graph connectivity, bioactivity, target prediction, pathway enrichment, or candidate ranking from being described as therapeutic efficacy or clinical utility. The no ethics-compliance boundary prevents governance metadata, access controls, attribution, stakeholder consultation, or audit records from being described as ethical certification, legal compliance, community approval, or commercialization authorization.

Table 3 presents the proposed responsible knowledge framework for traditional medicine knowledge graphs in drug discovery.

Table 3. Proposed Responsible Knowledge Framework for Traditional Medicine Knowledge Graphs in Drug Discovery: Source Screening, Cultural Context Preservation, Sensitive Knowledge Triage, Evidence Grading, Ontology Design, Mechanistic Link Mapping, Safety Evidence, Benefit-Sharing Logic, Access Governance, Expert and Stakeholder Review, Validation Planning, Feedback, and Claim Boundaries

Framework stage	Core purpose	Required input	Responsible knowledge function	Potential output	Validation need	Failure risk	Feedback mechanism	Decision boundary
Research question definition	Establish a legitimate, bounded scientific purpose.	Research objective, intended users, expected outputs, downstream-use assumptions, cultural risks,	Prevents purposeless aggregation and scope drift.	Documented research-purpose statement.	Scientific and governance review proportionate to risk.	Data accumulation without a defensible or authorized purpose.	Revise scope when assumptions or intended uses change.	No source ingestion before purpose and boundaries are documented.

		and governance context.						
Knowledge source screening	Determine whether each source is suitable for the intended graph use.	Source type, provenance, publication status, methodological quality, access status, sensitivity, material identity, and relevance.	Includes eligible sources and rejects or restricts unsuitable ones.	Include, controlled-include, minimize, defer, or exclude decision.	Source-level verification.	Ingestion of unreliable, restricted, or contextually inappropriate material.	Re-screen after corrections, retractions, or governance changes.	Excluded or deferred sources cannot enter downstream graph layers.
Cultural context preservation	Retain the original meaning and setting of represented knowledge.	Original terminology, use context, preparation, medical-system concepts, source narrative, and mapping rationale.	Links standardized concepts to original contextual records.	Context-preserving knowledge object.	Cultural, linguistic, and ethnopharmacological review.	Decontextualization and biomedical reductionism.	Correct mappings and add missing context.	No exact-equivalence claim without documented justification.
Sensitive knowledge triage	Identify knowledge requiring restriction, abstraction, deferral, or exclusion.	Sensitivity indicators, disclosure status, harm assessment, knowledge-holder expectations, and permitted-use information where available.	Classifies information before graph ingestion.	Open, controlled, restricted, summarized, deferred, or excluded status.	Authorized governance and security review.	Disclosure or reconstruction of protected knowledge.	Reclassify after stakeholder or risk review.	Restricted content cannot enter open graph views.
Evidence grading	Distinguish evidence types and the claims each can support.	Study design, source proximity, identity evidence, methods, replication, relevance, consistency, and uncertainty.	Assigns domain-specific evidence classes and limitations.	Evidence-grade record with rationale.	Domain-expert assessment and audit.	False equivalence among traditional use, assays, predictions, and clinical findings.	Update grades as new evidence emerges.	No universal score may erase evidence-type differences.
Ontology and graph schema design	Define concepts, relations, constraints, qualifiers, and	Competency questions, domain vocabulary, standard ontologies, local terms, evidence	Creates an extensible semantic structure.	Versioned ontology and graph schema.	Logical, technical, cultural, and domain review.	Schema bias, category collapse, or absent governance concepts.	Versioned change requests and mapping revision.	Material schema changes require documented review.

	governance fields.	classes, and access requirements.						
Node and relation construction	Create auditable entities and claims.	Identity criteria, identifiers, source assertions, evidence objects, relation definitions, and uncertainty.	Separates entities from claims and claims from evidence.	Typed nodes and evidence-qualified relations.	Entity-resolution and claim-to-source audit.	Merged entities, ambiguous predicates, and unsupported edges.	Correct entities or relations while retaining change history.	No unsupported high-impact relation.
Mechanistic link mapping	Connect preparation, compound, target, pathway, phenotype, and outcome evidence.	Evidence-qualified intermediate relations and competing explanations.	Builds transparent mechanistic paths.	Mechanistic hypothesis with evidence chain.	Validation of critical links and directionality.	Graph connectivity interpreted as proven mechanism.	Update or retract paths after experimental findings.	No mechanism-confirmed status from connectivity alone.
Safety evidence integration	Make toxicity, interactions, contraindications, exposure, and uncertainty visible.	Toxicology, ADMET, interaction, pharmacovigilance, and exposure evidence where available.	Builds a safety and evidence-gap subgraph.	Safety signal, interaction hypothesis, or documented evidence gap.	Toxicological, pharmacological, and clinical review where appropriate.	Absence of information interpreted as absence of risk.	Add new safety findings and retire superseded claims.	No safety or suitability claim from incomplete graph evidence.
Benefit-sharing logic mapping	Connect knowledge and material provenance to prospective recognition and benefit pathways.	Stakeholder interests, contribution records, intended use, collaboration arrangements, and use-transition triggers.	Makes benefit-sharing considerations visible at decision points.	Benefit-sharing consideration record.	Stakeholder and institutional governance outside the graph.	Symbolic metadata substituted for substantive arrangements.	Reassess when the purpose, value, or downstream use changes.	Graph metadata cannot constitute a benefit-sharing agreement.
Access governance review	Match access privileges to purpose, sensitivity, and governance status.	User role, purpose, access tier, permitted operations, sensitivity, and review date.	Authorizes, limits, logs, or denies graph operations.	Time- and purpose-bounded access decision.	Human governance, security testing, and periodic review.	Excessive access, secondary use, or inferential disclosure.	Access recertification and incident review.	Denied or expired access blocks querying and export.

ISSN (Online) 2249-6084 (Print) 2250-1029

Expert and stakeholder review	Test scientific, cultural, safety, and governance fit.	Graph records, evidence grades, conflicts, uncertainty, intended output, and access conditions.	Adds human judgment, records dissent, and imposes conditions.	Approved-for-research, revision-required, restricted, or rejected status.	Appropriate expertise, role clarity, and conflict management.	Superficial consultation or decisions made outside reviewer competence.	Structured comments, appeal, correction, and re-review.	Review scope must be explicit; no blanket endorsement.
Drug discovery hypothesis generation	Produce testable and bounded research hypotheses.	Quality-controlled graph, evidence paths, novelty criteria, uncertainty, safety flags, and governance status.	Ranks hypotheses with transparent rationales and limitations.	Prioritized compound, preparation, target, pathway, or interaction hypothesis.	Computational robustness and expert plausibility review.	Ranking interpreted as efficacy, validation, or readiness.	Experimental findings update hypothesis status.	Hypothesis status only.
Validation planning	Define the evidence needed to test each hypothesis.	Material definition, hypothesis, evidence gaps, assay relevance, safety concerns, and decision criteria.	Links graph outputs to staged analytical and experimental plans.	Validation roadmap.	Independent methodological review.	Convenient or weak assays selected without construct validity.	Supporting, refuting, or inconclusive results feed back into the graph.	No progression without predefined criteria.
Feedback updating	Correct errors and incorporate new scientific or governance information.	New evidence, experimental outcomes, corrections, retractions, stakeholder decisions, and access changes.	Versions nodes, relations, grades, and governance status.	Updated graph release and change log.	Regression testing and provenance audit.	Silent overwriting and loss of accountability.	Formal correction, appeal, and update channels.	Material changes trigger re-review.
No therapeutic-claim boundary	Prevent unsupported efficacy, treatment, or clinical-use conclusions.	Claim type, evidence hierarchy, validation status, intended audience, and downstream context.	Restricts predicates, labels, exports, and interface language.	Research-only statement with evidence limitations.	Pharmacological, toxicological, and clinical review before stronger claims.	Therapeutic overclaiming and clinical misinformation.	Boundary changes only after appropriate evidence review.	No therapeutic recommendation or efficacy conclusion.
No ethics-compliance boundary	Prevent technical governance	Review records, access conditions,	Uses descriptive governan	Governance-status record	Continuing context-specific	Ethics washing, implied community	Reassess after changes in purpose,	No ethical, legal, community-approval, or

	features from being represented as ethical authorization.	stakeholder input, unresolved issues, scope, and date.	ce statuses rather than certification labels.	with explicit limitations.	human governance.	endorsement, or false legal assurance.	access, stakeholders, or use.	commercialization certification.

Figure 1 illustrates the proposed responsible knowledge framework for traditional medicine knowledge graphs in drug discovery, connecting cultural context, evidence grading, mechanistic links, benefit-sharing logic, governance, validation planning, and feedback updating.

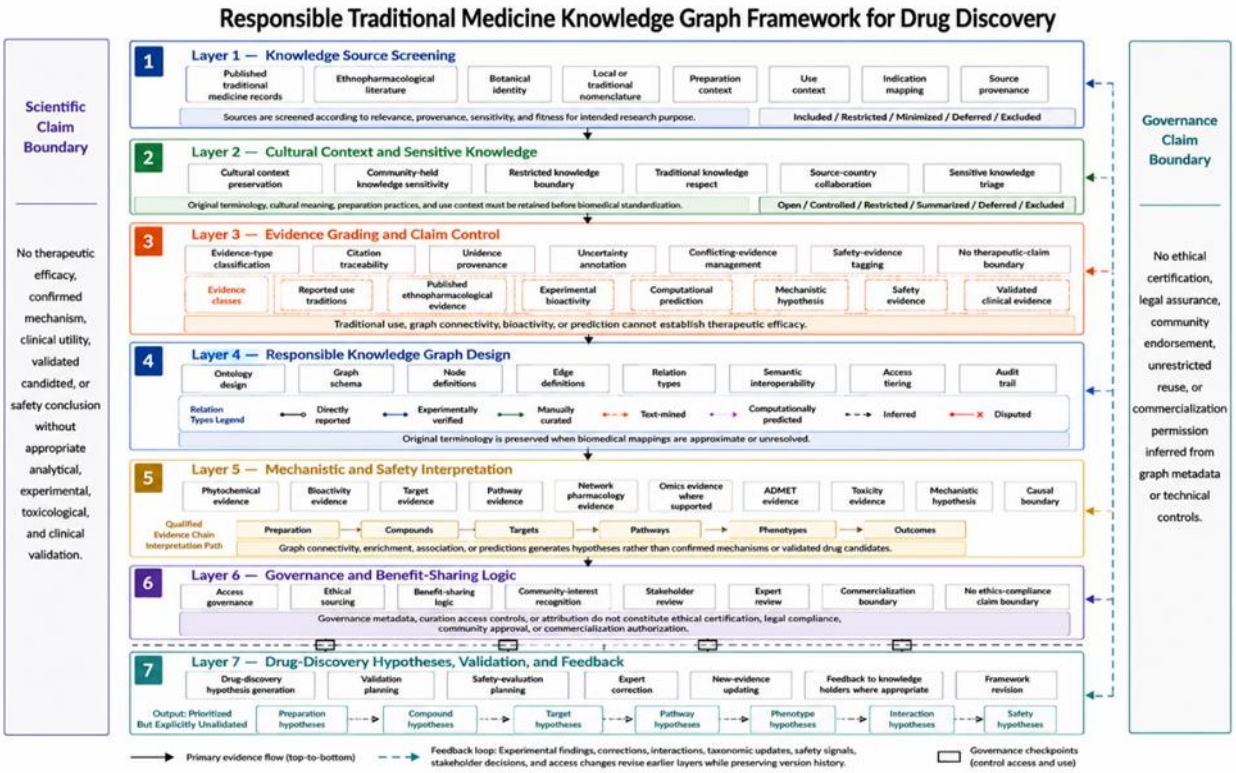

Figure 1. Responsible Traditional Medicine Knowledge Graph Framework for Drug Discovery

Alt-text description

A conceptual responsible knowledge framework showing traditional medicine knowledge flowing through source screening, cultural context preservation, sensitive knowledge triage, evidence grading, graph schema design, mechanistic link mapping, benefit-sharing logic, access governance, expert review, validation planning, feedback loops, and claim boundaries.

CONCLUSION

This article proposes a responsible knowledge framework for traditional medicine knowledge graphs in drug discovery that integrates source screening, cultural context preservation, sensitive knowledge triage, evidence grading, citation traceability, ontology design, mechanistic interpretation, safety evidence, benefit-sharing logic, access governance, expert and stakeholder review, validation planning, and feedback updating. Its central premise is that traditional medicine knowledge should not be reduced to decontextualized data points or converted directly into therapeutic, mechanistic, ethical, legal, or commercial conclusions. Knowledge graphs may support respectful and evidence-aware hypothesis generation when original terminology, preparation context, source provenance, uncertainty, conflicting evidence, sensitivity, and governance conditions remain visible. Their outputs must nevertheless remain bounded by the quality and meaning of the underlying evidence and by continuing human responsibilities. The proposed framework therefore treats computationally generated associations as research hypotheses, not validated drug candidates; governance metadata as accountability information, not ethical

certification; and attribution as recognition, not a substitute for benefit sharing. Responsible traditional medicine knowledge graphs should strengthen traceability, interpretation, and collaboration while resisting extractive reuse, therapeutic overclaiming, ethical shortcutting, exposure of restricted knowledge, and progression toward commercialization without appropriate validation and governance.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

[1] Newman DJ, Cragg GM. Natural products as sources of new drugs over the nearly four decades from 01/1981 to 09/2019. J Nat Prod. 2020;83(3):770-803. doi:10.1021/acs.jnatprod.9b01285

[2] Heinrich M, Appendino G, Efferth T, Fürst R, Izzo AA, Kayser O, et al. Best practice in research—overcoming common challenges in phytopharmacological research. J Ethnopharmacol. 2020;246:112230. doi:10.1016/j.jep.2019.112230

[3] 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

[4] Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. Nat Mach Intell. 2019;1(9):389-99. doi:10.1038/s42256-019-0088-2

[5] Buenz EJ, Verpoorte R, Bauer BA. The ethnopharmacologic contribution to bioprospecting natural products. Annu Rev Pharmacol Toxicol. 2018;58:509-30. doi:10.1146/annurev-pharmtox-010617-052703

[6] Fang SS, Dong L, Liu L, Guo JC, Zhao LH, Zhang JY, et al. HERB: A high-throughput experiment- and reference-guided database of traditional Chinese medicine. Nucleic Acids Res. 2021;49(D1):D1197-D1206. doi:10.1093/nar/gkaa1063

[7] Meijer D, Beniddir MA, Coley CW, Mejri YM, Öztürk M, van der Hooft JJJ, et al. Empowering natural product science with AI: Leveraging multimodal data and knowledge graphs. Nat Prod Rep. 2025;42(4):654-62. doi:10.1039/D4NP00008K

[8] Bhadra S, Pundir C, Mukherjee M, Kar A, Banerjee S, Charoensup R, et al. Reimagining ethnopharmacology with generative AI: Towards inclusive, ethical, and data-driven traditional medicine. Pharmacol Res. 2025;221:108002. doi:10.1016/j.phrs.2025.108002

[9] Heinrich M, Lardos A, Leonti M, Weckerle CS, Willcox M, Applequist W, et al. Best practice in research: Consensus statement on ethnopharmacological field studies—ConSEFS. J Ethnopharmacol. 2018;211:329-39. doi:10.1016/j.jep.2017.08.015

[10] Weckerle CS, de Boer HJ, Puri RK, van Andel T, Bussmann RW, Leonti M. Recommended standards for conducting and reporting ethnopharmacological field studies. J Ethnopharmacol. 2018;210:125-32. doi:10.1016/j.jep.2017.08.018

[11] Schultz F, Garbe LA. How to approach a study in ethnopharmacology? Providing an example of the different research stages for newcomers to the field today. Pharmacol Res Perspect. 2023;11(4):e01109. doi:10.1002/prp2.1109

[12] Jalil B, Schultz F, Heinrich M. Where to begin? The best publications for newcomers to ethnopharmacology. Front Pharmacol. 2023;14:1141502. doi:10.3389/fphar.2023.1141502

[13] Heinrich M, Jalil B, Abdel-Tawab M, Echeverria J, Kulić Ž, McGaw LJ, et al. Best practice in the chemical characterisation of extracts used in pharmacological and toxicological research—the ConPhyMP guidelines. Front Pharmacol. 2022;13:953205. doi:10.3389/fphar.2022.953205

[14] Heinrich M, Jalil B. From the CONSORT to the ConPhyMP statement and beyond—how to ascertain best practice. Front Pharmacol. 2023;14:1338710. doi:10.3389/fphar.2023.1338710

[15] Qu X, Tian Z, Cui J, Li R, Li D, Zhang X. A review of knowledge graph in traditional Chinese medicine: Analysis, construction, application and prospects. Comput Mater Contin. 2024;81(3):3583-616. doi:10.32604/cmc.2024.055671

[16] Zhu Y, Liu L, Gao B, Liu J, Qiao X, Lian C, et al. TCDO: A community-based ontology for integrative representation and analysis of traditional Chinese drugs and their properties. Evid Based Complement Alternat Med. 2021;2021:6637810. doi:10.1155/2021/6637810

[17] Weng H, Chen J, Ou A, Lao Y. Leveraging representation learning for the construction and application of a knowledge graph for traditional Chinese medicine: Framework development study. JMIR Med Inform. 2022;10(9):e38414. doi:10.2196/38414

[18] Taneja SB, Callahan TJ, Paine MF, Kane-Gill SL, Kilicoglu H, Joachimiak MP, et al. Developing a knowledge graph for pharmacokinetic natural

product–drug interactions. J Biomed Inform. 2023;140:104341. doi:10.1016/j.jbi.2023.104341

[19] 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

[20] 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):bbac404. doi:10.1093/bib/bbac404

[21] Liu C, Li Z, Li JM, Qu YQ, Chang Y, Han Q, et al. Research on traditional Chinese medicine: Domain knowledge graph completion and quality evaluation. JMIR Med Inform. 2024;12:e55090. doi:10.2196/55090

[22] Heinrich M, Scotti F, Andrade-Cetto A, Berger-Gonzalez M, Echeverría J, Friso F, et al. Access and benefit sharing under the Nagoya Protocol—quo vadis? Six Latin American case studies assessing opportunities and risk. Front Pharmacol. 2020;11:765. doi:10.3389/fphar.2020.00765

[23] Golan J, Riddle K, Hudson M, Anderson J, Kusabs N, Coltman T. Benefit sharing: Why inclusive provenance metadata matter. Front Genet. 2022;13:1014044. doi:10.3389/fgene.2022.1014044

[24] Liggins L, Hudson M, Anderson J. Creating space for Indigenous perspectives on access and benefit-sharing: Encouraging researcher use of the Local Contexts Notices. Mol Ecol. 2021;30(11):2477-82. doi:10.1111/mec.15918

[25] Davis CC, Choisy P. Medicinal plants meet modern biodiversity science. Curr Biol. 2024;34(4):R158-R173. doi:10.1016/j.cub.2023.12.038

[26] Alum EU. The role of indigenous knowledge in advancing the therapeutic use of medicinal plants: Challenges and opportunities. Plant Signal Behav. 2024;19(1):2439255. doi:10.1080/15592324.2024.2439255

[27] 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

[28] 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

[29] 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

[30] Gierend K, Krüger F, Genehr S, Hartmann F, Siegel F, Waltemath D, et al. Provenance information for biomedical data and workflows: Scoping review. J Med Internet Res. 2024;26:e51297. doi:10.2196/51297

[31] 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:e26726. doi:10.7554/eLife.26726

[32] Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6

[33] Morley J, Floridi L, Kinsey L, Elhalal A. From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. Sci Eng Ethics. 2020;26(4):2141-68. doi:10.1007/s11948-019-00165-5

ISSN (Online) 2249-6084 (Print) 2250-1029