Reference Architectures for AI-Enabled Phytopharmaceutical Platforms: Data Layers, Model Layers, Evidence Layers, and Decision-Support Interfaces

James O'Leary[1]*, Aisling Dunne[1], Sean O'Brien[2]

[1]Department of Artificial Intelligence for CNS-Active Phytochemicals, School of Pharmacy, Trinity College Dublin, Dublin, Ireland.

[2]Department of Machine Learning for Nutraceuticals, School of Pharmacy, University College Cork, Cork, Ireland.

ABSTRACT

AI-enabled phytopharmaceutical platforms are emerging as computational infrastructures for organizing natural product data, interpreting pharmacological mechanisms, screening safety-relevant signals, and supporting translational research prioritization. However, many current computational workflows remain fragmented across databases, predictive models, network analyses, literature-mining tools, and dashboard outputs, making it difficult to distinguish curated evidence, predicted outputs, validation status, and decision-support boundaries. This article proposes an original reference architecture for AI-enabled phytopharmaceutical platforms. The architecture organizes platform design across data acquisition, ingestion, curation, provenance tracking, knowledge integration, model execution, evidence grading, explainability, validation gates, governance controls, expert review, and decision-support interfaces. Its data layer focus includes botanical identity, preparation context, phytochemical composition, bioactivity evidence, target and pathway data, omics information, ADMET and toxicity data, herb–drug interaction signals, clinical evidence where available, literature evidence, and metadata provenance. Its model and evidence layer focus separates machine learning, natural language processing, knowledge graphs, network pharmacology, QSAR or ADMET prediction, target prediction, pathway inference, safety screening, uncertainty communication, and validation status. Its interface focus emphasizes research prioritization, mechanism-evidence review, safety-alert triage, translational readiness assessment, and validation planning without replacing expert judgment. The main contribution is a modular, reusable, governance-aware architecture for designing transparent, auditable, explainable, and validation-aware phytopharmaceutical research platforms.

Key Words: *Phytopharmaceuticals, Reference architecture, Artificial intelligence, Natural products, Decision support, Evidence integration*

eIJPPR 2025; 15(4):53-65

HOW TO CITE THIS ARTICLE: O'Leary J, Dunne A, O'Brien S. Reference Architectures for AI-Enabled Phytopharmaceutical Platforms: Data Layers, Model Layers, Evidence Layers, and Decision-Support Interfaces. Int J Pharm Phytopharmacol Res. 2025;15(4):53-65. https://doi.org/10.51847/UUAOtzPFCj

INTRODUCTION

AI-enabled drug discovery has shifted computational pharmacology from isolated predictive tasks toward integrated workflows that connect chemical information, biological evidence, model outputs, and downstream validation requirements. In natural product research, this shift is especially important because phytopharmaceutical candidates often arise from chemically complex materials, heterogeneous biological evidence, traditional-use contexts, incomplete target information, and safety uncertainties that cannot be resolved by a single model or database query [1-3].

A reference architecture is needed because the central problem is not only whether individual computational tools

Corresponding author: James O'Leary
Address: Department of Artificial Intelligence for CNS-Active Phytochemicals, School of Pharmacy, Trinity College Dublin, Dublin, Ireland.
E-mail: ✉ james.oleary@tcd.ie
Received: 03 April 2025; **Revised:** 14 August 2025; **Accepted:** 16 August 2025

can predict a target, classify an ADMET property, extract literature evidence, or rank a compound. The larger design challenge is how raw data, curated records, linked biomedical entities, model-ready features, predicted outputs, validated evidence, expert-reviewed evidence, and decision-support displays should be organized into a reusable platform structure. Knowledge graphs provide one important architectural logic for connecting compounds, targets, diseases, pathways, evidence records, and hypotheses in drug discovery contexts [4].

For phytopharmaceutical research, platform architecture must also preserve distinctions that are often blurred in computational studies. Botanical identity is not equivalent to resolved chemical composition; predicted targets are not confirmed mechanisms; network proximity is not proof of pharmacological causality; ADMET prediction is not toxicological validation; and evidence aggregation is not clinical actionability. A platform that fails to encode these distinctions can make computational outputs appear more mature than the supporting evidence allows.

The aim of this article is to develop an original reference architecture for AI-enabled phytopharmaceutical platforms. The proposed architecture is intended as a conceptual, reusable system design pattern rather than a finished software implementation, deployed platform, clinical decision-support tool, or regulatory framework. It clarifies how data layers, model layers, evidence layers, validation gates, governance controls, expert review, and decision-support interfaces can be arranged to support transparent research prioritization, mechanism interpretation, safety assessment, translational readiness assessment, and validation planning.

AI-enabled phytopharmaceutical platforms

An AI-enabled phytopharmaceutical platform can be defined as an integrated computational environment that connects phytochemical data, botanical identity, biological evidence, predictive models, safety-relevant information, pathway interpretation, literature evidence, and user-facing decision-support workflows. Existing traditional medicine and herbal informatics resources show the value of structuring herb, ingredient, target, disease, symptom, and evidence relationships, while also illustrating the need for careful interpretation of database coverage, evidence provenance, and model boundaries [5-7].

Such a platform differs from a standalone natural product database because it does not merely store records; it organizes data into computationally usable, evidence-aware, and interface-ready layers. It also differs from a single machine-learning model because model outputs are only one component within a broader architecture that includes data quality control, provenance tracking, uncertainty communication, validation gates, and expert review. Multimodal natural product AI and knowledge-graph approaches support this broader platform view by emphasizing the need to connect chemical, biological, textual, and relational evidence rather than treating predictions as isolated results [8].

The platform concept is also distinct from a conventional docking pipeline or network pharmacology analysis. Docking, target prediction, network inference, and pathway mapping may be useful model-layer components, but a reference architecture must define where these outputs enter the evidence system, how their assumptions are documented, how their uncertainty is communicated, and what validation steps are required before interpretation. In this sense, the platform is not a model container; it is a structured environment for converting heterogeneous evidence into traceable, reviewable, and decision-bounded research outputs.

Modularity is therefore a core design principle. A phytopharmaceutical platform should allow data acquisition, curation, knowledge integration, model execution, evidence grading, explainability, validation, governance, and interface modules to evolve without collapsing into a single opaque workflow. This modularity supports auditing, updating, source replacement, model comparison, expert review, and context-specific implementation. It also helps prevent a common failure mode in AI-enabled pharmacology: treating computational sophistication as a substitute for pharmacological plausibility, experimental validation, or clinical judgment.

Data layer requirements

The data acquisition layer should begin with explicit source selection rather than indiscriminate aggregation. Natural product data are distributed across heterogeneous resources with differences in chemical coverage, taxonomic scope, metadata completeness, update practices, and curation depth, so platform ingestion should evaluate whether each source is suitable for the intended architecture function [9]. For phytopharmaceutical systems, this means distinguishing sources used for botanical identity, chemical structures, bioactivity, targets, pathways, omics, safety, interactions, clinical evidence, and literature evidence.

The chemical and phytochemical data layer should capture compound names, structures, identifiers, source organisms, preparation context where available, and versioned links to curated records. Curated natural product repositories illustrate the importance of chemical identity, source metadata, and update-aware database design for platform-level reuse [10]. In phytopharmaceutical contexts, chemical records should not be treated as complete representations of a plant material unless the

preparation, extraction method, plant part, batch context, and measured constituents are also available.

The pharmacological interoperability layer should connect phytochemical records with target, drug, interaction, and pharmacology information. Drug-centered resources support the need for structured links among compounds, targets, pharmacological actions, and interaction information, although phytopharmaceutical platforms must avoid assuming that drug database evidence automatically applies to complex botanical preparations [11]. Target-prediction resources further demonstrate why platforms must label predicted targets separately from experimentally supported targets, because computational target identification can support prioritization but does not establish mechanism confirmation [12].

The bioactivity, omics, and pathway layers should preserve assay context, target annotation, biological system, pathway interpretation, and biosynthetic or molecular-origin information. Curated bioactivity resources support the need to retain compound–assay–target relationships rather than reducing bioactivity to a generic activity label [13]. Biosynthetic gene-cluster repositories show how omics and pathway information can provide source and biosynthetic context for natural product interpretation, but such data should be connected through provenance-aware architecture rather than presented as direct pharmacological evidence unless supported by appropriate validation [14].

Table 1 summarises the core data layers required for AI-enabled phytopharmaceutical platforms.

Table 1. Core Data Layers for AI-Enabled Phytopharmaceutical Platforms: Botanical Identity, Phytochemical Composition, Bioactivity Evidence, Target and Pathway Data, Omics Data, ADMET and Safety Data, Herb–Drug Interaction Data, Literature Evidence, and Provenance Metadata

Data layer	Core content	Primary data sources	Curation requirement	Quality-control concern	Provenance requirement	Platform function	Failure risk if absent
Botanical identity data	Scientific name, synonym status, plant part, taxonomic context, voucher or source information where available	Botanical records, herbal databases, literature reports, curated phytomedicine resources	Taxonomic normalization and synonym control	Misidentified plant source or ambiguous naming	Source record, curator note, version date	Anchors all downstream phytopharmaceutical interpretation	Chemical, safety, or mechanism evidence may be linked to the wrong botanical entity
Preparation and extraction data	Processing method, solvent, extraction context, formulation type, batch context where available	Experimental articles, phytomedicine databases, formulation records	Standardized preparation descriptors	Incomplete preparation metadata	Original source and extraction-method trace	Connects botanical material to actual exposure context	Platform may treat non-equivalent preparations as interchangeable
Phytochemical composition data	Measured constituents, predicted constituents, structures, identifiers, relative abundance where available	Natural product databases, analytical studies, curated chemical repositories	Identifier mapping and structure normalization	Duplicate names, unresolved stereochemistry, missing concentration context	Compound source, measurement status, version record	Supports compound-level modeling and composition review	Model-ready data may misrepresent actual phytochemical exposure
Bioactivity evidence data	Assays, biological systems, targets, activity labels, experimental conditions	Bioactivity databases, pharmacology literature	Assay-context preservation	Assay heterogeneity and non-comparable endpoints	Assay source and curation trail	Links compounds or extracts to biological effects	Activity claims may become detached from experimental context
Target and pathway data	Protein targets, genes, disease links, pathways, biological processes	Target databases, pathway resources, computational	Entity normalization and relation typing	Predicted and confirmed targets may be mixed	Relation source and evidence type	Supports mechanism-hypothesis generation	Predicted links may be misread as validated mechanisms

		target tools, literature					
Omics and biosynthetic data	Transcriptomic, proteomic, metabolomic, microbiome, and biosynthetic information where available	Omics studies, biosynthetic repositories, pathway datasets	Cross-omics mapping and metadata harmonization	Platform may overconnect noisy or context-specific omics signals	Dataset origin, organism, condition, method	Supports systems-level interpretation	Pathway claims may be biologically under-specified
ADMET and toxicity data	Absorption, distribution, metabolism, excretion, toxicity, alerts, safety evidence	ADMET predictors, toxicology studies, pharmacology databases, safety literature	Evidence-type labeling	Prediction may be confused with toxicological validation	Model or study source and applicability note	Supports safety screening and alert triage	Safety uncertainty may be hidden from users
Herb–drug interaction data	Interaction mechanisms, co-exposure concerns, metabolic enzymes, transporters, evidence reports	Drug databases, interaction literature, pharmacology studies	Mechanism and evidence classification	Anecdotal, predicted, and validated interactions may be conflated	Interaction source and evidence status	Supports safety-alert interface design	Users may miss interaction uncertainty or overstate weak alerts
Literature evidence data	Claims, study type, methods, outcomes, limitations, extracted entities	Peer-reviewed literature, structured abstracts, text-mining outputs	Human-verifiable extraction and claim mapping	NLP extraction errors and claim drift	Article source, extraction method, reviewer status	Feeds evidence grading and mechanism review	Literature-derived claims may lose context
Provenance metadata	Source, timestamp, transformation history, curator, model version, evidence status	All platform layers	Persistent provenance tagging	Broken lineage or undocumented transformations	Complete audit trail	Enables transparency, auditability, and evidence review	Platform outputs become difficult to verify or reproduce

Model and evidence layers

The model layer should be designed as a set of task-specific computational modules rather than a single intelligence layer. In drug design, AI can assist representation learning, compound prioritization, molecular design, and hypothesis generation, but the architecture must place these functions within an evidence system that distinguishes prediction, plausibility, feasibility, and validation [15]. For phytopharmaceutical platforms, this means that each model output should carry information about input data, intended task, applicability domain, uncertainty, and required follow-up.

Deep learning modules can support molecular representation learning, structure–activity modeling, feature extraction, and other discovery tasks, but their role depends on the quality and relevance of the data used for training and inference [16]. Generative molecular design modules may propose candidate structures or analogues, yet such outputs should be routed through feasibility, developability, safety, and pharmacological plausibility gates before being treated as research priorities [17]. In a reference architecture, generated or predicted outputs are therefore intermediate artifacts, not evidence of therapeutic value.

Natural language processing can support evidence extraction from biomedical literature, including entity recognition, relation extraction, and retrieval of relevant claims. Biomedical language models demonstrate the feasibility of text-mining support for literature evidence workflows, but extracted information still requires traceability, human verification, and evidence grading before it can support mechanism interpretation or decision-support displays [18]. This is especially important in phytopharmaceutical research, where claims may involve extracts, preparations, mixtures, isolated compounds, animal models, in vitro assays, observational evidence, or traditional-use contexts that are not interchangeable.

Target prediction, interaction prediction, ADMET prediction, pathway inference, and safety screening should feed an evidence layer rather than bypass it. Drug–target interaction models can support prioritization of compound–protein hypotheses, but their outputs require uncertainty labeling and validation gates before mechanistic interpretation [19]. ADMET prediction platforms can support early safety-oriented screening and developability review, but predicted ADMET properties should not be presented as definitive safety conclusions [20]. Similar ADMET screening tools reinforce the need to treat safety predictions as evidence-graded alerts that require context-specific interpretation rather than final risk determinations [21].

Table 2 maps the model and evidence layers needed to connect phytopharmaceutical data with explainable and validation-aware computational outputs.

Table 2. Model and Evidence Layers in AI-Enabled Phytopharmaceutical Platforms: Machine Learning, Natural Language Processing, Knowledge Graphs, Network Pharmacology, ADMET Prediction, Pathway Inference, Evidence Grading, Explainability, Uncertainty, and Validation Gates

Layer or module	Core function	Input requirements	Output type	Evidence interpretation role	Explainability requirement	Validation requirement	Main limitation
Machine-learning module	Predict properties, classify candidates, rank hypotheses, or learn molecular and biological representations	Model-ready chemical, biological, omics, or evidence features	Predicted labels, rankings, embeddings, candidate lists	Supports research prioritization, not proof	Task-specific explanation or transparent model documentation	Benchmarking, applicability-domain review, data leakage checks	Predictions depend on training data relevance and feature quality
Natural language processing module	Extract entities, relations, claims, and evidence fragments from literature	Full text or abstracts, controlled vocabularies, entity dictionaries	Extracted claims, entity links, relation candidates	Feeds evidence curation and literature review	Traceable source sentence or passage and extraction confidence	Human verification and extraction-quality audit	Text-mined claims may lose methodological context
Knowledge graph module	Connect compounds, herbs, targets, diseases, pathways, evidence records, and safety signals	Curated entities, relation types, provenance metadata	Linked evidence paths and graph neighborhoods	Supports mechanism navigation and hypothesis generation	Path-level explanation and relation provenance	Relation review and graph-quality assessment	Spurious or weakly supported links may appear meaningful
Network pharmacology module	Explore multi-target and pathway-level relationships	Compound–target, target–pathway, and disease association data	Network modules, candidate mechanisms, pathway associations	Supports systems-level interpretation	Display of relation type, evidence source, and model assumptions	Experimental or literature-based follow-up	Network connectivity does not establish causality
Target prediction module	Prioritize possible compound–target relationships	Chemical structures, pharmacophore features, protein features, known interactions	Candidate target list or interaction prediction	Generates target hypotheses	Input similarity, model basis, and uncertainty label	Orthogonal target validation and evidence comparison	Predicted targets may not be biologically relevant in the preparation context
Pathway inference module	Map target or omics signals to biological processes	Target lists, omics data, pathway annotations	Pathway hypotheses and process-level interpretation	Supports mechanism framing	Display of pathway source and mapping assumptions	Biological-context review and experimental follow-up	Pathway enrichment can overgeneralize from incomplete inputs
QSAR or ADMET module	Estimate developability and safety-	Chemical structures and descriptor features	Predicted ADMET or toxicity flags	Supports screening and triage	Applicability-domain and	Comparison with experimental ADMET or	Prediction is not equivalent to safety validation

	relevant molecular properties				feature-basis reporting	toxicology evidence	
Safety prediction module	Identify potential toxicity, adverse-event, or interaction concerns	Chemical, pharmacological, toxicology, and interaction data	Safety alerts or concern categories	Supports cautionary review	Alert rationale and evidence source display	Toxicological and pharmacological validation	May produce false alarms or miss preparation-specific risks
Evidence grading layer	Classify evidence by type, strength, provenance, and validation status	Literature claims, assay records, model outputs, expert notes	Evidence-grade labels and summaries	Separates prediction from validated evidence	Clear grade definitions and evidence trace	Independent review of grading rules	Grading may appear more precise than evidence allows
Explainability and uncertainty layer	Communicate why outputs were generated and how uncertain they are	Model outputs, graph paths, confidence indicators, provenance data	Explanation panels and uncertainty notes	Prevents overconfident interpretation	User-specific explanations tied to decision context	Review by intended expert users	Explanations can be persuasive even when models are wrong
Validation gate layer	Prevent premature advancement of weak outputs	Evidence grades, model outputs, expert review, provenance status	Gate-pass, gate-hold, or gate-reject status	Converts evidence maturity into workflow control	Transparent gate criteria	Gate-specific validation and audit	Overly rigid gates may block useful hypotheses; weak gates may overstate readiness
Expert review layer	Add domain judgment to computational outputs	Evidence summaries, uncertainty notes, provenance records, model explanations	Reviewed evidence status and reviewer comments	Adds contextual interpretation	Reviewer rationale should be documented	Reviewer consistency and conflict checks	Expert review can be inconsistent without structured criteria

Decision-support interfaces

Decision-support interfaces are the visible layer through which platform outputs become usable for researchers, pharmacologists, toxicologists, formulation scientists, and translational teams. In AI-enabled phytopharmaceutical platforms, these interfaces should not simply display ranked compounds, predicted targets, or pathway diagrams. They should present evidence provenance, uncertainty, validation status, and decision boundaries so that computational outputs can support research prioritization without being mistaken for clinical recommendations. Decision-support scholarship emphasizes that such systems require workflow fit, risk management, staged evaluation, and expert oversight before their outputs can responsibly influence high-consequence decisions [22-24].

A research prioritization interface should help users compare candidates by composition evidence, bioactivity support, predicted target plausibility, pathway context, safety alerts, interaction concerns, and validation readiness. A mechanism-evidence interface should show whether a proposed mechanism is based on direct assay evidence, target prediction, network inference, literature extraction, omics association, or expert review. Explainability must be matched to the user and decision context; an explanation useful for a computational scientist may not be sufficient for a pharmacologist evaluating biological plausibility or a safety expert reviewing interaction concerns [25].

A safety-alert interface should separate predicted ADMET alerts, literature-derived toxicity signals, interaction warnings, contraindication concerns, and experimentally supported safety findings. It should also communicate uncertainty in a way that discourages both false reassurance and exaggerated alarm. A translational-readiness interface should display what evidence is available, what evidence is missing, which validation gates remain unresolved, and whether the output is appropriate only for research planning. Lifecycle thinking is essential because AI-enabled health decision systems require monitoring, auditability, governance, and reassessment as data, models, users, and implementation contexts change [26].

Figure 1 illustrates how AI-enabled phytopharmaceutical decision-support interfaces can connect data evidence, model outputs, uncertainty communication, expert review, and research decision boundaries.

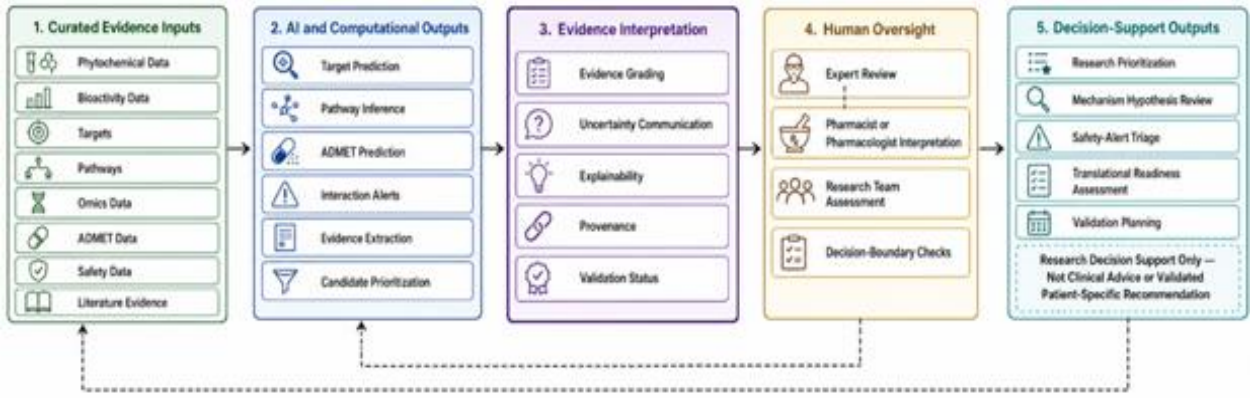

Figure 1. Decision-Support Interface Logic for AI-Enabled Phytopharmaceutical Platforms

Alt-text description

A conceptual workflow diagram showing curated phytopharmaceutical data flowing into AI model outputs, evidence grading, uncertainty indicators, safety alerts, translational readiness displays, expert review, and research decision-support boundaries.

Decision-boundary safeguards should be explicit in the interface rather than hidden in documentation. Outputs should be labeled as research prioritization, mechanism hypothesis review, safety-alert triage, translational readiness assessment, or validation planning. The interface should avoid language that implies treatment recommendation, clinical suitability, regulatory acceptability, or validated patient-specific action. It should support expert interpretation by showing why a candidate was prioritized, what evidence was used, what uncertainty remains, and what validation would be required before stronger claims could be made.

Table 3 summarises decision-support interface requirements for AI-enabled phytopharmaceutical platforms.

Table 3. Decision-Support Interface Requirements for AI-Enabled Phytopharmaceutical Platforms: Research Prioritization, Mechanism Evidence Review, Safety Alerts, Translational Readiness, Expert Review, Uncertainty Communication, and Decision-Boundary Safeguards

Interface function	Primary user	Decision question supported	Required evidence display	Uncertainty communication need	Human oversight requirement	Decision boundary	Risk if poorly designed
Research prioritization dashboard	Natural product researchers, pharmacologists, discovery teams	Which candidates should be reviewed or validated next?	Composition evidence, bioactivity evidence, predicted targets, pathway links, safety alerts, provenance	Candidate ranking should show uncertainty and missing evidence	Expert review before prioritization is acted upon	Supports research planning only	Rankings may be mistaken for validated therapeutic potential
Mechanism evidence viewer	Pharmacologists, systems biologists, computational biologists	What evidence supports a proposed mechanism?	Assay evidence, target prediction, network paths, pathway annotations, literature claims	Must distinguish direct evidence from inferred relationships	Mechanistic plausibility review	Supports hypothesis interpretation, not mechanism confirmation	Predicted targets may be presented as confirmed mechanisms
Safety-alert interface	Toxicologists, pharmacologists, formulation scientists	What safety concerns require review?	ADMET predictions, toxicity reports, interaction warnings, contraindication	Alerts should indicate evidence type and confidence limits	Safety expert interpretation	Supports triage and follow-up planning	Weak alerts may cause overreaction, while missing context may

			signals where available				cause false reassurance
Translational readiness interface	Translational researchers, project teams	What evidence is needed before further development?	Data completeness, validation status, safety evidence, mechanism support, expert notes	Must show unresolved gates and evidence gaps	Multidisciplinary review	Supports readiness assessment for research progression	Outputs may be confused with clinical or regulatory readiness
Validation planning module	Experimental pharmacologists, assay teams	Which experiments or reviews should be prioritized?	Candidate rationale, predicted mechanism, safety flags, evidence gaps	Should identify uncertainty sources driving validation need	Experimental design review	Supports validation planning only	Experiments may be planned from unsupported model outputs
Expert review workspace	Pharmacologists, botanists, toxicologists, data curators	How should computational outputs be interpreted?	Full provenance, model explanation, evidence grade, reviewer history	Should show disagreement, ambiguity, and unresolved evidence	Named or role-based expert annotation	Supports documented interpretation	Review may become informal or unauditable
Uncertainty communication panel	All expert users	How reliable is the output for its stated purpose?	Model confidence where appropriate, applicability notes, missing data, conflicting evidence	Must avoid false precision	Review of uncertainty language	Supports cautious use of outputs	Users may overtrust precise-looking displays
Decision-boundary safeguard	Platform administrators, governance teams, all users	What should the platform not be used to decide?	Output category, validation status, permitted use, restricted use	Must state limits plainly	Governance approval and periodic review	Excludes clinical advice, product recommendation, dose selection, and regulatory determination	Platform outputs may be misapplied beyond evidence maturity

Proposed reference architecture

The proposed reference architecture is a modular, layered design pattern for AI-enabled phytopharmaceutical platforms. It begins with data acquisition and curation, then moves through provenance tracking, linked-data or knowledge-graph construction, model execution, evidence interpretation, validation gates, expert review, governance controls, and decision-support interfaces. Biomedical knowledge integration has shown how heterogeneous entities can be connected to support hypothesis prioritization [27]. Precision-medicine knowledge graph work further supports the use of linked-data layers as intermediate architecture components between curated data and user-facing decision outputs [28].

The bottom layer of the architecture is the data acquisition and curation layer. It includes botanical identity, preparation context, phytochemical composition, compound identifiers, bioactivity evidence, target data, pathway data, omics data, ADMET information, toxicity evidence, herb–drug interaction data, clinical evidence where available, literature evidence, and provenance metadata. Before any data are treated as model-ready, they should pass quality checks for source reliability, entity normalization, duplicate resolution, missing metadata, evidence type, and version history. FAIR-oriented biopharma data principles support the need for quality-managed, reusable, traceable data before computational reuse [29].

The model layer is positioned above the knowledge integration layer rather than directly above raw data. This placement reflects the need for models to operate on curated, normalized, provenance-aware, and context-labeled inputs. Machine-learning applications in drug discovery support task-specific modeling across discovery and development workflows, but a reference architecture should map each model to a defined function, input requirement, output type, and validation need [30]. In phytopharmaceutical platforms, machine learning, natural language processing, network pharmacology, target prediction, pathway inference, QSAR or ADMET models, safety prediction, interaction mapping, and candidate prioritization should remain separable modules.

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

The evidence layer converts model outputs and curated records into interpretable research artifacts. It should classify evidence as raw, curated, predicted, inferred, experimentally supported, literature-derived, expert-reviewed, or validation-ready. Recent AI drug discovery scholarship supports the need to interpret model outputs with attention to uncertainty, limitations, and the difference between computational prediction and actionable evidence [31]. The architecture therefore routes each output through explainability, uncertainty communication, evidence grading, provenance display, and validation gates before it reaches a decision-support interface.

The decision-support and expert-review layers define what users can do with platform outputs. AI-assisted drug discovery should be framed as support for discovery workflows rather than autonomous proof of biological activity, safety, or translational readiness [32]. Where outputs may influence high-stakes prioritization, the architecture should prefer interpretable or auditable reasoning pathways and should make decision boundaries visible to users [33]. Expert review is not an optional decoration in this architecture; it is the layer that checks whether computational outputs are biologically plausible, safety-aware, evidence-consistent, and appropriate for the stated research decision.

Table 4 presents the proposed reference architecture for AI-enabled phytopharmaceutical platforms.

Table 4. Proposed Reference Architecture for AI-Enabled Phytopharmaceutical Platforms: Data Layers, Model Layers, Evidence Layers, Validation Gates, Governance Controls, Expert Review, and Decision-Support Interfaces

Architecture layer	Core purpose	Required inputs	Processing function	Output generated	Governance requirement	Validation gate	Implementation risk	Platform value
Data acquisition layer	Identify and select relevant data sources	Botanical, chemical, biological, safety, interaction, clinical, and literature sources	Source screening and eligibility assessment	Source inventory	Source-quality policy	Source fitness check	Low-quality or irrelevant sources enter the platform	Creates transparent data foundation
Data ingestion and curation layer	Transform heterogeneous inputs into curated records	Raw database records, literature records, experimental records	Normalization, deduplication, annotation, entity mapping	Curated platform records	Curator accountability and version control	Curation completeness check	Entity mismatch or metadata loss	Produces reliable model-ready records
Data provenance layer	Track origin and transformation history	Source identifiers, timestamps, extraction methods, curator notes	Provenance tagging and lineage recording	Traceable evidence records	Audit trail and versioning	Provenance completeness check	Claims become unverifiable	Enables reproducibility and review
Knowledge integration layer	Link entities and evidence across domains	Curated compounds, plants, targets, pathways, diseases, evidence, safety signals	Linked-data construction, semantic mapping, knowledge graph assembly	Evidence-linked graph structure	Relation governance and access control	Relation-support check	Unsupported links appear meaningful	Supports mechanism navigation and hypothesis generation
Model layer	Generate task-specific computational outputs	Model-ready curated data and graph features	Machine learning, NLP, network pharmacology, target prediction,	Predictions, rankings, extracted claims, inferred paths	Model registry and monitoring	Model applicability check	Predictions are treated as proof	Supports scalable research prioritization

			pathway inference, ADMET prediction					
Evidence grading layer	Classify evidence maturity and strength	Curated evidence, model outputs, literature claims, assay records	Evidence labeling and grading	Evidence-grade summaries	Transparent grading rules	Evidence-grade review	Weak evidence is overstated	Separates prediction from validated evidence
Explainability and uncertainty layer	Make outputs interpretable and cautious	Model outputs, graph paths, uncertainty indicators, provenance	Explanation generation and uncertainty display	Explanation panels and uncertainty notes	Explanation-use safeguards	Explanation adequacy check	Persuasive explanations create false confidence	Improves expert interpretability
Validation gate layer	Control movement between hypothesis and stronger claims	Evidence grades, expert comments, model outputs, provenance records	Gate logic for composition, target, mechanism, safety, and readiness	Gate-pass, gate-hold, or gate-reject status	Gate audit and change tracking	Gate-specific validation requirement	Premature advancement of weak outputs	Prevents overclaiming
Expert review layer	Add human pharmacological, botanical, toxicological, and translational judgment	Evidence summaries, uncertainty notes, safety alerts, mechanism hypotheses	Expert annotation and adjudication	Expert-reviewed evidence status	Reviewer role control and auditability	Expert review completion check	Informal review becomes inconsistent	Adds contextual scientific judgment
Decision-support interface layer	Present bounded outputs to users	Evidence-graded and expert-reviewed outputs	Dashboard display, query, alerting, readiness review	Research prioritization, mechanism review, safety triage, validation planning	Role-based access and use-boundary controls	Decision-boundary check	Users misapply outputs as clinical recommendations	Converts architecture outputs into usable research support
Governance and audit layer	Control lifecycle, security, accountability, and monitoring	Logs, model versions, data versions, access records, review records	Monitoring, audit reporting, access management, lifecycle review	Governance record	Security, access control, auditability, human oversight	Governance audit	Drift, unauthorized access, or undocumented changes	Maintains trustworthiness and accountability

Figure 2 presents the proposed reference architecture for AI-enabled phytopharmaceutical platforms across data layers, model layers, evidence layers, governance controls, and decision-support interfaces.

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

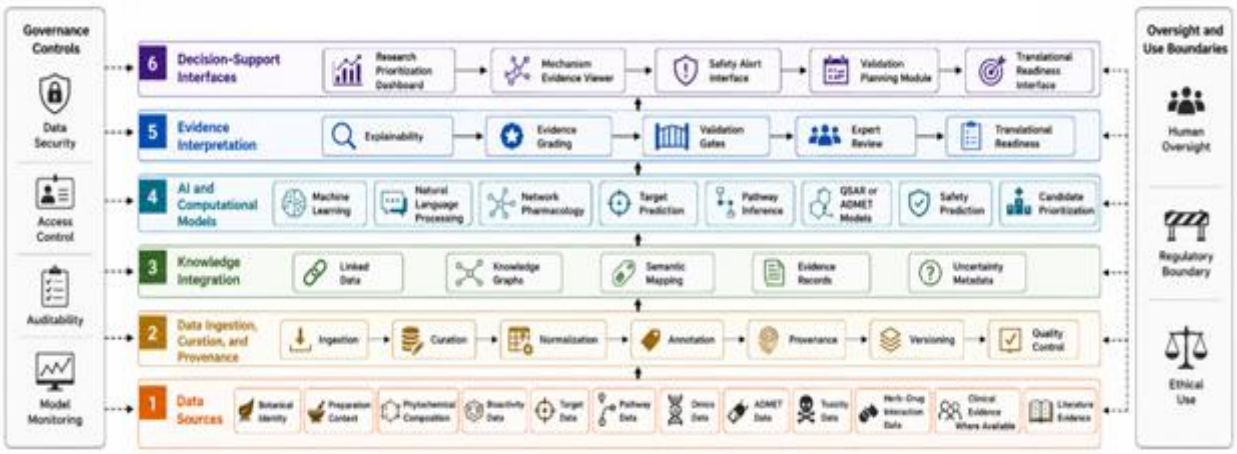

Figure 2. Reference Architecture for AI-Enabled Phytopharmaceutical Platforms

Alt-text description

A layered architecture diagram showing phytopharmaceutical data sources feeding into data curation, provenance tracking, knowledge integration, AI model layers, evidence grading, explainability, uncertainty, validation gates, governance, expert review, and decision-support interfaces.

Implementation and governance implications

Implementation of the proposed reference architecture requires more than technical integration. It requires governance controls for data quality, semantic interoperability, access management, audit trails, model versioning, human oversight, and responsible AI use. Trustworthy AI scholarship emphasizes that responsible systems require accountability, transparency, robustness, fairness, human oversight, and governance mechanisms that connect principles to operational controls [34]. In phytopharmaceutical platforms, these controls should apply to data ingestion, model deployment, evidence grading, interface design, and user permissions.

Explainability must also be governed carefully. Explanations can help users understand why a model or graph produced an output, but explanation displays should not be interpreted as proof that the output is correct, safe, clinically relevant, or pharmacologically validated. Critical evaluations of explainable AI in healthcare warn that explanation methods can create false confidence if they are treated as substitutes for rigorous validation, uncertainty assessment, and domain-specific review [35]. A phytopharmaceutical platform should therefore pair explanation panels with provenance, uncertainty, evidence grade, and validation status.

Future implementation work should focus on benchmark datasets, transparent reporting templates, evidence-grade definitions, validation-gate criteria, user-centered interface studies, and governance frameworks tailored to phytopharmaceutical research. Important unresolved challenges include incomplete botanical metadata, preparation-dependent chemical variability, heterogeneous assay evidence, sparse safety data, inconsistent interaction evidence, limited external validation of predictive models, and the need to preserve clinical and regulatory decision boundaries. The architecture proposed here is intended to guide platform design and evaluation, not to claim that any specific implementation is ready for deployment or clinical use.

CONCLUSION

This article proposed a reference architecture for AI-enabled phytopharmaceutical platforms that integrates data layers, model layers, evidence layers, validation gates, governance controls, expert review, and decision-support interfaces. The architecture clarifies how botanical identity, phytochemical composition, bioactivity evidence, target and pathway data, omics information, safety data, interaction evidence, literature evidence, model outputs, explainability, uncertainty, and provenance can be organized into a transparent and auditable platform structure. AI-enabled phytopharmaceutical platforms can support natural product research when their outputs are evidence-graded, explainable, provenance-aware, validation-aware, and bounded by expert judgment. They should not be treated as substitutes for experimental validation, clinical evaluation, regulatory review, or professional decision-making.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

1. Paul D, Sanap G, Shenoy S, Kalyane D, Kalia K, Tekade RK. Artificial intelligence in drug discovery and development. Drug Discov Today. 2021;26(1):80-93. doi:10.1016/j.drudis.2020.10.010
2. 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
3. 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
4. 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
5. Zhang R, Zhu X, Bai H, Ning K. Network pharmacology databases for traditional Chinese medicine: Review and assessment. Front Pharmacol. 2019;10:123. doi:10.3389/fphar.2019.00123
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). doi:10.1093/nar/gkaa1063
7. Wu Y, Zhang F, Yang K, Fang S, Bu D, Li H, et al. SymMap: An integrative database of traditional Chinese medicine enhanced by symptom mapping. Nucleic Acids Res. 2019;47(D1). doi:10.1093/nar/gky1021
8. Meijer D, Beniddir MA, Coley CW, Mejri YM, Ozturk 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
9. Sorokina M, Steinbeck C. Review on natural products databases: Where to find data in 2020. J Cheminform. 2020;12(1):20. doi:10.1186/s13321-020-00424-9
10. van Santen JA, Poynton EF, Iskakova D, McMann E, Alsup TA, Clark TN, et al. The Natural Products Atlas 2.0: A database of microbially-derived natural products. Nucleic Acids Res. 2022;50(D1). doi:10.1093/nar/gkab941
11. Wishart DS, Feunang YD, Guo AC, Lo EJ, Marcu A, Grant JR, et al. DrugBank 5.0: A major update to the DrugBank database for 2018. Nucleic Acids Res. 2018;46(D1). doi:10.1093/nar/gkx1037
12. Wang X, Shen Y, Wang S, Li S, Zhang W, Liu X, et al. PharmMapper 2017 update: A web server for potential drug target identification with a comprehensive target pharmacophore database. Nucleic Acids Res. 2017;45(W1). doi:10.1093/nar/gkx374
13. Gaulton A, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, et al. The ChEMBL database in 2017. Nucleic Acids Res. 2017;45(D1). doi:10.1093/nar/gkw1074
14. Kautsar SA, Blin K, Shaw S, Navarro-Muñoz JC, Terlouw BR, van der Hooft JJJ, et al. MIBiG 2.0: A repository for biosynthetic gene clusters of known function. Nucleic Acids Res. 2020;48(D1). doi:10.1093/nar/gkz882
15. Schneider P, Walters WP, Plowright AT, Sieroka N, Listgarten J, Goodnow RA, et al. Rethinking drug design in the artificial intelligence era. Nat Rev Drug Discov. 2020;19(5):353-64. doi:10.1038/s41573-019-0050-3
16. Chen H, Engkvist O, Wang Y, Olivecrona M, Blaschke T. The rise of deep learning in drug discovery. Drug Discov Today. 2018;23(6):1241-50. doi:10.1016/j.drudis.2018.01.039
17. Elton DC, Boukouvalas Z, Fuge MD, Chung PW. Deep learning for molecular design—a review of the state of the art. Mol Syst Des Eng. 2019;4(4):828-49. doi:10.1039/C9ME00039A
18. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: A pre-trained biomedical language representation model for biomedical text mining. Bioinformatics. 2020;36(4):1234-40. doi:10.1093/bioinformatics/btz682
19. Huang K, Fu T, Glass LM, Zitnik M, Xiao C, Sun J. DeepPurpose: A deep learning library for drug–target interaction prediction. Bioinformatics. 2020;36(22-23):5545-7. doi:10.1093/bioinformatics/btaa1005
20. Xiong G, Wu Z, Yi J, Fu L, Yang Z, Hsieh C, et al. ADMETlab 2.0: An integrated online platform for accurate and comprehensive predictions of ADMET properties. Nucleic Acids Res. 2021;49(W1). doi:10.1093/nar/gkab255
21. Yang H, Lou C, Sun L, Li J, Cai Y, Wang Z, et al. admetSAR 2.0: Web-service for prediction and optimization of chemical ADMET properties. Bioinformatics. 2019;35(6):1067-9. doi:10.1093/bioinformatics/bty707
22. Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: Benefits, risks, and strategies for success. NPJ Digit Med. 2020;3:17. doi:10.1038/s41746-020-0221-y
23. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC

Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2
24. Park Y, Jackson GP, Foreman MA, Gruen D, Hu J, Das AK. Evaluating artificial intelligence in medicine: Phases of clinical research. JAMIA Open. 2020;3(3):326-31. doi:10.1093/jamiaopen/ooaa033
25. Amann J, Blasimme A, Vayena E, Frey D, Madai VI. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. doi:10.1186/s12911-020-01332-6
26. Ng MY, Kapur S, Blizinsky KD, Hernandez-Boussard T. The AI life cycle: A holistic approach to creating ethical AI for health decisions. Nat Med. 2022;28(11):2247-9. doi:10.1038/s41591-022-01993-y
27. Himmelstein DS, Lizee A, Hessler C, Brueggeman L, Chen SL, Hadley D, et al. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. Elife. 2017;6. doi:10.7554/eLife.26726
28. 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
29. Harrow I, Balakrishnan R, Küçük McGinty H, Plasterer T, Romacker M. Maximizing data value for biopharma through FAIR and quality implementation: FAIR plus Q. Drug Discov Today. 2022;27(5):1441-7. doi:10.1016/j.drudis.2022.01.006
30. 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
31. Jiménez-Luna J, Grisoni F, Weskamp N, Schneider G. Artificial intelligence in drug discovery: Recent advances and future perspectives. Expert Opin Drug Discov. 2021;16(9):949-59. doi:10.1080/17460441.2021.1909567
32. Cerchia C, Lavecchia A. New avenues in artificial-intelligence-assisted drug discovery. Drug Discov Today. 2023;28(4):103516. doi:10.1016/j.drudis.2023.103516
33. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
34. Díaz-Rodríguez N, Del Ser J, Coeckelbergh M, López de Prado M, Herrera-Viedma E, Herrera F. Connecting the dots in trustworthy artificial intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. Inf Fusion. 2023;99:101896. doi:10.1016/j.inffus.2023.101896
35. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9