Future-Ready Pharmaceutical Sciences and Phytopharmacology: From Molecular Intelligence to Responsible Natural Product-Based Therapeutics

David Thompson[1]*, Sarah Mitchell[1], Rachel Adams[2], James Brown[3]

[1]Department of AI for Nutraceutical-Gut Microbiome Interactions, Faculty of Pharmacy, University of Toronto, Toronto, Canada.
[2]Department of Machine Learning for Bioactive Metabolite Identification, Faculty of Pharmaceutical Sciences, University of British Columbia, Vancouver, Canada.
[3]Department of Cheminformatics for Prebiotic Phytochemicals, Faculty of Pharmacy, McGill University, Montreal, Canada.

ABSTRACT

Pharmaceutical sciences and phytopharmacology are being reshaped by artificial intelligence, computational pharmacology, multimodal evidence integration, and increasingly connected molecular data. However, discovery acceleration alone cannot establish mechanistic validity, safety, clinical utility, sustainability, or regulatory readiness. This article proposes a future research agenda for integrating AI-era pharmaceutical sciences with responsible natural product-based therapeutic research. The agenda is organized around high-integrity data ecosystems, molecular intelligence, natural product identity and characterization, multimodal discovery methods, mechanistic validation, safety-aware design, exposure and response modeling, translational evidence, regulatory science, sustainability, ethical sourcing, benefit-sharing, and accountable governance. It emphasizes benchmarking, reproducibility, uncertainty estimation, bias assessment, model documentation, external validation, and multidisciplinary human oversight as foundational research capabilities rather than optional additions. Molecular intelligence is positioned as a means of connecting chemical structures, phytochemical profiles, assays, omics, textual knowledge, and supported clinical signals while maintaining provenance and evidence-level distinctions. Responsible translation is framed as a stage-gated research process requiring empirical validation, safety evaluation, exposure feasibility, clinical evidence where relevant, and regulatory assessment where applicable. Strategic priorities include interoperable infrastructures, natural product knowledge graphs, validation pipelines, federated collaboration, causal and digital-twin research where justified, sustainability safeguards, and workforce development. The principal contribution is an evidence-bounded agenda that links discovery innovation with validation and governance while maintaining a clear separation between research prioritization and therapeutic decision-making.

Key Words: *Pharmaceutical sciences, Phytopharmacology, Molecular intelligence, Natural product therapeutics, AI drug discovery, Responsible translation*

eIJPPR 2026; 16(3):80-99

HOW TO CITE THIS ARTICLE: Thompson D, Mitchell S, Adams R, Brown J. Future-Ready Pharmaceutical Sciences and Phytopharmacology: From Molecular Intelligence to Responsible Natural Product-Based Therapeutics. Int J Pharm Phytopharmacol Res. 2026;16(3):80-99. https://doi.org/10.51847/CHWXTuYMdO

INTRODUCTION

Pharmaceutical sciences are entering a period in which artificial intelligence can influence molecular representation, target identification, property prediction, compound prioritization, medicinal chemistry, preclinical research, and development planning. These capabilities create opportunities to connect previously fragmented chemical and biological evidence, but their scientific value depends on how well computational outputs are integrated

Corresponding author: David Thompson
Address: Department of AI for Nutraceutical-Gut Microbiome Interactions, Faculty of Pharmacy, University of Toronto, Toronto, Canada.
E-mail: ✉ david.thompson@utoronto.ca
Received: 02 March 2026; **Revised:** 03 June 2026; **Accepted:** 08 June 2026

with pharmaceutical expertise, experimental pharmacology, and medicinal-chemistry judgement. Future-ready research should therefore treat AI as part of a multidisciplinary evidence system rather than as an autonomous discovery mechanism or a substitute for domain specialists [1, 2].

The expansion of AI-enabled discovery also changes the requirements for research design. Faster prediction or candidate generation does not resolve deficiencies in material identity, dataset quality, assay relevance, model generalizability, mechanistic validation, exposure feasibility, toxicology, or clinical evidence. A model may produce a technically plausible ranking while remaining unsuitable for a pharmaceutical decision because its training data, evaluation design, applicability domain, uncertainty, or intended use is insufficiently characterized. The central challenge is consequently not only to increase computational capability but also to build research systems in which model outputs are traceable, testable, reproducible, and proportionate to the claims being considered [3].

This challenge is particularly important in phytopharmacology and natural product discovery. Natural products have made substantial contributions to pharmaceutical innovation and continue to offer structurally and biologically diverse starting points for research. Nevertheless, natural origin does not establish chemical consistency, pharmacological efficacy, clinical utility, safety, sustainability, or ethical legitimacy. Extracts, fractions, purified constituents, analogues, and multicomponent preparations require distinct evidentiary treatment, while botanical authentication, chemical characterization, dose, exposure, metabolism, and source provenance can materially affect interpretation [4].

The objective of this article is to propose an original future research agenda that connects AI-era pharmaceutical sciences, molecular intelligence, natural product discovery, responsible therapeutic translation, regulatory science, sustainability, ethical sourcing, benefit-sharing, and accountable governance. Rather than presenting a generic catalogue of technologies, the agenda specifies the capabilities, validation gates, research priorities, feedback mechanisms, and claim boundaries needed to organize future work. It is intended to support research planning and interdisciplinary coordination; it does not constitute an implemented framework, an accepted standard, a clinical guideline, a regulatory pathway, or evidence that any natural product-based intervention is safe or therapeutically effective.

Pharmaceutical sciences in the AI era

AI-era pharmaceutical sciences require coordinated capabilities that extend beyond model selection. Predictive and generative methods may assist with target assessment, molecular design, property estimation, formulation questions, safety prioritization, and development planning, but their contribution depends on whether they address a clearly defined pharmaceutical problem and outperform appropriate baselines under realistic conditions. Research programs should begin by specifying the intended scientific function, relevant decision context, acceptable uncertainty, required evidence, and conditions under which a computational output should be revised, rejected, or escalated for expert assessment [5].

Molecular representation learning and deep-learning methods can identify complex patterns across chemical structures and biological measurements, including relationships that may be difficult to encode through manually selected descriptors. Their use nevertheless introduces research questions concerning chemical-domain coverage, stereochemical representation, endpoint reliability, data sparsity, task transferability, and interpretability. A future-ready research environment should therefore preserve the relationship between the model input, the underlying pharmaceutical entity, the experimental context, and the intended prediction rather than treating learned representations as context-independent molecular knowledge [6].

Data ecosystems are the operational foundation of this environment. Pharmaceutical datasets frequently combine records produced for different purposes, under different assay conditions, with different identifiers, quality controls, access restrictions, and degrees of missingness. Increasing dataset size without resolving duplication, provenance loss, label inconsistency, assay incompatibility, temporal drift, or selective reporting may scale error rather than knowledge. High-integrity infrastructures should support entity resolution, metadata standards, version control, lineage tracking, quality grading, access governance, leakage prevention, and documentation of excluded or uncertain records [7].

Explainability can help scientists interrogate model behaviour, identify potentially influential features, examine unexpected predictions, and formulate follow-up experiments, but an explanation is not proof that a model has learned a valid biological mechanism [8]. Transparent and responsible pharmaceutical AI also requires reproducibility, predefined benchmarks, uncertainty estimation, bias assessment, external validation, model documentation, multidisciplinary oversight, and explicit evaluation of whether a result is effective for its intended research purpose. These requirements should be considered throughout study design rather than added after model development, and neither technical performance nor transparent reporting alone establishes clinical benefit or implementation readiness [9]. **Table 1** summarises future-

facing pharmaceutical sciences capabilities required for responsible AI-era research and translation.

Table 1. Pharmaceutical Sciences Capabilities in the AI Era: Data Infrastructure, High-Quality Datasets, Model Documentation, Reproducibility, Benchmarking, Uncertainty, Bias Assessment, Human Oversight, Safety Evaluation, Regulatory Science, Interdisciplinary Collaboration, and Claim Boundaries

Research domain	Core future question	Required capability	AI-era function	Potential research output	Risk of overinterpretation	Validation or governance need	Agenda role
AI-enabled drug discovery	Which pharmaceutical questions materially benefit from AI rather than simpler or established methods?	Prospective task definition, baseline comparison, and intended-use specification	Prediction, representation learning, prioritization, and constrained generation	Ranked hypotheses or experimentally testable candidates	Benchmark improvement interpreted as discovery success	Comparative evaluation, prospective testing, and expert review	Cross-cutting enabling capability
Molecular intelligence	How can chemical and biological evidence be connected without obscuring differences in evidence quality?	Provenance-aware multimodal representation	Evidence linkage, pattern discovery, and hypothesis generation	Traceable compound–target–pathway hypothesis	Integrated evidence presented as a validated mechanism	Evidence grading, source attribution, uncertainty analysis, and experimental validation	Scientific integration layer
Data infrastructure	How can heterogeneous pharmaceutical data remain identifiable, versioned, governed, and reusable?	Interoperable repositories, identifiers, ontologies, audit trails, and access controls	Data discovery, harmonization, and controlled integration	Governed research-ready data environment	Assuming data aggregation automatically improves reliability	Stewardship, security, access policy, and lifecycle management	Foundational infrastructure
High-quality datasets	What minimum identity, assay, metadata, provenance, and quality criteria should precede model training?	Domain-specific curation, deduplication, and quality-control standards	Training, validation, and benchmarking input	Fit-for-purpose dataset with documented limitations	Scaling inaccurate labels, duplicates, or incompatible assays	Independent quality audit and leakage assessment	Foundational quality priority
Model documentation	What information is required to reproduce and independentl	Versioned model dossier, data description, intended-use	Transparency, auditability, and comparison	Reproducible model package	Documentation treated as evidence that a model is valid	Change control, independent review, and update logs	Mandatory governance layer

	y assess a pharmaceutical model?	statement, and limitation record					
Reproducibility	Can independent teams reconstruct the data, splits, analysis, training, and evaluation?	Archived code, environments, preprocessing records, and reproducible workflows	Verification of computational findings	Reproducible analysis and evaluation record	Computational repeatability confused with biological reproducibility	Independent replication and experimental confirmation	Mandatory evidence condition
Benchmarking	Do evaluation tasks measure generalization relevant to pharmaceutical research?	Scaffold, temporal, external, and prospective test designs	Comparable model assessment	Performance profile across relevant domains and baselines	Leaderboard optimization, hidden leakage, or narrow task success	Locked tests, predefined metrics, and domain-shift assessment	Comparative-evaluation priority
Uncertainty estimation	When should a model abstain or request additional evidence?	Calibrated uncertainty, applicability-domain analysis, and abstention rules	Confidence-aware research prioritization	Prediction accompanied by uncertainty and domain status	Uncertainty output interpreted as a guaranteed probability of correctness	Calibration testing and monitoring under distribution shift	Risk-management priority
Bias assessment	Which chemical, biological, geographic, demographic, or publication biases shape the evidence?	Coverage analysis, subgroup testing, and bias audits	Identification of systematic performance limitations	Bias and representativeness report	Assuming a single fairness metric resolves contextual bias	Domain-specific interpretation and remediation planning	Responsible-research priority
Human oversight	Which outputs require review by chemists, pharmacologists, toxicologists, clinicians, data scientists, or other experts?	Defined competencies, decision rights, and escalation procedures	Contextual interpretation, exception handling, and stop decisions	Reviewed research recommendation	Symbolic human participation without authority or relevant expertise	Decision logs, conflict-of-interest management, and accountable review	Mandatory control
Safety evaluation	How can foreseeable risks be investigated before candidate progression?	Integrated ADMET, toxicology, interaction, and exposure workflow	Early hazard identification and experiment prioritization	Risk hypothesis and testing plan	In-silico screening interpreted as safety clearance	Tiered experimental testing and specialist toxicology review	Safety-by-design priority

Regulatory science	What evidence and documentation may later be needed for regulatory assessment?	Early evidence mapping, product definition, model traceability, and quality planning	Organization of development evidence and uncertainties	Regulatory research-gap map	Research alignment interpreted as predicted approval	Jurisdiction-specific expert assessment and evidence updating	Translational planning priority
Interdisciplinary collaboration	How can computational, chemical, pharmacological, clinical, ethical, and sustainability expertise remain connected?	Shared terminology, evidence models, responsibilities, and review milestones	Coordinated research design and interpretation	Integrated research protocol	Parallel workstreams presented as genuine integration	Joint governance and multidisciplinary milestone review	Organizational enabler
Implementation-readiness boundary	What evidence is required before moving from research prototype to operational evaluation?	Readiness taxonomy, security review, external validation, and stage gates	Classification of research maturity	Readiness-gap assessment	Prototype performance presented as deployability	Workflow evaluation, governance review, and prospective testing	Protective boundary
No clinical-claim boundary	Which conclusions remain unsupported before appropriate clinical evidence is available?	Claim taxonomy and evidence-linked communication rules	Controls interpretation and dissemination	Explicit evidence-bounded statement	Molecular, computational, or preclinical findings portrayed as clinical efficacy	Clinical evidence assessment and independent expert review	Non-negotiable claim boundary

Molecular intelligence and natural product discovery

Molecular intelligence in natural product research begins with the definition of the entity being studied. A botanical extract, microbial fraction, purified metabolite, semisynthetic analogue, and standardized preparation cannot be treated as interchangeable data objects. Their taxonomic source, tissue, geography, collection conditions, processing history, extraction procedure, batch identity, analytical profile, purity, stereochemistry, and uncertainty must remain connected to every downstream assay and computational representation. Existing natural product resources demonstrate the value of structured chemical and biological data, but they also reveal fragmentation, uneven coverage, inconsistent identifiers, variable provenance, and differences in annotation quality that future infrastructures must address [10].

Analytical computation can improve the efficiency with which complex mixtures and metabolomic signals are organized. Tandem mass-spectrometry interpretation can generate molecular-formula and structure candidates, assign confidence levels, and prioritize features for further investigation, but computational annotation should not be presented as definitive structural confirmation when authentic standards or orthogonal analytical evidence are absent [11]. Molecular networking can complement this process by grouping related spectral features and revealing metabolite families, chemical similarities, or distribution patterns across samples. These relationships can support discovery prioritization, although network proximity alone does not establish identity, shared bioactivity, target engagement, or pharmacological mechanism [12].

Knowledge graphs provide one architecture for connecting natural product sources, chemical entities, assays, targets, pathways, diseases, literature statements, omics observations, and supported clinical signals. Their principal value lies not simply in the number of connected

entities but in whether each consequential relationship retains source attribution, evidence type, confidence, temporal status, and contradiction information [13]. Network pharmacology and systems pharmacology can build on these structures to examine multicomponent and multitarget hypotheses relevant to phytopharmacology. However, database-derived target lists, enrichment analyses, and highly connected networks can reproduce literature bias or create circular interpretations unless the underlying material is chemically characterized and the predicted relationships are tested experimentally [14].

Multimodal AI may further integrate chemical structures, analytical spectra, assay results, omics profiles, textual evidence, source metadata, and exposure information. Its research value should be assessed against single-modality baselines, with attention to missingness, shortcut learning, data imbalance, source duplication, and the possibility that one dominant modality masks weak evidence elsewhere. Generative chemistry may be used where appropriate to explore natural product-inspired chemical space or optimize defined molecular objectives, but generated structures should be treated as design hypotheses whose plausibility depends on chemical validity, novelty, synthesizability, stability, accessibility, and compatibility with the intended pharmacological question [15]. Medicinal-chemistry assessment and experimental testing remain necessary because numerical optimization does not by itself establish that a generated candidate is useful, safe, or biologically relevant [16].

Within the proposed agenda, molecular intelligence should operate through explicit validation and claim-control gates. Target prediction should lead to testable target-engagement questions; pathway hypotheses should lead to perturbational and temporal experiments; bioactivity should be interpreted alongside assay metadata and material identity; and mechanistic propositions should remain distinguishable from correlations or model explanations. Safety-by-design constraints should be introduced during prioritization and generation rather than deferred until a preferred candidate has emerged. Exposure modeling should test whether observed activity is plausible at achievable concentrations, while ADMET and toxicity evidence should be developed progressively through applicability-aware computation and fit-for-purpose experiments. These elements are proposed as an integrated research architecture and not as evidence that a candidate possesses therapeutic efficacy, clinical utility, or an acceptable safety profile. **Table 2** maps molecular intelligence capabilities required for future-ready natural product discovery.

Table 2. Molecular Intelligence and Natural Product Discovery Capabilities: Natural Product Identity, Phytochemical Profiles, Chemical Structure Representation, Bioactivity Evidence, Omics Evidence, Text Evidence, Knowledge Graphs, Network Pharmacology, Multimodal AI, Generative Chemistry, Safety-by-Design, Mechanistic Evidence, Exposure Modeling, and Claim Boundaries

Discovery component	Core discovery question	Required evidence or input	Molecular intelligence function	Potential output	Failure risk	Validation requirement	Agenda role
Natural product identity	Is the studied material chemically and biologically unambiguous?	Voucher or strain information, analytical profile, purity, batch data, and stable identifiers	Entity resolution and identity tracking	Versioned natural product research entity	Misidentification or conflation of extracts, fractions, and compounds	Expert authentication and orthogonal analytical confirmation	Entry gate
Botanical or biological source	Which organism, tissue, location, season, and processing history produced the material?	Taxonomy, geography, collection, cultivation, processing, and provenance metadata	Source–chemistry linkage	Traceable source record	Loss of ecological, biological, or cultural context	Source documentation and provenance audit	Provenance foundation
Phytochemical profile	Which constituents are present, at what abundance,	Chromatography, mass spectrometry, spectroscopy, standards, and	Feature annotation and batch comparison	Confidence-graded composition map	Overannotation, batch inconsistency, or undetected contaminants	Reference standards, orthogonal analysis, and	Composition layer

	and with what identification confidence?	replicate profiles				replicate testing	
Chemical structure representation	Which representation preserves stereochemistry, uncertainty, mixtures, and relevant molecular features?	Canonical structures, stereochemical notation, descriptors, fingerprints, and molecular graphs	Computable molecular encoding	Model-ready representation linked to identity confidence	Loss of stereochemistry or uncertain structures treated as confirmed	Representation audit and structure verification	Computational foundation
Bioactivity evidence	What activity was observed under which concentration, time, and assay conditions?	Protocol, controls, concentration range, endpoint, replicates, and raw measurements	Compound–assay and mixture–assay linkage	Contextualized bioactivity record	Context-free potency values or assay artifacts	Orthogonal assays, concentration–response analysis, and appropriate controls	Functional evidence layer
Assay metadata	Are target construct, cell system, platform, timing, quality criteria, and interference risks documented?	Complete experimental metadata and quality-control results	Measurement contextualization	Comparable and filterable assay records	Incompatible assays treated as equivalent	Assay-quality review and interference testing	Data-integrity gate
Omics evidence	Which molecular states or pathways change following exposure?	Transcriptomic, proteomic, metabolomic, epigenomic, dose, and time-course data	Pattern detection and pathway hypothesis generation	Mechanistic hypothesis with evidence provenance	Correlation interpreted as causation	Independent replication, perturbation, and temporal validation	Mechanism-enrichment layer
Text evidence	What relevant knowledge is available in peer-reviewed literature and curated records?	Source-linked text corpora, controlled vocabularies, and retrieval metadata	Entity extraction, relation extraction, and evidence retrieval	Traceable literature-derived relation	Hallucinated, duplicated, outdated, or decontextualized statements	Retrieval of underlying sources and expert verification	Evidence-discovery layer
Knowledge graphs	How can heterogeneous entities and relationships remain queryable and provenance-aware?	Ontologies, normalized entities, relation evidence, confidence, and versioning	Multimodal relation integration	Natural product knowledge graph	False relationships, circular inference, or popularity bias	Relation-level provenance review and competency testing	Core integration architecture
Molecular networking	Which analytical features form chemically related families	Processed spectra, similarity metrics, sample metadata, and	Spectral relationship mapping	Prioritized metabolite families	Network proximity interpreted as structural identity or	Orthogonal structural and biological confirmation	Analytical prioritization layer

	worthy of follow-up?	confidence thresholds			shared activity		
Network pharmacology	Which multicomponent–multitarget relationships merit experimental investigation?	Characterized composition, credible target data, disease context, and pathway evidence	Network-based hypothesis generation	Ranked target and pathway hypotheses	Database bias, degree bias, and circular pathway interpretation	Target-engagement and pathway-validation experiments	Conditional discovery method
Systems pharmacology	How might molecular interactions connect with dynamic biological responses?	Mechanistic networks, dose, time, exposure, and physiological context	Multiscale response modeling	Testable systems-level hypothesis	Static associations presented as dynamic causal mechanisms	Perturbational, temporal, and exposure-informed validation	Mechanistic integration layer
Multimodal AI	Does combining evidence types improve generalization and scientific usefulness?	Aligned chemical, analytical, assay, omics, textual, and contextual data	Cross-modal representation and prediction	Integrated prediction with modality attribution	Missingness, shortcut learning, modality dominance, or hidden duplication	Ablation studies, external validation, uncertainty analysis, and baseline comparison	Advanced integration priority
Generative chemistry	Can natural product-informed chemical space be explored under realistic pharmaceutical constraints?	Valid structures, objective functions, novelty rules, synthesizability, and safety constraints	Candidate generation and molecular optimization	Ranked molecular design hypotheses	Invalid, unstable, inaccessible, trivial, or unsafe structures	Medicinal-chemistry review, synthesis, and experimental assays	Hypothesis-generation layer
Safety-by-design	Which foreseeable hazards should constrain prioritization from the beginning?	Toxicity alerts, reactive motifs, ADMET evidence, interaction data, and exposure assumptions	Risk-aware ranking and constrained generation	Candidate hypothesis with preliminary risk profile	False reassurance, incomplete hazard coverage, or optimisation around imperfect predictors	Tiered toxicology, interaction testing, and expert review	Mandatory safety layer
Target prediction	Which targets are plausible for a defined compound, fraction, or characterized mixture?	Structural, bioactivity, omics, textual, and network evidence	Target-ranking and evidence integration	Ranked target hypotheses	Popular-target bias, chemical-domain mismatch, and unsupported mixture attribution	Biophysical, biochemical, cellular, genetic, or pharmacological validation	Mechanistic hypothesis stage
Mechanistic evidence	What causal sequence connects exposure with the reported biological response?	Target engagement, perturbation, rescue, temporal evidence, pathway	Evidence-chain organization	Confidence-graded mechanistic model	Correlation chains or model explanations labeled as mechanisms	Orthogonal perturbation, rescue studies, and independent replication	Central validation objective

		evidence, and phenotype					
Exposure modeling	Are active concentrations achievable, sustained, and relevant to the intended biological site?	Physicochemical properties, formulation, absorption, metabolism, tissue distribution, and concentration data	Exposure-feasibility assessment	Exposure-informed research prioritization	Ignoring bioavailability, metabolites, protein binding, or tissue distribution	Qualified analytical measurements and PK evaluation	Translation bridge
ADMET evidence	Which absorption, distribution, metabolism, excretion, and transport liabilities require investigation?	Applicability-aware predictions, in-vitro assays, and in-vivo evidence where justified	Multi-endpoint developability assessment	ADMET evidence profile with uncertainty	Prediction beyond chemical domain or unsupported endpoint extrapolation	Experimental confirmation and endpoint-specific qualification	Development gate
Toxicity evidence	Which acute, chronic, organ-specific, genotoxic, reproductive, or immunological hazards are plausible?	Structural alerts, mechanistic evidence, validated assays, models, and exposure context	Hazard prioritization	Toxicity hypothesis and testing plan	Absence of predicted alerts interpreted as absence of risk	Dose- and exposure-relevant toxicological evaluation	Safety gate
Therapeutic-claim boundary	What may legitimately be concluded from discovery-stage evidence?	Evidence hierarchy, intended-use statement, and claim taxonomy	Controls interpretation and communication	Research hypothesis or prioritization statement	Computational, molecular, or preclinical evidence presented as therapeutic efficacy	Independent preclinical and clinical evidence where applicable	Non-negotiable boundary

Responsible therapeutic translation

Responsible therapeutic translation begins by distinguishing predictive performance from evidence that can support a pharmaceutical or clinical interpretation. Molecular models should be evaluated with task definitions, datasets, splitting strategies, metrics, and baselines that reflect the intended research question rather than convenience alone. Standardized molecular-learning benchmarks can improve comparability, but benchmark performance should be treated as one component of an evidence package and not as proof that a model will generalize to new natural products, biological systems, development settings, or patient populations [17].

Early developability assessment should combine molecular-intelligence outputs with exposure and safety questions. Computational ADMET models may help prioritize absorption, distribution, metabolism, excretion, transporter, and physicochemical investigations, but their outputs remain dependent on endpoint definition, training-domain coverage, and the quality of available chemical data [18]. Integrated platforms that include toxicity-related endpoints can broaden preliminary risk triage, yet predictions should direct experiments rather than function as safety determinations. For chemically complex extracts, uncertain metabolites, or underrepresented natural product scaffolds, applicability-domain limitations and constituent-specific effects require particular attention [19].

Uncertainty should be reported alongside molecular predictions so that researchers can distinguish relatively supported estimates from outputs that warrant abstention, additional data collection, or alternative methods. Comparative studies of uncertainty approaches show that no single technique can be assumed to perform reliably

across every molecular task, making endpoint-specific calibration and external testing necessary [20]. Benchmark construction also requires protection against chemical similarity leakage and memorization. When closely related ligands appear across training and test sets, apparent performance may overstate the ability to generalize to structurally distinct compounds, new botanical chemotypes, or prospectively generated candidates [21].

Natural product translation must also account for interaction risk, material variability, and the evidentiary limitations of phytopharmacological studies. Herb–drug interaction signals should be evaluated through structured causality reasoning that considers composition, dose, timing, exposure, concomitant treatments, alternative explanations, and the completeness of case information rather than relying on association alone [22]. More broadly, reproducible phytopharmacology requires authenticated source materials, adequate chemical characterization, appropriate controls, biologically relevant models, transparent dosing, and cautious alignment between observed findings and the claims being made [23]. Target validation, pathway validation, and exposure–response assessment should therefore be planned as connected but distinct activities, with PK/PD modeling used only where the available data and assumptions can support it.

The proposed translation logic places safety evaluation, model documentation, regulatory-science planning, sustainability, ethical sourcing, benefit-sharing, and human oversight alongside mechanistic and exposure evidence. Responsibility should be considered across the research lifecycle, including problem formulation, data selection, model development, experimental validation, interpretation, dissemination, and feedback updating, rather than being reduced to an endpoint review [24]. Clinical evidence is required where clinical efficacy or utility is at issue, while regulatory assessment remains jurisdiction-, product-, and evidence-specific. Neither a molecular hypothesis, a favorable preclinical result, a documented model, nor a responsibly framed research program establishes clinical utility or predicts regulatory approval. **Figure 1** illustrates how molecular intelligence, natural product discovery, safety evaluation, regulatory science, sustainability, and responsible governance can be connected for therapeutic translation without converting research hypotheses into clinical claims.

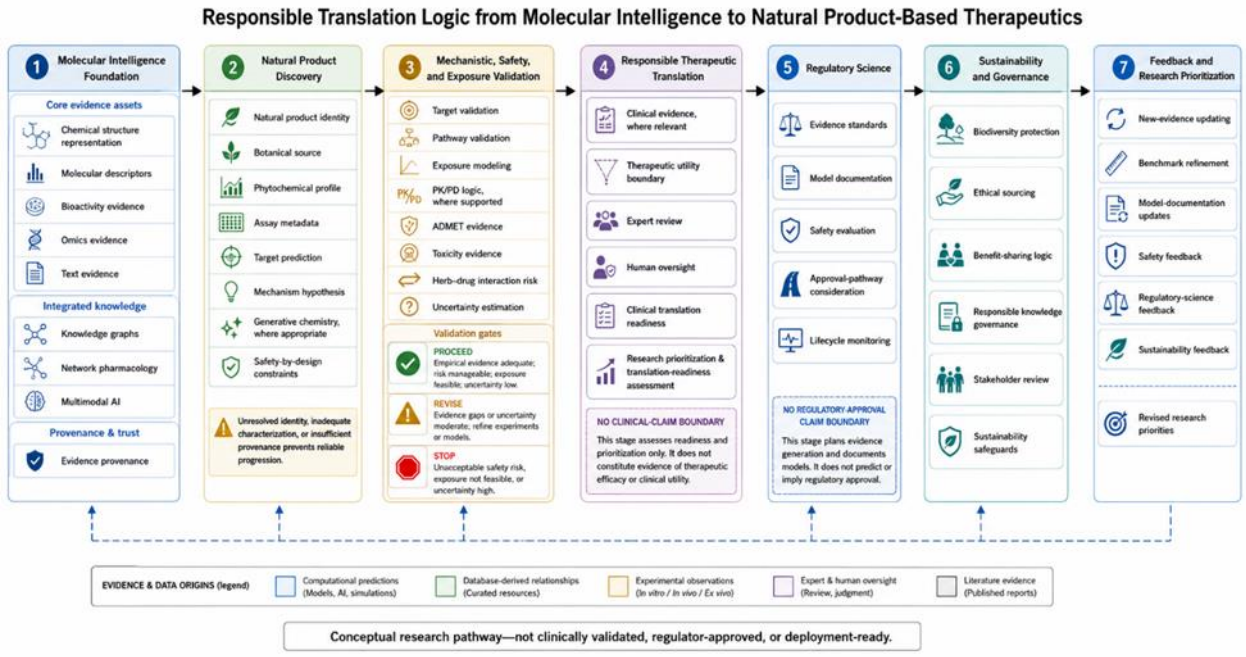

Figure 1. Responsible Translation Logic from Molecular Intelligence to Natural Product-Based Therapeutics

Alt-text description

A conceptual responsible translation diagram showing molecular intelligence and natural product discovery connected to mechanism evidence, safety evaluation, exposure modeling, regulatory science, sustainability, ethical sourcing, benefit-sharing, expert review, validation planning, and clinical-claim boundaries.

Strategic research priorities

The first strategic priority is to develop evaluation systems that measure whether molecular models provide reproducible and pharmaceutically meaningful value. Generative chemistry requires more than validity percentages or visually plausible structures; it requires standardized assessments of distribution learning, novelty,

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

diversity, goal-directed optimization, chemical filters, and task-specific usefulness. Complementary molecular-generation benchmarks provide a foundation for comparison, but their metrics should be combined with natural product-specific applicability analysis, medicinal-chemistry review, safety constraints, and prospective experiments [25, 26]. Benchmark governance should also include fixed test sets, transparent preprocessing, relevant baselines, and protections against repeated tuning to public evaluation data.

A second priority is governed collaboration across institutions that hold heterogeneous, confidential, culturally sensitive, or proprietary data. Federated learning may permit model training or evaluation without routinely centralizing raw records, potentially enabling collaboration across natural product laboratories, pharmaceutical groups, omics centers, clinical research environments, and other authorized partners. However, distributed training does not automatically guarantee privacy, security, fairness, data compatibility, or reliable performance. Site-level data quality, non-identically distributed populations, secure update exchange, contribution governance, withdrawal rights, and independent security assessment remain necessary [27].

Causal inference and digital-twin research should be pursued selectively where the scientific question and evidence density justify them. Causal approaches may help researchers formulate intervention-relevant questions, distinguish predictive associations from causal hypotheses, and expose assumptions that require testing, but they remain vulnerable to confounding, selection bias, measurement error, and unsupported identification conditions [28]. Digital twins may offer a future means of integrating dynamic mechanistic, exposure, biomarker, and patient-level information, yet their value must be demonstrated against simpler models and through prospective validation. They should not be presented as individualized therapeutic decision systems while their data, mechanistic fidelity, update rules, and clinical consequences remain unvalidated [29].

Model documentation and reproducibility infrastructure constitute a fourth priority. Every consequential model should have a clear purpose, intended users, data lineage, preprocessing record, endpoint definition, architecture and training description, evaluation design, performance limits, uncertainty method, version history, and prohibited-use statement. Structured reporting frameworks can improve independent assessment and comparison, although documentation alone cannot establish scientific or clinical validity [30]. Research programs must also prevent leakage between training, validation, and testing; preserve temporal and institutional separation where relevant; archive code and environments; and distinguish repeatable computation from reproducible biological evidence. Leakage and inadequate evaluation design can otherwise generate persuasive but unreliable findings that do not survive independent examination [31].

The fifth priority is to connect technical development with mechanistic validation, safety-by-design, regulatory science, sustainability, ethical sourcing, benefit-sharing, and workforce development. Future infrastructures should support natural product knowledge graphs, multimodal evidence systems, exposure and PK/PD research, prospective safety workflows, causal and federated methods where appropriate, and external validation across relevant chemical and biological domains. At the organizational level, interdisciplinary teams require shared evidence terminology, defined decision rights, specialist training, and explicit stop criteria. Clinical translation readiness should be treated as an evidence-gap assessment rather than a declaration of utility, and no proposed platform, model, consortium, or roadmap should be described as deployment-ready without operational, security, governance, and prospective validation. **Table 3** summarises strategic research priorities for future-ready pharmaceutical sciences and phytopharmacology.

Table 3. Strategic Research Priorities for Future-Ready Pharmaceutical Sciences and Phytopharmacology: Data Integrity, Multimodal Evidence Integration, Knowledge Graphs, Benchmarking, Mechanistic Validation, Safety-by-Design, Exposure Modeling, Federated Collaboration, Regulatory Science, Sustainability, Ethical Sourcing, Benefit-Sharing, Clinical Translation Readiness, and Workforce Development

Strategic priority	Core research need	Required infrastructure or method	Expected research contribution	Failure risk	Success condition	Responsible oversight need	Priority boundary
Data integrity infrastructure	Preserve identity, provenance, metadata quality, versioning, and traceability	Governed repositories, ontologies, entity resolution, audit trails, and quality-control pipelines	Reusable and independently auditable evidence ecosystem	Scaling ambiguous entities, mislabeled assays, and duplicated records	Data lineage and transformations can be independently reconstructed	Data stewards, analytical scientists, pharmacologists, and governance specialists	Data integrity does not establish biological validity

Multimodal evidence integration	Link chemical, assay, omics, textual, exposure, and contextual evidence without collapsing evidence levels	Provenance-aware schemas, modality alignment, missingness handling, and interpretable integration methods	More complete and testable discovery hypotheses	Shortcut learning, modality dominance, hidden duplication, and incompatible evidence	Each modality's contribution, uncertainty, and limitations remain visible	Computational, experimental, and domain-specific review	Integration is not mechanistic validation
Natural product knowledge graphs	Represent sources, compounds, extracts, assays, targets, pathways, literature, and governance context	Curated ontologies, normalized identifiers, relation provenance, and version control	Queryable natural product evidence network	Circular inference, relation hallucination, and popularity bias	Consequential relationships are source-traceable and independently reviewable	Knowledge engineers, pharmacologists, analytical scientists, and ethnopharmacologists	Graph connectivity is not causality
Benchmarking standards	Measure generalization and practical research usefulness	Scaffold, temporal, external, institutional, and prospective benchmarks	Comparable evidence on model performance and limitations	Leakage, narrow-task optimization, and leaderboard overfitting	Locked tests, meaningful baselines, predefined metrics, and external evaluation	Independent benchmark governance and domain review	Benchmark success is not therapeutic success
Mechanistic validation pipelines	Convert predicted associations into experimentally tested causal explanations	Target-engagement studies, perturbation, rescue, temporal experiments, and pathway analysis	Validated, revised, or rejected mechanism hypotheses	Selective validation and confirmation bias	Orthogonal methods reproduce the relationship and plausible alternatives are tested	Pharmacologists, chemical biologists, and experimental-design specialists	Model explanation is not a validated mechanism
Safety-by-design workflows	Integrate foreseeable hazards before candidate optimization and progression	Structural alerts, ADMET, toxicity, interaction, exposure, and formulation assessments	Earlier identification of avoidable risk and defined testing priorities	False reassurance from incomplete or out-of-domain predictors	Safety hypotheses produce explicit tests, margins, and stop criteria	Toxicologists, clinical pharmacologists, chemists, and safety-governance reviewers	Screening does not establish safety
Exposure and PK/PD modeling	Determine whether biological activity is achievable at relevant sites and times	Bioanalysis, formulation data, metabolism studies, exposure models, and response biomarkers	Exposure-informed prioritization and study design	Potency-based ranking at biologically unrealistic concentrations	Model predictions are qualified and compared with measured exposure and response	Pharmacometricians, bioanalytical scientists, and pharmacologists	Modeling does not establish a clinical dose
Digital-twin research where appropriate	Test whether dynamic individualized models add value over	Longitudinal multimodal data, mechanistic models, update rules, and	Research hypotheses concerning individual response trajectories	Underidentified models, overpersonalization, and unstable updates	Prospective evaluation demonstrates added value and acceptable uncertainty	Clinical, statistical, computational, and ethics oversight	No therapeutic use before appropriate validation

	simpler approaches	prospective protocols					
Causal inference research where appropriate	Distinguish intervention effects from predictive associations	Causal diagrams, appropriate designs, sensitivity analyses, and explicit assumptions	Better-defined intervention and mechanism questions	Residual confounding and unsupported identification assumptions	Assumptions are transparent and findings remain robust across sensitivity analyses	Causal-methods experts and relevant domain scientists	Causal language is limited to what the design supports
Federated learning collaboration where appropriate	Support research across protected, culturally governed, or proprietary data environments	Secure aggregation, harmonization, authentication, site-level validation, and consortium governance	Multi-site learning without routine raw-data centralization	Privacy leakage, incompatible labels, non-IID bias, and institutional conflict	Collaboration improves external evidence without unacceptable privacy or governance risk	Privacy, security, legal, community, and scientific governance	Federated learning does not itself guarantee privacy or generalizability
Regulatory science alignment	Anticipate evidence, quality, documentation, and product-classification questions	Product definition, model documentation, quality planning, and staged regulatory research	Earlier identification of translational evidence gaps	Treating regulatory planning as predicted approval	Evidence gaps are explicit, jurisdiction-specific, and updated as research changes	Qualified regulatory scientists and product-development specialists	No approval or submission-readiness claim
Sustainability and ethical sourcing	Investigate environmental, biodiversity, labor, cultural, and supply-chain consequences	Biodiversity assessment, traceable sourcing, lifecycle methods, and stakeholder engagement	More responsible candidate, process, and sourcing decisions	Assuming natural origin is inherently sustainable or ethical	Source and process safeguards are verifiable and monitored over time	Ecologists, sourcing experts, ethics specialists, and source communities	No sustainability claim without relevant evidence
Benefit-sharing governance	Define fair participation in knowledge, value creation, and research benefits	Context-specific agreements, decision roles, accountability, and revision procedures	More legitimate and durable research partnerships	Boilerplate agreements, token consultation, and extractive knowledge use	Benefits and decision rights are negotiated early and remain revisable	Community representatives and qualified legal and ethical experts	The research agenda is not a legally sufficient agreement
Clinical translation readiness	Specify evidence needed before clinical or implementation research	Stage gates, external validation, safety evidence, manufacturing information, and prospective protocols	Transparent translation-readiness gap map	Premature efficacy, utility, or workflow claims	Independent evidence meets predefined progression criteria	Clinical, pharmaceutical, regulatory, and methodological experts	Readiness assessment is not clinical utility
Research workforce development	Build expertise across computation, chemistry, pharmacology, clinical science,	Interdisciplinary education, shared competencies, rotations, and joint review structures	Better-integrated study design and interpretation	Superficial collaboration and unrecognized expertise gaps	Roles, competencies, and decision authority are explicit and demonstrable	Institutional leadership and specialist professional communities	Training does not replace qualified expert review

	regulation, sustainability, and governance						
No deployment-readiness claim boundary	Prevent prototypes and proposed infrastructures from being described as operational systems	Readiness taxonomy, security review, governance assessment, and publication controls	Accurate communication of maturity and evidence gaps	Research prototype presented as deployable or decision-ready	Every output is labeled by evidence level, intended use, and readiness status	Independent scientific, governance, and editorial review	No deployment claim without operational validation

Future research agenda

The proposed agenda begins with a research question that is pharmaceutically meaningful, empirically testable, and explicit about the decision it is intended to inform. It then develops a governed data ecosystem linking natural product identity, chemical representation, assays, omics, text, provenance, and supported translational evidence. An integrated therapeutic-science foundation can help connect molecular and biological learning tasks across scales, but such integration should remain coupled to experimental feedback rather than treated as a self-validating computational layer [32]. Prospective discovery studies illustrate that AI prioritization becomes scientifically consequential when predicted candidates undergo laboratory evaluation [33], while generative design studies similarly show that generated molecules require synthesis and biochemical or cellular testing before their relevance can be assessed [34].

After data and molecular-intelligence development, the agenda advances through candidate prioritization, mechanistic evidence generation, safety-aware design, and exposure–response investigation. Progression should depend on predefined evidence gates rather than the novelty of the model or the apparent strength of a single result. Clinical AI research has shown how biased designs, incomplete reporting, weak comparators, retrospective evaluation, and insufficient external validation can lead to overstated claims [35]. The same caution applies to pharmaceutical and phytopharmacological translation: technical performance does not establish clinical impact, which also depends on workflow relevance, prospective evaluation, safety, stakeholder participation, and an appropriate regulatory and organizational context [36].

Responsible therapeutic translation is followed by regulatory-science planning and explicit governance of model documentation, evidence standards, intended use, safety evaluation, and lifecycle updating. Ethical principles such as transparency, responsibility, fairness, privacy, and non-maleficence can orient research, but international ethics frameworks differ in scope, interpretation, and operational detail [37]. Principles alone cannot ensure responsible implementation; institutions need defined responsibilities, review procedures, enforcement mechanisms, expertise, resources, and routes for contesting or revising decisions [38]. Accordingly, responsible innovation should be assessed through observable governance practices rather than through labels attached to a project.

For natural product research, sustainability, sourcing, and knowledge governance should be integrated before data collection and resource acquisition. Ethnopharmacological field research requires documentation of research context, source communities, permissions, methods, taxonomy, knowledge transmission, and relevant ethical and legal arrangements [39]. Access and benefit-sharing questions likewise require early, context-sensitive engagement because rights, expectations, institutions, resource status, and feasible benefits vary across settings [40]. The proposed agenda therefore treats biodiversity protection, ethical sourcing, responsible knowledge governance, and benefit-sharing as research-design responsibilities, while recognizing that an academic framework cannot certify legal compliance, community approval, ecological sustainability, or fairness.

The agenda concludes with collaborative infrastructure, multidisciplinary expert review, feedback updating, and periodic reprioritization. Validation findings, negative results, safety signals, data corrections, model failures, regulatory-science questions, sourcing concerns, and stakeholder feedback should return the research program to the relevant earlier stage. Continue, revise, and stop decisions should be documented and proportionate to uncertainty. The agenda is proposed as an adaptive planning structure rather than an implementation standard: it can organize questions, identify evidence gaps, coordinate expertise, and clarify progression criteria, but it cannot validate a therapy, establish safety, define a clinical dose, guarantee sustainability, confer regulatory status, or authorize therapeutic decision-making. **Table 4** presents the proposed future research agenda for responsible natural product-based therapeutics.

Table 4. Proposed Future Research Agenda for Responsible Natural Product-Based Therapeutics: Data Ecosystems, Molecular Intelligence, Discovery Acceleration, Mechanistic Evidence, Safety-Aware Design, Exposure and Response Modeling, Responsible Translation, Regulatory Science, Sustainability, Benefit-Sharing, Collaborative Infrastructure, Expert Review, Feedback, and Decision Boundaries

Agenda stage	Core purpose	Required input	Strategic function	Potential output	Validation need	Failure risk	Feedback mechanism	Decision boundary
Research question definition	Define a meaningful, bounded pharmaceutical or phytopharmacological problem	Evidence gap, biological rationale, stakeholder need, intended use, and prohibited claims	Prevent technology-first or method-driven research	Testable research question and success criteria	Multidisciplinary expert and stakeholder review	Retrofitting a question to an available model or dataset	Revise the question after feasibility and evidence review	No therapeutic recommendation
Data ecosystem development	Build trustworthy, reusable, and governed evidence foundations	Authenticated materials, stable identifiers, assay metadata, provenance, access rules, and quality criteria	Enable reliable modeling, integration, and replication	Versioned research dataset with documented limitations	Quality audit, entity-resolution checks, and leakage assessment	Heterogeneous low-quality aggregation and provenance loss	Continuous curation, error correction, and version updates	Data availability does not establish fitness for every purpose
Molecular intelligence integration	Connect structure, bioactivity, omics, text, source context, and exposure evidence	Harmonized modality-specific records and relation provenance	Generate traceable cross-modal research hypotheses	Evidence graph, multimodal representation, or integrated prediction	Ablation analysis, uncertainty estimation, baseline comparison, and external evaluation	Shortcut learning, unsupported relations, and hidden modality bias	Contradictions and new evidence update relations and models	Integrated evidence does not establish mechanism
Natural product discovery acceleration	Prioritize compounds, fractions, extracts, targets, or analogues for testing	Confirmed identity, composition, bioactivity, source data, and task definition	Reduce experimental search space	Ranked research candidates or experiments	Prospective analytical and biological testing	Bias toward well-studied compounds, targets, organisms, or assays	Experimental findings update rankings and data quality	Prioritization is not efficacy
Mechanistic evidence generation	Test target, pathway, and systems hypotheses	Target predictions, exposure context, perturbation tools, and relevant models	Convert associations into testable causal evidence	Validated, revised, or rejected mechanistic hypothesis	Target engagement, perturbation, rescue, temporal studies, and replication	Confirmation bias, model-system mismatch, and pathway overinterpretation	Positive and negative results revise the evidence model	Mechanism claims apply only to tested conditions
Safety-aware design	Identify and reduce foreseeable liabilities early	Structural alerts, ADMET, toxicity, interaction, formulation, and exposure evidence	Constrain generation, selection, and progression	Risk-informed candidate set and testing plan	Tiered toxicology and interaction studies	Unknown hazards, out-of-domain prediction, and false reassurance	New safety findings change objectives, rankings, and stop criteria	No safety clearance from computational evidence

Exposure and response modeling	Determine whether observed activity is achievable and relevant	Concentration–response, formulation, metabolism, bioanalysis, tissue exposure, and biomarkers	Link discovery evidence with translational plausibility	Exposure–response hypothesis or preliminary model	Qualified measurement, model diagnostics, and prospective assessment	Unrealistic concentrations, ignored metabolites, and unsupported extrapolation	Measured exposure and response update assumptions	No dose recommendation
Responsible therapeutic translation	Organize quality, mechanism, exposure, safety, and potential-utility evidence	Validated preclinical evidence and clearly defined development context	Support evidence-based progression decisions	Translation evidence map and gap assessment	Independent multidisciplinary stage-gate review	Premature progression and claim inflation	Gate failure returns the program to earlier stages	Research planning rather than therapeutic decision-making
Regulatory science planning	Anticipate product definition, evidence, quality, documentation, and monitoring questions	Product concept, model dossier, source and manufacturing information, and evidence map	Identify avoidable translational gaps	Regulatory-science research plan	Qualified, jurisdiction-specific assessment	Treating planning, dialogue, or documentation as regulatory acceptance	Update after classification, evidence, or development changes	No approval or submission-readiness claim
Sustainability and sourcing governance	Investigate ecological, biodiversity, process, and supply-chain consequences	Resource status, chain of custody, extraction or production information, and stakeholder knowledge	Embed environmental and sourcing safeguards in research choices	Sustainability and sourcing evidence plan	Independent verification and longitudinal monitoring	Resource depletion, substitution, and unverifiable supply chains	Monitoring results change sources, methods, or candidate choices	Natural origin is not evidence of sustainability
Benefit-sharing and knowledge governance	Establish fair and culturally appropriate relationships around materials and knowledge	Provenance, community engagement, negotiated expectations, roles, and legal context	Protect contributors and distribute research participation and benefits more fairly	Context-specific governance and benefit-sharing plan	Ongoing review with affected partners and qualified experts	Extractive knowledge use, token consultation, and static agreements	Agreements and priorities update as value, risks, and expectations change	The agenda is not legal certification or community authorization
Collaborative infrastructure	Enable secure and reproducible multi-site research	Shared standards, technical interoperability, access governance, and site-level quality evidence	Increase evidence diversity and external evaluation	Governed collaborative research environment	Security, privacy, harmonization, and site-performance assessment	Non-IID bias, privacy failure, and institutional role conflict	Site-level results and governance findings modify collaboration	Collaboration does not guarantee generalizability
Expert review	Examine evidence, uncertainty, alternative explanations, conflicts, and	Complete evidence package, declared limitations,	Provide accountable interpretation and stop authority	Documented continue, revise, or stop recommendation	Independent multidisciplinary review	Rubber-stamp oversight or expertise mismatch	Dissent, uncertainty, and new evidence trigger	Expert opinion cannot waive evidence requirements

	progression criteria	and specialist assessments				reassessment		
Research priority updating	Re-rank questions as evidence, feasibility, risk, and societal context change	Validation outcomes, failures, surveillance, stakeholder feedback, and resource constraints	Keep the agenda adaptive	Revised research-priority portfolio	Transparent update criteria and conflict review	Trend-driven, sponsor-driven, or opaque reprioritization	Scheduled and event-triggered revision	Priority change is not proof of therapeutic value
No therapeutic-claim boundary	Preserve separation between research hypotheses and therapeutic claims	Evidence hierarchy, claim taxonomy, and intended-use statement	Control interpretation and communication	Explicit evidence-bounded conclusion	Independent scientific and editorial review	Computational, molecular, or preclinical findings portrayed as efficacy	Claims are revised when evidence status changes	No efficacy, safety, or clinical-utility claim
No implementation-standard boundary	Clarify that the agenda is a proposed research framework	Scope statement, governance limitations, and readiness taxonomy	Prevent institutional, commercial, clinical, or regulatory overclaiming	Explicit framework limitation statement	External methodological and stakeholder review	Proposed agenda portrayed as implemented, adopted, or authoritative	Empirical findings and stakeholder assessment revise the agenda	No implementation, certification, deployment, or authority claim

Figure 2 presents the proposed future research agenda for pharmaceutical sciences and phytopharmacology, integrating data ecosystems, molecular intelligence, responsible translation, regulatory science, sustainability, governance, collaboration, and feedback updating.

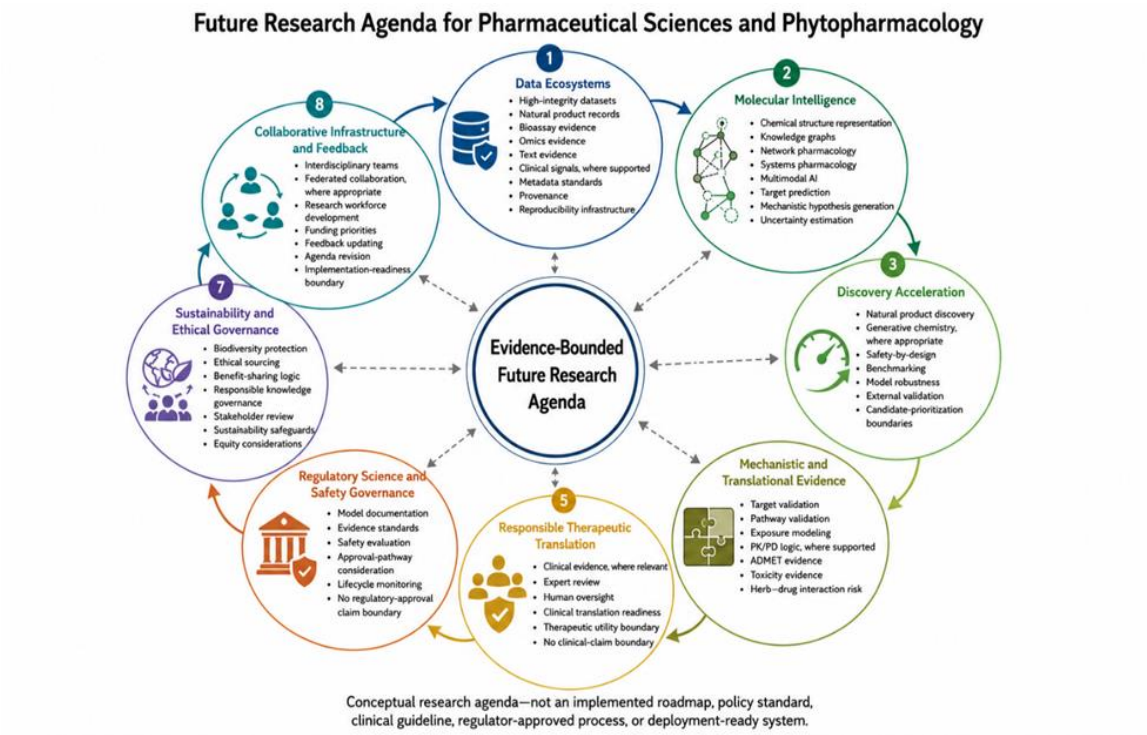

Figure 2. Future Research Agenda for Pharmaceutical Sciences and Phytopharmacology

Alt-text description

A conceptual future research agenda framework showing data ecosystems, molecular intelligence, discovery acceleration, mechanistic validation, safety-aware design, exposure modeling, responsible translation, regulatory science, sustainability, governance, collaboration, expert review, and feedback loops.

CONCLUSION

This article proposes a future research agenda for pharmaceutical sciences and phytopharmacology that connects high-integrity data ecosystems, molecular intelligence, natural product discovery, mechanistic validation, safety-aware design, exposure and response research, responsible therapeutic translation, regulatory science, sustainability, ethical sourcing, benefit-sharing, collaboration, expert oversight, and feedback updating. Its central contribution is to position these elements as an interdependent research system rather than as separate technical, translational, or governance workstreams. Future-ready natural product-based therapeutic research will require authenticated and characterized materials, reproducible evidence, robust and uncertainty-aware models, external and experimental validation, exposure-relevant pharmacology, staged safety evaluation, accountable governance, and interdisciplinary expertise. The proposed agenda can guide research priorities, identify evidence gaps, and clarify progression boundaries, but it cannot establish therapeutic efficacy, clinical utility, safety, sustainability, implementation readiness, or regulatory approval. Maintaining this separation between research prioritization and therapeutic claims is essential if innovation in AI-enabled pharmaceutical sciences and phytopharmacology is to remain scientifically credible and socially responsible.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

[1] Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5

[2] 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

[3] 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

[4] 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

[5] Mak KK, Pichika MR. Artificial intelligence in drug development: Present status and future prospects. Drug Discov Today. 2019;24(3):773-80. doi:10.1016/j.drudis.2018.11.014

[6] 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

[7] Zhu H. Big data and artificial intelligence modeling for drug discovery. Annu Rev Pharmacol Toxicol. 2020;60:573-89. doi:10.1146/annurev-pharmtox-010919-023324

[8] Jiménez-Luna J, Grisoni F, Schneider G. Drug discovery with explainable artificial intelligence. Nat Mach Intell. 2020;2(10):573-84. doi:10.1038/s42256-020-00236-4

[9] Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368. doi:10.1136/bmj.l6927

[10] 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

[11] Dührkop K, Fleischauer M, Ludwig M, Aksenov AA, Melnik AV, Meusel M, et al. SIRIUS 4: A rapid tool for turning tandem mass spectra into metabolite structure information. Nat Methods. 2019;16(4):299-302. doi:10.1038/s41592-019-0344-8

[12] Quinn RA, Nothias LF, Vining O, Meehan M, Esquenazi E, Dorrestein PC. Molecular networking as a drug discovery, drug metabolism, and precision medicine strategy. Trends Pharmacol Sci. 2017;38(2):143-54. doi:10.1016/j.tips.2016.10.011

[13] 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

[14] Noor F, Tahir ul Qamar M, Ashfaq UA, Albutti A, Alwashmi ASS, Aljasir MA. Network pharmacology approach for medicinal plants: Review and assessment. Pharmaceuticals (Basel). 2022;15(5):572. doi:10.3390/ph15050572

[15] Sánchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: Generative models for matter engineering. Science. 2018;361(6400):360-65. doi:10.1126/science.aat2663

[16] Walters WP, Murcko MA. Assessing the impact of generative AI on medicinal chemistry. Nat

Biotechnol. 2020;38(2):143-5. doi:10.1038/s41587-020-0418-2

[17] Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: A benchmark for molecular machine learning. Chem Sci. 2018;9(2):513-30. doi:10.1039/C7SC02664A

[18] 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

[19] 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

[20] 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

[21] Wallach I, Heifets A. Most ligand-based classification benchmarks reward memorization rather than generalization. J Chem Inf Model. 2018;58(5):916-32. doi:10.1021/acs.jcim.7b00403

[22] Awortwe C, Makiwane M, Reuter H, Muller C, Louw J, Rosenkranz B. Critical evaluation of causality assessment of herb–drug interactions in patients. Br J Clin Pharmacol. 2018;84(4):679-93. doi:10.1111/bcp.13490

[23] 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

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

[25] Brown N, Fiscato M, Segler MHS, Vaucher AC. GuacaMol: Benchmarking models for de novo molecular design. J Chem Inf Model. 2019;59(3):1096-108. doi:10.1021/acs.jcim.8b00839

[26] Polykovskiy D, Zhebrak A, Sánchez-Lengeling B, Golovanov S, Tatanov O, Belyaev S, et al. Molecular Sets (MOSES): A benchmarking platform for molecular generation models. Front Pharmacol. 2020;11:565644. doi:10.3389/fphar.2020.565644

[27] Rieke N, Hancox J, Li W, Milletarì F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. doi:10.1038/s41746-020-00323-1

[28] Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. Nat Mach Intell. 2020;2(7):369-75. doi:10.1038/s42256-020-0197-y

[29] Björnsson B, Borrebaeck C, Elander N, Gasslander T, Gawel DR, Gustafsson M, et al. Digital twins to personalize medicine. Genome Med. 2020;12(1):4. doi:10.1186/s13073-019-0701-3

[30] Norgeot B, Quer G, Beaulieu-Jones BK, Torkamani A, Dias R, Gianfrancesco M, et al. Minimum information about clinical artificial intelligence modeling: The MI-CLAIM checklist. Nat Med. 2020;26(9):1320-4. doi:10.1038/s41591-020-1041-y

[31] Kapoor S, Narayanan A. Leakage and the reproducibility crisis in machine-learning-based science. Patterns (N Y). 2023;4(9):100804. doi:10.1016/j.patter.2023.100804

[32] Huang K, Fu T, Gao W, Zhao Y, Roohani Y, Leskovec J, et al. Artificial intelligence foundation for therapeutic science. Nat Chem Biol. 2022;18(10):1033-6. doi:10.1038/s41589-022-01131-2

[33] 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

[34] 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

[35] Nagendran M, Chen Y, Lovejoy CA, Gordon AC, Komorowski M, Harvey H, et al. Artificial intelligence versus clinicians: Systematic review of design, reporting standards, and claims of deep learning studies. BMJ. 2020;368. doi:10.1136/bmj.m689

[36] 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

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

[38] Mittelstadt B. Principles alone cannot guarantee ethical AI. Nat Mach Intell. 2019;1(11):501-7. doi:10.1038/s42256-019-0114-4

[39] 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

[40] 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