From Ethnopharmacological Knowledge to Drug Candidates: A Decision-Intelligence Model for AI-Guided Natural Product Discovery

Mikko Lahtinen[1]*, Elina Salo[1], Juhani Virtanen[2]

[1]Department of Artificial Intelligence for Bioactive Compound Screening, Faculty of Pharmacy, University of Helsinki, Helsinki, Finland.

[2]Department of Pharmacoinformatics and Herbal Drug Development, Faculty of Pharmacy, University of Eastern Finland, Kuopio, Finland.

ABSTRACT

Natural products remain important sources of pharmacologically relevant chemical diversity, yet the translation of ethnopharmacological knowledge into drug-discovery candidates is constrained by heterogeneous provenance, taxonomic ambiguity, preparation variability, uncertain disease mapping, incomplete chemical resolution, and uneven experimental support. This article proposes an original decision-intelligence model for converting ethnopharmacological knowledge into transparent, revisable, and uncertainty-sensitive candidate-selection decisions supported by artificial intelligence and computational pharmacology. The model treats reported use, documented practice, historical persistence, cross-source recurrence, practitioner knowledge, and preparation context as structured discovery inputs rather than independent evidence of efficacy, safety, or mechanism. AI-guided functions—including knowledge extraction, entity resolution, molecular representation, similarity analysis, candidate identification, target inference, disease-context matching, and multimodal evidence integration—are positioned as hypothesis-generating and decision-support functions rather than validation mechanisms. Heterogeneous evidence is translated into explicit evidence states, decision variables, uncertainty classes, and conditional rules governing progression, pause, evidence acquisition, reanalysis, escalation, rejection, reprioritization, experimental validation, and translational assessment. The proposed architecture distinguishes knowledge claims, observations, curated evidence, computational inference, orthogonal support, mechanism-level validation, and translational contextualization while preserving provenance and evidence independence. Candidate prioritization is therefore framed as a context-dependent decision problem rather than a universal ranking task. The model contributes a stage-based architecture linking knowledge provenance, source and preparation resolution, phytochemical candidate generation, AI-guided identification, target and disease inference, cross-layer evidence translation, uncertainty assessment, decision gates, experimental feedback, and translational reassessment. Its principal research implication is a shift from opaque AI ranking systems toward auditable decision-intelligence platforms capable of documenting why a candidate progresses, pauses, requires additional evidence, undergoes expert adjudication, or is rejected.

Key Words: *Ethnopharmacology, Natural products, Artificial intelligence, Decision intelligence, Candidate prioritization, Evidence integration*

eIJPPR 2025; 15(3):51-68

HOW TO CITE THIS ARTICLE: Lahtinen M, Salo E, Virtanen J. From Ethnopharmacological Knowledge to Drug Candidates: A Decision-Intelligence Model for AI-Guided Natural Product Discovery. Int J Pharm Phytopharmacol Res. 2025;15(3):51-68. https://doi.org/10.51847/IPzwvZpuU5

INTRODUCTION

Natural products continue to provide chemically distinctive scaffolds, privileged structural motifs, biosynthetically evolved interactions, and experimentally tractable starting points for therapeutic discovery, while contemporary analytical, computational, genomic, and

Corresponding author: Mikko Lahtinen
Address: Department of Artificial Intelligence for Bioactive Compound Screening, Faculty of Pharmacy, University of Helsinki, Helsinki, Finland.
E-mail: ✉ mikko.lahtinen@helsinki.fi
Received: 01 March 2025; **Revised:** 20 June 2025; **Accepted:** 21 June 2025

pharmacological methods increasingly expand the ways in which such materials can be detected, characterized, and investigated [1].

Their historical contribution to medicines further demonstrates that natural-product space remains consequential to drug discovery, although retrospective contribution does not imply that any individual organism, preparation, traditional-use claim, or isolated constituent is intrinsically drug-like or therapeutically valid [2]. Taken together, natural products remain consequential to drug discovery, although historical contribution does not make any individual source or ethnopharmacological signal a validated candidate [1, 2]. This distinction is particularly important for ethnopharmacology, where therapeutic use reports, preparation practices, route information, cross-source recurrence, contextual knowledge, and symptom or disease associations may offer structured hypotheses about biologically relevant systems but may also reflect historical contingency, cultural specificity, selective documentation, survivorship effects, or repeated transmission of the same underlying claim. Ethnopharmacological knowledge is therefore most useful when treated as a structured discovery input whose evidentiary meaning depends on provenance, identity, preparation, context, and subsequent empirical challenge.

The scientific value of an ethnopharmacological signal is constrained by the validity of the material and inference chain that connects a knowledge claim to a pharmacological hypothesis. Weak botanical authentication, insufficient chemical characterization, inappropriate controls, poorly specified preparations, and overinterpretation of preliminary bioactivity can undermine downstream conclusions even when the original research question is compelling [3]. A vernacular name may refer to more than one taxon; an accepted species name may conceal geographic or chemotypic variation; different plant parts may contain materially different chemical profiles; and decoctions, infusions, powders, fermented preparations, multi-herb formulations, and solvent extracts should not be treated as interchangeable representations of "the same plant." Similarly, a historical indication does not necessarily map cleanly onto a modern disease ontology, and a repeated traditional use may represent related symptoms, culturally situated explanatory models, or distinct pathological processes. These limitations do not make ethnopharmacological knowledge irrelevant. They make explicit source validation, contextual reconstruction, and evidence-state separation necessary before computational scaling can be scientifically defensible.

Artificial intelligence offers a potentially important set of tools for handling this complexity because machine learning can support discovery functions ranging from representation and property prediction to target identification, screening, and development-oriented analysis [4]. Within an ethnopharmacology-informed workflow, AI may assist natural-language processing of heterogeneous records, recognition of botanical, chemical, disease, and preparation entities, resolution of synonyms, mapping of claims to controlled concepts, generation of molecular representations, similarity analysis, target inference, disease-context matching, integration of omics and network evidence, and prioritization of candidates for further study. Yet these functions operate on data whose meaning may already be uncertain. Scaling extraction does not repair weak provenance; entity resolution does not guarantee taxonomic identity; a graph edge does not establish causation; a high target-prediction score does not establish engagement; and computational similarity does not establish therapeutic equivalence. AI may reduce some information-processing burdens while amplifying database popularity, label bias, duplicated claims, annotation errors, and chemical-space imbalance if those problems are not explicitly governed.

The central challenge is therefore not simply how to apply AI to ethnopharmacological or natural-product data, but how to determine which computational outputs are decision-relevant, under what evidentiary conditions, and with what consequences for subsequent action. AI-enabled natural-product discovery now spans detection, annotation, biosynthetic interpretation, chemical representation, candidate identification, and biological inference, but these capabilities remain dependent on data quality, model applicability, uncertainty handling, and experimental confirmation [5]. The aim of this article is to propose an original decision-intelligence model for translating ethnopharmacological knowledge into prioritized natural-product drug-discovery candidates through a transparent architecture integrating knowledge provenance, source validation, preparation context, chemical evidence, AI-guided identification, target and disease inference, evidence translation, uncertainty, explicit decision rules, candidate prioritization, experimental validation, feedback, and translational assessment. Decision intelligence is used here to mean the structured integration of data, evidence, prediction, uncertainty, constraints, objectives, expert knowledge, decision rules, and feedback to support decisions that are transparent, context-sensitive, auditable, and revisable. It is not treated as equivalent to a prediction model, ranking algorithm, recommendation engine, network pharmacology workflow, classifier, or weighted scoring table.

Ethnopharmacological knowledge as discovery input

Ethnopharmacological knowledge comprises more than a plant–disease association. A decision-relevant record may

include a reported or documented therapeutic use, biological source, local or vernacular name, plant part, preparation practice, formulation context, processing method, route of administration, dose context where available, therapeutic indication, cultural setting, temporal persistence, source independence, practitioner interpretation, and information about who generated, transmitted, documented, or aggregated the claim. Recommended standards for ethnopharmacological field research emphasize the importance of transparent documentation and methodological context, supporting the view that a claim should not enter a computational discovery system as an unqualified fact merely because it appears in a publication or database [6]. In the proposed framework, a knowledge claim is therefore a distinct decision object. Reported use differs from documented observation; documented observation differs from community practice; community practice differs from experimentally observed pharmacology; and each differs from clinical evidence. Historical persistence may increase the relevance of a hypothesis for investigation but does not establish safety. Cross-cultural recurrence may motivate comparison but does not prove a shared molecular mechanism. Practitioner knowledge may contain important contextual interpretation but should not be converted into a decontextualized model label without preserving provenance and the conditions under which the knowledge was produced.

The ethnopharmacological literature itself is heterogeneous in geography, therapeutic focus, disciplinary orientation, terminology, and evidentiary maturity, which limits attempts to treat the field as a uniform database of candidate–indication relationships [7]. A recurrent use may arise from genuinely independent knowledge traditions, from historical diffusion, from repeated citation of a single secondary source, or from database aggregation that obscures common provenance. Consequently, "frequency" can be a misleading decision variable when independence is unknown. The same caution applies to plant reputation and apparent evidence volume: highly visible taxa may accumulate more studies, annotations, and database edges because they are already popular, creating a feedback process in which popularity is mistaken for independent corroboration. A decision-intelligence architecture must therefore retain the difference among primary documentation, secondary reporting, expert synthesis, database aggregation, and text-mined claims. Source recurrence becomes decision-relevant only when its lineage can be inspected and its dependence structure considered. Where independence cannot be established, recurrence should be represented as potentially correlated support rather than multiplied evidence.

Biological-source resolution is a second prerequisite because a plant name is not a sufficient identity guarantee. DNA barcoding and metabarcoding can support authentication of herbal materials and products, but their value depends on reference quality, marker resolution, sample composition, processing, and the biological question being asked; they complement rather than replace taxonomic, chemical, morphological, and contextual assessment [8]. Vernacular names may map to multiple taxa, accepted names may change, synonyms may fragment records, closely related species may be substituted, and species complexes may be difficult to resolve. Even correct species identification does not eliminate geographic, seasonal, developmental, ecological, or chemotypic variation. For decision purposes, taxonomic certainty must therefore be matched to the consequence of the intended action. Exploratory literature retrieval may tolerate bounded ambiguity that would be unacceptable for compound-specific attribution, mechanism testing, or translational progression. The proposed model treats unresolved source identity as U2 taxonomic uncertainty and permits progression only when the residual ambiguity is explicitly compatible with the next decision.

Preparation context creates a parallel problem because processing can alter both the detectability of biological-source information and the chemical meaning of a sample. DNA-based authentication methods for processed products illustrate why the degree of processing and matrix complexity can constrain source verification and why complementary approaches may be necessary [9]. In ethnopharmacological discovery, the same species may be represented as a decoction, infusion, powder, tincture, fermented preparation, fresh material, smoke, topical paste, multi-herb formulation, laboratory extract, enriched fraction, or purified constituent. These objects should not be collapsed into a single "plant candidate." The traditional preparation may selectively extract polar constituents, transform labile compounds, generate metabolites through fermentation or heating, or depend on combinations that are absent from a modern solvent extract. Route of administration further changes the relevance of exposure assumptions. The proposed architecture therefore distinguishes D3 herbal preparation or formulation from D4 extract or fraction and D5 candidate compound, and it treats missing or ambiguous preparation information as U3 preparation uncertainty rather than silently imputing equivalence.

The transition from ethnopharmacological knowledge to modern natural-product discovery is consequently a methodological translation problem rather than a direct validation pathway [10]. A defensible workflow must determine which claim is being considered, whether its provenance is interpretable, whether the biological source

is sufficiently resolved, whether the preparation and indication context are preserved, and whether subsequent chemical and biological evidence pertains to the same decision object. Discovery inputs require provenance and fit-for-purpose source authentication, with method limitations becoming especially important for processed or compositionally complex materials [6, 8, 9]. Ethical governance is part of this input architecture rather than an afterthought: Indigenous knowledge, local knowledge systems, community practices, and practitioner knowledge should not be treated as freely extractable raw material for AI systems. Attribution, ethical reuse, consent context where known, benefit-sharing relevance, decontextualization risk, data sovereignty, and unequal value capture may alter whether and how a knowledge claim should enter the model. These safeguards do not fabricate consent or legal entitlement; they require documentation of what is known, what is unknown, and what governance review is necessary. **Table 1** maps major ethnopharmacological knowledge inputs according to provenance, decision relevance, transformation requirements, uncertainty sources, validation needs, and risks of computational overinterpretation.

Table 1. Ethnopharmacological Knowledge as Input to AI-Guided Natural Product Discovery: Provenance, Decision Relevance, Formalization Requirements, Uncertainty Sources, Validation Needs, and Risks of Overinterpretation

Knowledge input	Primary source	Decision relevance	Required formalization	Key contextual metadata	Primary uncertainty	Validation requirement	Potential computational use	Risk of overinterpretation
Reported traditional use	Field record, historical source, practitioner report, ethnopharmacological publication	Generates a hypothesis about a possible therapeutic context	Separate claim, source, taxon, indication, preparation, route, and reporter	Place, time, knowledge holder or reporting context, terminology, original wording	U1 Source uncertainty; U7 Disease-context uncertainty	Provenance audit and contextual interpretation	Retrieval, semantic indexing, hypothesis generation	Treating reported use as efficacy
Documented community practice	Primary field documentation or community-linked record	May strengthen contextual relevance of a claim	Represent practice, actors, preparation, use conditions, and source lineage	Community context, preparation, route, therapeutic setting, documentation method	U1; U3 Preparation uncertainty	Methodological review; governance assessment where applicable	Structured knowledge representation, context-aware retrieval	Treating documentation as pharmacological confirmation
Practitioner knowledge	Practitioner interview or documented therapeutic tradition	May identify preparation logic, indication distinctions, contraindications, or use sequences	Preserve attribution, epistemic status, and practice context	Practitioner role, therapeutic framework, preparation, patient context where ethically reportable	U1; U7	Expert contextual review; governance check	Relation extraction, ontology extension, hypothesis generation	Converting expert-derived knowledge into unqualified fact
Historical persistence	Longitudinal documentary recurrence	May justify investigation of a durable knowledge signal	Time-stamp sources and preserve source dependence	Period, transmission pathway, changes in formulation or indication	U1	Source-lineage analysis	Temporal knowledge graphs, recurrence analysis	Treating persistence as safety or efficacy
Cross-source recurrence	Multiple records	May suggest repeated relevance when sources are independent	Deduplicate and link records to origin chains	Source identity, publication lineage, geographic and cultural relation	U1	Independence assessment	Evidence mapping, recurrence detection	Counting duplicated propagation as convergence
Cross-cultural recurrence	Distinct cultural or geographic traditions	May motivate comparative investigation	Preserve each context rather than merge labels	Independent provenance, local indications,	U1; U3; U7	Contextual and independence review	Comparative semantic mapping	Treating recurrence as shared molecular mechanism

				preparation differences				
Vernacular plant name	Local record, market record, ethnobotanical source	Entry point for biological-source resolution	Map to candidate taxa with uncertainty retained	Language, region, spelling variants, local taxonomy	U2 Taxonomic uncertainty	Taxonomic verification	Entity resolution, synonym mapping	Treating one vernacular name as one species
Accepted taxonomic identity	Voucher-linked or authenticated source	Supports biological-source admissibility	Normalize accepted name, synonyms, voucher or authentication evidence	Voucher, authority, geography, tissue, authentication method	U2	Fit-for-purpose orthogonal identity check	Entity normalization, record linkage	Assuming same species means same chemistry
Preparation practice	Field record, pharmacopoeial-type documentation, practitioner report	Defines the material actually associated with use	Encode plant part, processing, extraction, formulation, route	Temperature, solvent, duration, combination, storage, route	U3	Preparation reconstruction or verification	Preparation ontology, similarity of preparation contexts	Collapsing decoction, extract, powder, and purified compound
Dose context where available	Primary documentation or practice record	May influence exposure plausibility and experimental design	Preserve units, preparation basis, frequency, route, and uncertainty	Body context where ethically available, duration, co-administration	U3; U9 Exposure uncertainty	Unit/context verification	Exposure hypothesis generation	Inventing precision or extrapolating directly to modern dosing
Traditional indication	Field or historical record	Provides disease- or symptom-context hypothesis	Preserve original concept and separately map to modern terminology	Local disease concept, symptom cluster, therapeutic intent	U7	Domain-expert mapping	Ontology alignment, disease matching	Treating historical indication as a modern diagnosis
Database-aggregated claim	Curated or semi-curated database	Supports retrieval efficiency	Retain upstream sources and aggregation history	Original references, update date, curation method	U1	Back-trace to primary evidence	Large-scale retrieval, knowledge graphs	Treating database frequency as independent evidence
Text-mined claim	NLP-derived relation	Supports discovery of candidate associations	Label as computationally extracted; retain sentence and document provenance	Source document, extraction method, entity confidence, negation/context	U1; U8 Model uncertainty	Human or rule-based contextual verification	Relation extraction, literature-based discovery	Treating extracted co-occurrence as validated relationship
Knowledge-governance status	Community, project, publication, or governance documentation	Determines whether reuse requires review, restriction, attribution, or benefit-sharing consideration	Record known permissions, attribution expectations, unresolved governance questions	Knowledge origin, custodianship, reuse context, sensitivity	U1; U10 Translational uncertainty	Governance review appropriate to context	Access control, provenance-aware system design	Assuming publicly visible knowledge is freely appropriable

AI-guided candidate identification

AI-guided candidate identification should begin only after the relevant decision object and input limitations have been specified. In natural-product discovery, AI can assist candidate generation through cheminformatics, machine learning, representation learning, graph-based models, virtual screening, target inference, and data integration, but the scientific meaning of an output depends on what object is being ranked and what evidence underlies the representation [11]. A "candidate" may be a plant, preparation, extract, fraction, compound, metabolite, target, mechanism, disease indication, or experimental hypothesis. These objects are not interchangeable. Ranking a compound for predicted target activity does not rank the source plant for therapeutic efficacy; identifying a disease-associated target does not validate a traditional indication; and assigning an extract to a promising cluster does not establish which constituent, interaction, or exposure profile explains its activity. The proposed decision-intelligence approach therefore requires AI systems to state the decision object, intended decision, evidence inputs, applicable uncertainty, and downstream validation requirement before ranking is interpreted.

Natural-language processing and literature-based discovery are especially relevant to ethnopharmacology because the evidence landscape is distributed across heterogeneous terminology, narrative reporting, chemical names, taxonomic synonyms, preparation descriptions, disease concepts, and biological relationships. Computational literature-based discovery can support retrieval of latent relationships and accelerate identification of potentially useful natural-product associations, while also facing substantial challenges of corpus quality, domain terminology, relation validity, and biological interpretation [12]. Named-entity recognition may identify plant, compound, target, disease, and preparation mentions; relation extraction may propose links among them; entity resolution may reconcile synonyms; semantic mapping may connect historical indications to candidate modern concepts; and knowledge graphs may preserve relationships among knowledge claims, sources, taxa, preparations, compounds, targets, pathways, diseases, and experiments. However, extracted relations should retain their provenance and epistemic status. A relation obtained from text mining is not equivalent to a curated mechanistic relationship, and neither is equivalent to an experimentally observed interaction. Where generative systems are used, unsupported or fabricated associations should remain E0 unverified claims until independently checked rather than being absorbed into the evidence graph as facts.

Molecular representation is another major determinant of AI-guided identification because algorithms act on encodings of molecules rather than on chemistry in the abstract. Comparative evaluation of learned and engineered molecular representations shows that predictive performance and generalization depend on task, data regime, architecture, and representation choice [13]. For natural products, these issues may be intensified by stereochemical complexity, macrocycles, glycosylation, unusual scaffolds, mixtures, sparse labels, and chemical regions underrepresented in mainstream training sets. Molecular fingerprints may support similarity searching; graph representations may capture atom–bond topology; learned embeddings may support property prediction or retrieval; and graph neural networks may identify task-relevant structural patterns. Yet similarity remains representation-dependent. A candidate may be "close" under one fingerprint and distant under another, and structural resemblance does not establish equivalent exposure, selectivity, efficacy, or safety. The decision consequence of representation uncertainty should therefore be explicit, particularly when a candidate is outside the apparent applicability domain of the model.

Target prediction and drug–target interaction inference can extend candidate identification from chemical resemblance toward mechanistic hypotheses, but predicted targets must remain computationally inferred evidence unless independently supported. Large-scale comparisons of machine-learning approaches for target prediction demonstrate substantial method dependence and the importance of benchmark design, data composition, and evaluation assumptions [14]. An ethnopharmacology-informed candidate may be linked to a known target through curated evidence, assigned a predicted target through supervised learning, connected through a knowledge graph, or placed near a disease module through network analysis. These are distinct epistemic routes. Disease mapping adds further ambiguity because historical indications may not correspond directly to contemporary diagnoses and because target–disease associations can be context-specific. The proposed framework therefore separates D7 candidate target, D8 candidate mechanism, and D9 candidate disease indication and requires target–disease coherence to be assessed rather than assumed from network overlap. Model outputs may justify focused investigation, but they do not independently establish target engagement or causal mechanism.

AI-guided identification becomes decision-relevant when it is linked to a clear next action and an explicit validation requirement. Deep-learning studies that connect model-based prioritization to subsequent experimental testing demonstrate the general principle that computational identification can be used to select candidates for empirical challenge rather than treated as its own confirmation [15]. In an ethnopharmacology-informed setting, candidate ranking should preserve knowledge provenance, source

identity, preparation context, chemical resolution, model applicability, and the distinction among measured, curated, predicted, inferred, and expert-derived evidence. Major risks include training-data bias, popular-entity bias, database circularity, label leakage, chemical-space bias, target-annotation bias, disease-label ambiguity, correlated models that appear falsely independent, and unsupported relations produced by generative systems. Consequently, the model asks not only "Which candidate ranks highest?" but "Is candidate plausibility supported beyond model ranking, what uncertainty could reverse the decision, what evidence would be most informative next, and who must review the result before resources are committed?" AI supports identification; it does not confirm validity.

Evidence translation and decision rules

Decision intelligence begins where prediction alone becomes insufficient. A model output must be translated into an explicit decision variable whose meaning is tied to an evidence source, a decision object, an intended action, and a consequence of error. Uncertainty quantification in drug design is therefore relevant not merely as a technical property of predictive models but as an input to decisions about whether to progress, pause, test, or reconsider a candidate [16]. In the proposed architecture, heterogeneous evidence is not collapsed into a single undifferentiated score. Instead, knowledge provenance, source identity, preparation fidelity, chemical resolution, observed bioactivity, target coherence, disease relevance, evidence independence, model applicability, exposure plausibility, safety concerns, novelty, experimental feasibility, and translational plausibility are represented as distinct decision variables. A candidate can be strong on one dimension and weak on another. High ethnopharmacological relevance cannot compensate automatically for unresolved taxonomic identity; strong predicted target coherence cannot erase poor chemical resolution; and novelty cannot justify progression when the candidate is experimentally uninterpretable.

The model classifies evidence into exactly ten states. E0 Unverified claim includes assertions whose provenance or validity has not been established. E1 Documented ethnopharmacological evidence records a use or knowledge claim without converting it into pharmacological proof. E2 Source-validated evidence indicates that the biological source has been sufficiently verified for the intended decision. E3 Chemically resolved evidence indicates fit-for-purpose identification of a chemical object. E4 Experimentally observed bioactivity records measured biological activity without automatically establishing target, mechanism, or translation. E5 Curated mechanistic evidence captures established knowledge from curated sources or literature without assuming candidate-specific validation. E6 Computationally inferred evidence includes model predictions, network inferences, and other computational hypotheses. Empirical comparisons of uncertainty methods for molecular property prediction show why outputs within E6 require method-aware reliability assessment rather than automatic authorization of downstream action [17]. E7 Orthogonally supported evidence reflects convergence across meaningfully distinct evidence routes. E8 Experimentally validated mechanism-level evidence requires direct mechanism-relevant testing. E9 Translationally contextualized evidence incorporates exposure, pharmacokinetic, safety, disease, feasibility, and development context.

Evidence state is not equivalent to evidence quality, and higher state labels should not be interpreted as a universal ordinal score. A carefully documented E1 claim may be methodologically stronger than a poorly executed E4 assay; a well-calibrated E6 prediction may be more useful for a limited screening decision than a contextually irrelevant E5 association; and E8 mechanism-level support may remain translationally weak when exposure is implausible. Trustworthy AI-oriented drug-discovery reasoning therefore requires uncertainty to be visible in relation to the decision being contemplated [18]. The architecture also preserves evidence provenance categories—measured, curated, predicted, inferred, expert-derived, and clinically contextualized—because apparently convergent findings may share the same underlying data or assumption. A curated target relation, a model trained on that relation, and a knowledge graph derived from the same source are not three independent confirmations. Similarly, a database frequency, a review statement, and a text-mined extraction may reproduce one ancestral claim. Evidence translation must therefore assess lineage and independence before quantity is interpreted as convergence.

Uncertainty is represented through ten predefined classes: U1 source uncertainty, U2 taxonomic uncertainty, U3 preparation uncertainty, U4 chemical identity uncertainty, U5 bioactivity uncertainty, U6 target uncertainty, U7 disease-context uncertainty, U8 model uncertainty, U9 exposure uncertainty, and U10 translational uncertainty. Comparative work on scalable uncertainty estimation for molecular property prediction indicates that different methods capture different aspects of reliability and can behave differently across datasets and prediction settings [19]. The proposed model consequently does not convert uncertainty into invented probabilities or a universal confidence threshold. Instead, each class is linked to a source, decision consequence, detectability route, mitigation strategy, and progression condition. U2 may block compound attribution when taxonomic ambiguity

changes the source identity; U6 may permit computational prioritization while blocking a claim of validated engagement; U8 may trigger reanalysis when a candidate is outside the model's domain; U9 may prevent translation-oriented progression when plausible exposure is absent; and U10 may remain decisive even after mechanism-level support if safety, feasibility, or disease-context assumptions are unresolved.

Decision rules specify what action follows from the joint interpretation of evidence and uncertainty. Progress is permitted when the relevant gate is passed and the next-stage question is meaningful. Pause is used when a material gap is potentially recoverable. Acquire Evidence requests a specific observation capable of changing the decision. Reanalyse applies when new information, data error, model instability, applicability-domain concern, or revised mapping changes interpretation. Escalate routes high-impact conflict, model–expert disagreement, culturally sensitive reuse questions, or unresolved disease-context ambiguity to documented adjudication. Reject terminates the specified candidate pathway when failure is decisive. Reprioritize updates relative candidate positions after new evidence. Validate Experimentally moves a coherent, testable hypothesis toward empirical challenge. Assess Translation examines whether experimentally supported evidence is compatible with exposure, safety, disease context, reproducibility, and development feasibility. Critical analyses of AI in drug discovery show why chemical and biological data volume, benchmark performance, and apparent model success cannot be assumed to establish data validity, independence, or prospective decision value [20]. Accordingly, missing evidence, conflicting evidence, correlated evidence, weak source identity, poor exposure plausibility, and high uncertainty trigger different actions rather than being absorbed into a single lower score.

Progression is conditional because the appropriate decision depends on the object and the consequence of being wrong. A D1 knowledge claim with incomplete provenance may move from S0 Insufficient Evidence to S1 Evidence Acquisition Required; a source-resolved D5 compound may become S2 Eligible for Computational Analysis; a model-supported object may reach S3 Computationally Prioritized; unresolved conflict may require S4 Expert Adjudication; a coherent and falsifiable hypothesis may reach S5 Eligible for Experimental Validation; relevant empirical support may permit S6 Experimentally Supported; contradictory or destabilizing findings may trigger S7 Requires Reprioritization; integrated translational plausibility may support S8 Translation-Oriented Candidate; and decisive failure may lead to S9 Rejected. These trajectories are not mandatory or linear. Uncertainty should be treated as a decision-relevant property of AI-supported drug discovery rather than an optional confidence annotation [16-18]. The model therefore allows forward movement, backward movement, pause, escalation, reanalysis, evidence acquisition, and termination while requiring the reason for each transition to be documented. **Table 2** compares major evidence layers used in AI-guided natural-product discovery and specifies how each layer should influence decision rules without conflating prediction, observation, and validation.

Table 2. Evidence Translation and Decision Rules for AI-Guided Natural Product Discovery: Evidence Layers, Decision Variables, Uncertainty Constraints, Progression Logic, Rejection Conditions, and Validation Requirements

Evidence layer	Evidence state	Primary decision variable	Potential decision contribution	Major uncertainty	Progression condition	Pause or escalation condition	Rejection condition	Validation requirement	Interpretive safeguard
Unverified knowledge assertion	E0	Claim admissibility	Identifies a possible information need	U1	Only when provenance can be investigated	Pause for source acquisition; escalate if sensitive or contested	Non-recoverable fabrication or decisive provenance failure	Source verification	Unverified claim ≠ documented evidence
Documented ethnopharmacological evidence	E1	Provenance strength; contextual relevance	Supports structured hypothesis generation	U1; U3; U7	Claim context is sufficiently preserved for the next question	Missing preparation, ambiguous indication, uncertain independence	Claim pathway rejected if decisive provenance failure invalidates it	Primary-source and contextual review	Documented use ≠ efficacy
Source-validated evidence	E2	Biological-source credibility	Permits source-specific chemical or	U2	Identity confidence is fit for intended consequence	Conflicting taxonomy, substitution, unresolved	Decisive misidentification invalidates source-	Orthogonal authentication proportional to use	Plant name ≠ taxonomic certainty

			computational analysis			vernacular mapping	dependent pathway		
Chemically resolved evidence	E3	Chemical identity credibility	Supports compound-, metabolite-, fraction-, or structure-specific reasoning	U4	Resolution is sufficient for contemplated claim	Isomer ambiguity, weak annotation, inconsistent analytical evidence	Material identity failure invalidates object-specific inference	Fit-for-purpose analytical confirmation	Compound annotation ≠ exposure or activity
Experimentally observed bioactivity	E4	Bioactivity support	May justify mechanism investigation or comparative prioritization	U5	Activity is reproducible, context-relevant, and interpretable	Assay interference, weak controls, discordant replication	Repeated orthogonal evidence refutes central activity claim	Replication, counterscreens, orthogonal assays	Bioactivity ≠ mechanism or clinical effect
Curated mechanistic evidence	E5	Mechanistic plausibility	Connects candidates to established biological knowledge	U6; U7	Curated relation is relevant to the specific candidate and context	Relation is indirect, outdated, or context-mismatched	Rejection only when the central candidate logic depends on a decisively false relation	Candidate-specific mechanistic assessment	Curated association ≠ candidate-specific validation
Computationally inferred evidence	E6	Predicted candidate, target, property, or disease coherence	Supports prioritization, hypothesis generation, and experiment selection	U6; U7; U8	Model is applicable, inference is stable enough for the decision, and provenance is retained	OOD concern, model disagreement, leakage risk, unstable ranking	Reject prediction-dependent pathway when inference collapses under valid challenge and no alternative rationale remains	External challenge, sensitivity analysis, experimental testing	Prediction ≠ observation
Orthogonally supported evidence	E7	Evidence convergence and independence	Strengthens progression toward focused validation	U5–U8	Support arises from meaningfully distinct evidence routes	Apparent convergence shares source, labels, or model assumptions	Reject convergence claim if support collapses after dependency audit	Independence analysis and targeted confirmation	Evidence quantity ≠ evidence independence
Experimentally validated mechanism-level evidence	E8	Mechanism support	Supports advanced biological and translational assessment	U5; U6; U9	Mechanism-relevant evidence is replicated and contextually interpretable	Conflicting mechanism tests, uncertain exposure, poor model-system relevance	Reject mechanism pathway when decisive experimental evidence refutes it	Orthogonal perturbation and replication	Mechanism support ≠ clinical efficacy
Translationally contextualized evidence	E9	Translational plausibility	Informs translation-oriented progression, pause, or rejection	U9; U10	Exposure, safety, disease context, reproducibility, and feasibility are	Major unresolved PK, metabolism, toxicity, formulation, or disease-context issue	Decisive safety, exposure, feasibility, or contextual failure	External/prospective testing appropriate to stage	Translation-oriented ≠ clinically proven

					collectively plausible				
Knowledge-provenance layer	Cross-state	Traceability and attribution	Determines whether evidence is admissible and auditable	U1	Lineage is inspectable	Duplicated, circular, or contested provenance	Decisive non-recoverable provenance failure	Provenance audit	Repetition ≠ independent corroboration
Evidence-independence layer	Cross-state	Independence of support	Distinguishes genuine convergence from redundancy	U1; U5–U8	Dependency structure is sufficiently understood	Shared source, correlated models, derivative databases	Reject convergence rationale if all support is derivative	Evidence-lineage mapping	Multiple outputs ≠ multiple independent findings
Exposure and safety layer	Cross-state, often E7–E9	Exposure plausibility; safety concern	Modifies experimental and translational priority	U9; U10	Exposure and safety are compatible with intended use context	Missing PK, active-metabolite, toxicity, or formulation evidence	Decisive exposure or safety failure	Fit-for-purpose PK, metabolism, toxicity, and formulation assessment	Compound presence ≠ biologically relevant exposure
Expert-derived evidence	Cross-state	Contextual adjudication	Challenges mappings, resolves ambiguity, or identifies missing considerations	U1–U10	Expertise is relevant, rationale documented, and challengeable	Expert disagreement or undocumented override	No automatic rejection based on status or authority alone	Transparent adjudication record	Human judgment ≠ infallibility
Negative or contradictory evidence	Cross-state	Refutation strength; update consequence	Triggers reanalysis, reprioritization, pause, or rejection	U5–U10	Contradiction is credible and object-specific	Methodological ambiguity prevents interpretation	Decisive replicated contradiction invalidates pathway	Replication or alternative explanation testing where justified	Negative evidence must not be hidden
Experimental feedback	Cross-state	Evidence-update consequence	Revises candidate state, uncertainty, and next action	U5–U10	New evidence changes a material decision variable	Unexpected findings require adjudication	Decisive failure may terminate the pathway	Versioned evidence updating	Feedback changes decisions; it does not guarantee improvement

Candidate prioritization pathways

Candidate prioritization should not be reduced to one universal ranking scheme because the appropriate pathway depends on which decision object is sufficiently resolved, what evidence is available, and which uncertainty is most consequential. An ethnopharmacology-first pathway begins with a contextually strong knowledge signal and proceeds through provenance assessment, source resolution, preparation reconstruction, chemical investigation, and candidate generation. A chemistry-first pathway begins with a resolved natural-product compound or metabolite and uses similarity, representation learning, property prediction, target inference, or disease matching to identify plausible next questions. A bioactivity-first pathway begins with experimentally observed activity and asks whether the responsible candidate object can be resolved and whether mechanism-oriented testing is justified. A target-first pathway begins from a disease-relevant target and searches natural-product space for chemically and biologically plausible candidates. A network/omics-first pathway begins from a disease module, molecular signature, or systems-level perturbation and seeks candidate matches across heterogeneous evidence. Iterative selection methods such as active learning are especially relevant when the objective is not to rank an entire chemical universe once, but to choose the next informative experiment and revise priorities as evidence accumulates [21]. The pathway selected should therefore reflect the decision problem rather than the convenience of the available algorithm.

Pathway A—Ethnopharmacology-First is most appropriate when the input signal has interpretable provenance, preparation context, and source identity but the responsible chemistry or mechanism remains unresolved. Progression follows strong contextual knowledge → source resolution → chemical investigation → candidate generation, with possible branching toward S1 Evidence Acquisition Required when preparation or taxonomy is inadequate. Pathway B—Chemistry-First is appropriate when a natural-product compound or metabolite is structurally resolved independently of a strong traditional-use signal. Progression follows resolved natural-product compound → similarity or representation analysis → target and disease matching, but candidate ranking remains contingent on applicability, biological relevance, and exposure. Pathway C—Bioactivity-First is appropriate when reproducible activity is observed in an extract, fraction, compound, or metabolite and the immediate decision concerns identity and mechanism. The pathway follows observed bioactivity → candidate identity → mechanism inference → target validation. Empirical AI-guided discovery studies illustrate the broader principle that computational selection becomes more meaningful when it leads to explicit synthesis or testing rather than being interpreted as validation by ranking alone [22]. The decision safeguard is to define in advance what result would support, contradict, or leave the hypothesis unresolved.

Pathway D—Target-First begins with a biologically credible disease-relevant target and searches natural-product chemical space for candidates with plausible target interactions, structural relationships, or supporting biological evidence. It is useful when disease biology is comparatively well resolved but candidate chemistry is open. Pathway E—Network/Omics-First begins from disease modules, pathway perturbations, transcriptomic or other omics signatures, or heterogeneous biomedical relations and attempts to identify candidates whose evidence profiles intersect with those systems. Heterogeneous network integration can support candidate prioritization across biological entities and relations, but the resulting connectivity should be treated as evidence for ranking or hypothesis generation rather than proof that the inferred path is mechanistically causal [23]. This distinction is especially important when ethnopharmacological signals are mapped onto modern disease networks: an apparent overlap may arise from broadly connected proteins, literature popularity, shared annotations, or downstream consequences rather than a specific therapeutic mechanism.

Evidence convergence should therefore be evaluated according to independence, lineage, consistency, context relevance, and maturity. Independent support occurs when substantively distinct methods or observations contribute non-derivative evidence to the same decision question. Shared-source evidence occurs when different outputs ultimately derive from one dataset or publication. Duplicated evidence arises when the same relationship is repeated across databases or reviews. Circular evidence occurs when a prediction is trained on a relationship and later treated as independent confirmation of that same relationship. Correlated model outputs arise when multiple algorithms share labels, representations, features, or training sets and therefore cannot be interpreted as separate replications. Knowledge-graph frameworks can integrate diverse relations and support candidate ranking, yet graph-derived associations remain dependent on the provenance, completeness, bias, and semantics of the underlying nodes and edges [24]. The proposed model therefore uses G7 Evidence Independence Gate before apparently convergent support can justify stronger progression.

Candidate prioritization should balance ethnopharmacological relevance, source identity, preparation fidelity, chemical plausibility, observed bioactivity, target coherence, disease relevance, evidence independence, uncertainty, exposure plausibility, safety, novelty, and experimental feasibility without inventing universal numerical weights. An attractive candidate may be deprioritized because its chemistry is unresolved, its target inference is unstable, its evidence streams are derivative, or plausible exposure is absent. Conversely, a modestly ranked candidate may progress because the experiment capable of resolving its key uncertainty is feasible and decision-relevant. Reliability-aware ADMET modelling demonstrates why predicted developability properties should be interpreted together with information about predictive uncertainty rather than as deterministic filters [25]. The resulting priority is therefore an action recommendation under stated evidence and uncertainty conditions, not a declaration of candidate validity. **Table 3** synthesizes the principal candidate-prioritization pathways and distinguishes complementary evidence convergence from correlated, redundant, or potentially misleading signals.

Table 3. Candidate Prioritization Pathways in AI-Guided Natural Product Discovery: Evidence Inputs, Decision Objectives, Convergence Logic, Uncertainty, Experimental Readiness, Translational Constraints, and Failure Risks

Prioritization pathway	Primary candidate object	Key evidence inputs	Decision objective	Evidence-convergence logic	Primary uncertainty	Experimental readiness requirement	Translational constraint	Major failure risk	Decision safeguard

Pathway A — Ethnopharmacology-First	D1 knowledge claim; D2 source; D3 preparation; later D4–D6	Provenance, documented use, preparation, route, source identity, chemical evidence	Determine whether a contextual knowledge signal merits chemical and biological investigation	Independent contextual evidence may justify investigation only after source and preparation resolution	U1, U2, U3, U7	A defined source/preparation object and a testable chemical or biological question	Historical indication may not map directly to modern disease context	Treating recurrence or persistence as efficacy, safety, or mechanism	Apply G1–G4 before AI ranking; preserve original context
Pathway B — Chemistry-First	D5 compound; D6 metabolite	Structure, annotation, occurrence, similarity, molecular representation, property prediction	Identify plausible targets, properties, diseases, or comparator candidates	Chemical and biological evidence should converge without assuming that structural similarity establishes therapeutic equivalence	U4, U8, U9	Chemical identity sufficient for the intended experiment	Exposure, metabolism, and formulation may invalidate otherwise attractive predictions	Chemical-space bias; similarity overreach; annotation error	Apply applicability checks and G4–G6; retain representation dependence
Pathway C — Bioactivity-First	D4 extract/fraction; D5 compound; D6 metabolite; D10 hypothesis	Observed activity, controls, candidate identity, orthogonal assays, target evidence	Determine whether observed activity justifies mechanism-oriented validation	Reproducible activity plus independently resolved identity and mechanistic support strengthens progression	U4, U5, U6	Activity must be interpretable, reproducible, and linked to a falsifiable hypothesis	Potency without plausible exposure or selectivity may not translate	Assay interference; activity attribution error; post hoc mechanism assignment	Counterscreens, orthogonal testing, explicit G5–G9 logic
Pathway D — Target-First	D7 target; D5/D6 candidate	Disease biology, validated or curated target knowledge, candidate libraries, DTI prediction, structural and biological evidence	Search natural-product space for target-relevant candidates	Target relevance and candidate plausibility must be independently established	U6, U8, U9	A target-specific experiment capable of testing engagement or functional consequence	Target relevance may be disease-stage-, tissue-, or exposure-dependent	Treating predicted binding as engagement; annotation bias	Separate predicted from validated targets; apply G6
Pathway E — Network/Omics-First	D8 mechanism; D9 indication; D5–D7 candidates	Disease modules, pathway networks, omics signatures, knowledge graphs,	Identify candidates coherent with systems-level disease evidence	Convergence is meaningful only when graph, omics, literature, and experiment	U6, U7, U8	A mechanistically discriminating validation strategy	Network coherence may not imply achievable exposure or causal intervention	Centrality bias; literature bias; circular evidence; non-causal overlap	Apply G7 Evidence Independence Gate and causal caution

		candidate profiles		al sources are not merely derivative					
Hybrid ethnopharmacology–chemistry	D1–D6	Contextual knowledge plus resolved chemistry	Identify compounds that plausibly represent the reported preparation	Context and chemistry must refer to the same material and preparation	U2–U4	Verified material correspondence	Purified compound may not represent preparation-level exposure	Decontextualizing the traditional preparation	Maintain D3–D5 distinction
Hybrid chemistry–target	D5–D7	Structure, similarity, target prediction, curated biology	Select candidates for target-specific testing	Distinct computational and curated routes may strengthen plausibility if independent	U6, U8	Target-engagement or functional assay	Selectivity and exposure may remain limiting	Correlated model outputs presented as replication	Model-lineage audit and orthogonal target evidence
Hybrid bioactivity–omics	D4–D8	Phenotypic activity, transcriptomics or other omics, pathway analysis	Generate mechanism hypotheses from observed perturbation	Observed phenotype and systems response should be linked through testable mechanistic alternatives	U5–U8	Replication and perturbation-based discrimination	Model-system relevance may limit translation	Reverse inference from signature to mechanism	Predefine competing hypotheses
Exposure-aware prioritization	D5, D6, D9, D10	Chemical evidence, PK, metabolism, local/systemic exposure, safety	Determine whether biological hypotheses are exposure-plausible	Activity and mechanistic support gain decision value only when relevant exposure is plausible	U9, U10	Exposure-relevant test design	Metabolites, formulation, route, and tissue distribution may dominate	Compound presence mistaken for active exposure	Explicit U9 assessment before S8
Feedback-driven reprioritization	D5–D10	New experimental, negative, contradictory, or translational evidence	Update candidate status and choose next action	New evidence can strengthen, weaken, or reverse prior convergence	U5–U10	Evidence must be versioned and linked to the affected decision variable	Repeated updating may expose feasibility or safety limits	Hiding negative evidence; preserving sunk-cost candidates	Enter S7 and rerun relevant gates

Proposed decision-intelligence model

The core decision problem is: Which ethnopharmacologically informed natural-product candidates should progress to the next evidence-generating stage, under what uncertainty conditions, and with what validation requirements? This formulation deliberately

shifts the objective from prediction accuracy or rank maximization toward action selection under heterogeneous evidence. End-to-end applications of machine learning across discovery and development demonstrate the importance of linking computational functions across stages rather than treating one isolated model as the complete discovery process [26]. The proposed decision-intelligence model therefore integrates data, evidence, predictions, uncertainty, constraints, objectives, expert knowledge, decision rules, and feedback. It does not require every decision object to traverse the same pathway, and it does not authorize progression because an aggregate score is high. Each action must be tied to a specified object, evidence state, uncertainty profile, gate condition, and validation consequence.

The model contains exactly thirteen stages: M1 Knowledge Intake; M2 Knowledge Provenance Assessment; M3 Biological Source Resolution; M4 Preparation and Context Resolution; M5 Phytochemical Candidate Generation; M6 AI-Guided Candidate Identification; M7 Target and Disease-Context Inference; M8 Cross-Layer Evidence Translation; M9 Uncertainty and Conflict Assessment; M10 Candidate Prioritization; M11 Experimental Validation Planning; M12 Evidence Updating and Reprioritization; and M13 Translational Readiness Assessment. M1 captures the claim before endorsing it. M2 asks whether its provenance and contextual meaning are sufficient for further use. M3 resolves the biological source. M4 determines whether preparation and therapeutic context are sufficiently represented. M5 identifies chemically plausible objects. M6 applies AI to candidate identification rather than validation. M7 tests coherence among candidate, target, mechanism, and disease context. M8 translates heterogeneous evidence into explicit decision variables. M9 makes uncertainty and conflict visible. M10 determines relative action priority. M11 selects a falsifiable experiment. M12 updates evidence and priorities. M13 evaluates exposure, safety, reproducibility, disease relevance, and development plausibility. Critical analyses of AI in drug discovery reinforce why retrospective predictive performance should not be treated as equivalent to prospective decision value, making these stage boundaries necessary rather than decorative [27].

The architecture distinguishes ten decision objects—D1 knowledge claim, D2 biological source, D3 herbal preparation or formulation, D4 extract or fraction, D5 candidate compound, D6 candidate metabolite, D7 candidate target, D8 candidate mechanism, D9 candidate disease indication, and D10 candidate experimental hypothesis—and retains the evidence states E0 through E9 defined earlier. It also applies exactly ten uncertainty classes: U1 source, U2 taxonomic, U3 preparation, U4 chemical identity, U5 bioactivity, U6 target, U7 disease context, U8 model, U9 exposure, and U10 translational uncertainty. Data-centric limitations, including bias, inconsistency, benchmark weakness, and human influences on data generation and interpretation, justify treating uncertainty and conflict as explicit decision variables rather than hidden residual error [28]. Ten decision gates govern transitions: G1 Knowledge Provenance Gate; G2 Source Identity Gate; G3 Preparation Context Gate; G4 Chemical Resolution Gate; G5 Candidate Plausibility Gate; G6 Target–Disease Coherence Gate; G7 Evidence Independence Gate; G8 Uncertainty Acceptability Gate; G9 Experimental Testability Gate; and G10 Translational Plausibility Gate. The permitted states are S0 Insufficient Evidence, S1 Evidence Acquisition Required, S2 Eligible for Computational Analysis, S3 Computationally Prioritized, S4 Requires Expert Adjudication, S5 Eligible for Experimental Validation, S6 Experimentally Supported, S7 Requires Reprioritization, S8 Translation-Oriented Candidate, and S9 Rejected. Current assessments of AI drug discovery likewise emphasize the continuing importance of ground truth, realistic validation, and human intervention beyond computational ranking [29].

Figure 1 presents the proposed decision-intelligence model linking ethnopharmacological knowledge provenance, AI-guided candidate identification, evidence translation, uncertainty assessment, decision gates, candidate prioritization, experimental validation, and translational reassessment.

Figure 1. Decision-Intelligence Model for Translating Ethnopharmacological Knowledge into AI-Guided Natural Product Drug Candidates

Alt-text description

A wide left-to-right decision-intelligence architecture begins with ethnopharmacological knowledge claims and proceeds through provenance assessment, taxonomic and source resolution, preparation-context verification, phytochemical candidate generation, AI-guided candidate identification, target and disease-context inference, cross-layer evidence integration, uncertainty and conflict assessment, candidate prioritization, experimental validation, and translational assessment. Decision gates between stages allow progression, pause, evidence acquisition, expert escalation, reprioritization, or rejection. Feedback arrows return new evidence to earlier stages.

Human oversight is concentrated at high-consequence points rather than distributed as a ceremonial approval layer across every computational step. At M2, human review is needed when provenance, attribution, source independence, or ethical reuse is disputed. At M3 and M4, taxonomic and preparation expertise is required when ambiguity could alter the candidate object. At M7 and M9, disease-domain, pharmacological, and computational expertise may challenge mappings, model assumptions, or apparently convergent evidence. At M10 and M13, documented adjudication is necessary when uncertainty, feasibility, safety, exposure, or governance conflicts materially affect resource allocation. Expert judgment is itself fallible; therefore, an override must record the original model-supported decision, the challenged assumption, the evidence considered, the reviewer's rationale, unresolved dissent, and the condition under which the decision should be revisited. Human review is a challenge mechanism, not an unrecorded privilege to replace one opaque ranking with another.

Feedback learning is implemented through explicit evidence updating rather than by assuming that every new observation should improve the model. Experimental outcomes may support, partially support, contradict, or fail to resolve the candidate hypothesis. Each outcome updates the affected decision object, evidence state, uncertainty class, and gate status. A supported result may move S5 to S6 but does not automatically generate S8. A contradictory finding may move S5 or S6 to S7 for reprioritization or to S9 if the failure is decisive. An inconclusive result may return the object to S1 when a specific missing observation is identifiable. Negative evidence is therefore retained as an information asset: failed activity replication, absent target engagement, unstable predictions, poor exposure, unexpected toxicity, and unsuccessful preparation reconstruction can prevent repeated investment in invalid assumptions. M12 Evidence Updating and Reprioritization consequently requires versioned evidence lineage so that changes in candidate status can be traced to specific new information rather than unexplained rank drift.

The proposed model is a decision architecture rather than a predictive algorithm because it represents objects, evidence states, uncertainty, constraints, gates, actions, human challenge, feedback, and terminal outcomes in addition to predictions. It can incorporate machine-learning classifiers, knowledge graphs, graph neural networks, target-prediction systems, molecular similarity tools, multimodal models, or active-learning components, but none of those components is synonymous with the architecture. Auditability requires documentation of what entered the model, which transformations occurred, what evidence was measured or inferred, which uncertainty was considered, what gate condition was applied, why the candidate progressed or failed, who intervened, what experiment was requested, and how new evidence changed the decision. Ethical knowledge governance is likewise cross-cutting: a technically attractive candidate pathway should not proceed as though provenance, attribution, community interests, benefit-sharing relevance, or data sovereignty were external to discovery quality. By making these elements explicit, the model seeks to support transparent and revisable portfolio decisions without presenting itself as experimentally or prospectively validated.

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

Research implications

For ethnopharmacology research, the model implies a shift from collecting plant–use associations toward constructing evidence objects that can survive downstream computational interpretation. Taxonomic provenance, preparation metadata, original indication wording, source lineage, route information, context of use, evidence independence, and unresolved ambiguity should be recorded in forms that remain accessible when records are aggregated or transformed. This requirement is scientific and ethical. Peer-reviewed analysis of access and benefit sharing in biodiversity-related research illustrates the practical complexity of translating biological resources and associated knowledge into research and development settings while addressing opportunity, risk, and unequal participation [30]. Ethnopharmacology-informed AI platforms should therefore include provenance and governance fields rather than treating attribution or benefit-sharing considerations as administrative metadata added after candidate selection. The research implication is not that one universal governance rule applies across all contexts, but that knowledge admissibility and reuse conditions can materially affect a discovery decision.

For Indigenous and other collectively held knowledge systems, platform design should also resist the assumption that public visibility makes knowledge freely extractable or that individual-level consent concepts exhaust collective governance questions. The CARE principles emphasize collective benefit, authority to control, responsibility, and ethics in Indigenous data governance [31]. Applied cautiously to ethnopharmacology-informed discovery, these principles support system requirements for attribution, provenance visibility, decision authority where applicable, documentation of unresolved reuse conditions, and scrutiny of unequal value capture. Such requirements should not be reduced to a generic ethics score. They are decision constraints that may trigger pause, evidence acquisition, expert or governance escalation, restricted use, or rejection of a proposed data pathway. Future research should evaluate how provenance-aware computational infrastructures can preserve culturally situated meaning while enabling legitimate scientific investigation without fabricating community permissions or assuming that all knowledge systems can be normalized into one ontology without loss.

For natural-product informatics and AI model development, the principal data requirement is not merely more observations but better-typed observations. Systems should distinguish taxonomically verified from unverified sources, measured compounds from database-reported constituents, curated targets from predicted targets, direct bioactivity from inferred mechanism, independent experiments from derivative evidence, and negative results from absent reporting. Versioning is essential because accepted taxonomy, chemical annotations, disease mappings, model outputs, and evidence states can change. Uncertainty fields should be stored alongside predictions rather than generated only for publication figures. Negative evidence should be represented in machine-readable form where ethically and scientifically appropriate. External validation, calibration, applicability-domain analysis, and prospective evaluation should be designed around decision consequences rather than benchmark performance alone. The uncertainty literature reviewed earlier supports this direction because reliable prioritization depends on knowing when model confidence is poorly aligned with prediction error, especially outside familiar data regimes [16].

For computational pharmacology, candidate screening, and experimental design, the model suggests that active learning and feedback should be organized around decision impact. The next experiment should be selected because it can resolve a material uncertainty or discriminate among competing hypotheses, not simply because the model predicts high activity. Prospective studies should therefore evaluate whether a decision-intelligence system improves the choice of experiments, reduces avoidable progression of invalid objects, surfaces uncertainty earlier, records negative evidence, or changes portfolio decisions in transparent ways. Methodological priorities include external validation, calibration, human–AI interaction, structured expert adjudication, causal inference where scientifically appropriate, multi-omics integration, exposure-aware prioritization, metabolism-aware modelling, and feedback loops connecting experimental evidence to model and evidence-state updates. AI-enabled natural-product research is particularly suited to such evaluation because its inputs span text, taxonomy, spectra, structures, assays, targets, networks, omics, and translational constraints, making hidden evidence conflation a persistent risk [5].

At the platform level, the long-term implication is a movement from AI ranking systems toward auditable decision-intelligence platforms. Such platforms should preserve the provenance of each knowledge claim, source resolution status, preparation context, chemical identity, evidence state, uncertainty class, dependency structure, model version, decision rationale, human intervention, requested validation, negative outcome, and subsequent update. Knowledge graphs may provide useful infrastructure for representing heterogeneous relations, but graph connectivity should remain typed by evidence source and epistemic status rather than presented as uniform truth. Portfolio interfaces should allow decision makers to inspect why two candidates with similar ranks receive different actions—for example, why one

progresses to experiment while another pauses for taxonomic resolution, why a third is escalated because of conflicting evidence, and why a fourth is rejected because exposure failure is decisive. The research priority is therefore not simply better candidate ranking, but empirical evaluation of whether transparent decision architectures improve the quality, reproducibility, ethical defensibility, and revisability of natural-product discovery choices.

CONCLUSION

This article proposes an original decision-intelligence model for translating ethnopharmacological knowledge into AI-guided natural-product drug candidates through a conditional and auditable architecture integrating knowledge provenance, biological-source resolution, preparation context, chemical evidence, AI-guided identification, target and disease-context inference, cross-layer evidence translation, uncertainty, explicit decision gates, candidate prioritization, experimental validation, feedback, and translational assessment. The model treats ethnopharmacological knowledge as a structured discovery input rather than proof of efficacy and distinguishes historical persistence from safety, cross-cultural recurrence from mechanistic confirmation, vernacular naming from taxonomic certainty, compound occurrence from exposure, predicted targets from validated engagement, network connectivity from causality, and AI ranking from candidate validity. Its central contribution is to specify how decision objects, evidence states, uncertainty classes, progression rules, pause conditions, evidence-acquisition actions, escalation pathways, rejection rules, human oversight, and experimental feedback can be combined without inventing universal numerical thresholds. Candidate trajectories may move forward, backward, pause, undergo adjudication, be reprioritized, or terminate when failure is decisive. AI should therefore support transparent and revisable decisions rather than substitute for pharmacological validation, experimental challenge, ethical knowledge governance, or expert responsibility. The proposed architecture remains conceptual and requires prospective evaluation of its decision impact, implementation feasibility, human–AI interaction, and translational value.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

[1] Atanasov AG, Zotchev SB, Dirsch VM, Orhan IE, Banach M, Rollinger JM, et al. Natural products in drug discovery: Advances and opportunities. Nat Rev Drug Discov. 2021;20(3):200-16. doi:10.1038/s41573-020-00114-z

[2] 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

[3] 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

[4] 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

[5] Mullowney MW, Duncan KR, Elsayed SS, Garg N, van der Hooft JJJ, Martin NI, et al. Artificial intelligence for natural product drug discovery. Nat Rev Drug Discov. 2023;22(11):895-916. doi:10.1038/s41573-023-00774-7

[6] 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

[7] Yeung AWK, Heinrich M, Kijjoa A, Tzvetkov NT, Atanasov AG. The ethnopharmacological literature: An analysis of the scientific landscape. J Ethnopharmacol. 2020;250:112414. doi:10.1016/j.jep.2019.112414

[8] Raclariu AC, Heinrich M, Ichim MC, de Boer H. Benefits and limitations of DNA barcoding and metabarcoding in herbal product authentication. Phytochem Anal. 2018;29(2):123-8. doi:10.1002/pca.2732

[9] Lo YT, Shaw PC. DNA-based techniques for authentication of processed food and food supplements. Food Chem. 2018;240:767-74. doi:10.1016/j.foodchem.2017.08.022

[10] Pirintsos S, Panagiotopoulos A, Bariotakis M, Daskalakis V, Lionis C, Sourvinos G, et al. From traditional ethnopharmacology to modern natural drug discovery: A methodology discussion and specific examples. Molecules. 2022;27(13):4060. doi:10.3390/molecules27134060

[11] Saldívar-González FI, Aldas-Bulos VD, Medina-Franco JL, Plisson F. Natural product drug discovery

in the artificial intelligence era. Chem Sci. 2022;13(6):1526-46. doi:10.1039/D1SC04471K

[12] Lardos A, Aghaebrahimian A, Koroleva A, Sidorova J, Wolfram E, Anisimova M, et al. Computational literature-based discovery for natural products research: Current state and future prospects. Front Bioinform. 2022;2:827207. doi:10.3389/fbinf.2022.827207

[13] Yang K, Swanson K, Jin W, Coley CW, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370-88. doi:10.1021/acs.jcim.9b00237

[14] Mayr A, Klambauer G, Unterthiner T, Steijaert M, Wegner JK, Ceulemans H, et al. Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. Chem Sci. 2018;9(24):5441-51. doi:10.1039/C8SC00148K

[15] Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, et al. A deep learning approach to antibiotic discovery. Cell. 2020;180(4):688-702.e13. doi:10.1016/j.cell.2020.01.021

[16] Mervin LH, Johansson S, Semenova E, Giblin KA, Engkvist O. Uncertainty quantification in drug design. Drug Discov Today. 2021;26(2):474-89. doi:10.1016/j.drudis.2020.11.027

[17] Hirschfeld L, Swanson K, Yang K, Barzilay R, Coley CW. Uncertainty quantification using neural networks for molecular property prediction. J Chem Inf Model. 2020;60(8):3770-80. doi:10.1021/acs.jcim.0c00502

[18] Yu J, Wang D, Zheng M. Uncertainty quantification: Can we trust artificial intelligence in drug discovery? iScience. 2022;25(8):104814. doi:10.1016/j.isci.2022.104814

[19] Scalia G, Grambow CA, Pernici B, Li YP, Green WH. Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction. J Chem Inf Model. 2020;60(6):2697-717. doi:10.1021/acs.jcim.9b00975

[20] Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 2: A discussion of chemical and biological data. Drug Discov Today. 2021;26(4):1040-52. doi:10.1016/j.drudis.2020.11.037

[21] Reker D. Practical considerations for active machine learning in drug discovery. Drug Discov Today Technol. 2019;32-33:73-9. doi:10.1016/j.ddtec.2020.06.001

[22] Zhavoronkov A, Ivanenkov YA, Aliper A, Veselov MS, Aladinskiy VA, Aladinskaya AV, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. Nat Biotechnol. 2019;37(9):1038-40. doi:10.1038/s41587-019-0224-x

[23] 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

[24] Gao Z, Ding P, Xu R. KG-Predict: A knowledge graph computational framework for drug repurposing. J Biomed Inform. 2022;132:104133. doi:10.1016/j.jbi.2022.104133

[25] Li P, Hua L, Ma Z, Hu W, Liu Y, Zhu J. Conformalized graph learning for molecular ADMET property prediction and reliable uncertainty quantification. J Chem Inf Model. 2024;64(23):8705-17. doi:10.1021/acs.jcim.4c01139

[26] Ekins S, Puhl AC, Zorn KM, Lane TR, Russo DP, Klein JJ, et al. Exploiting machine learning for end-to-end drug discovery and development. Nat Mater. 2019;18(5):435-41. doi:10.1038/s41563-019-0338-z

[27] Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 1: Ways to make an impact, and why we are not there yet. Drug Discov Today. 2021;26(2):511-24. doi:10.1016/j.drudis.2020.12.009

[28] Ghislat G, Hernandez-Hernandez S, Piyawajanusorn C, Ballester PJ. Data-centric challenges with the application and adoption of artificial intelligence for drug discovery. Expert Opin Drug Discov. 2024;19(11):1297-307. doi:10.1080/17460441.2024.2403639

[29] Hasselgren C, Oprea TI. Artificial intelligence for drug discovery: Are we there yet? Annu Rev Pharmacol Toxicol. 2024;64:527-50. doi:10.1146/annurev-pharmtox-040323-040828

[30] 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

[31] Carroll SR, Garba I, Figueroa-Rodríguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE principles for Indigenous data governance. Data Sci J. 2020;19:43. doi:10.5334/dsj-2020-043