Causal Inference for Natural Product Therapeutics: Moving Beyond Association Toward Mechanism, Intervention Logic, and Translational Confidence

Charlotte Dupuis[1]*, Hugo Perrin[1], Elise Morel[2]

[1]Department of AI for Anti-aging and Skin-whitening Phytocosmetics, Faculty of Pharmacy, University of Lille, Lille, France.

[2]Department of Cheminformatics for Melanin Regulation, Faculty of Pharmacy, University of Montpellier, Montpellier, France.

ABSTRACT

Natural product therapeutics generate diverse evidence from phytochemical profiling, molecular docking, bioactivity screening, target prediction, network pharmacology, omics studies, traditional-use records, observational analyses, experimental models, and safety surveillance. These evidence streams are valuable for discovery but frequently support association or prediction rather than causal conclusions about mechanism or therapeutic benefit. This article proposes an original causal AI framework for organizing natural product evidence around explicit causal questions, intervention definitions, comparator strategies, outcomes, causal estimands, directed acyclic graphs, structural causal reasoning, potential outcomes, confounding control, mediation, effect modification, target-trial logic, experimental perturbation, and evidence triangulation. The framework separates hypothesis-generating signals from evidence capable of supporting progressively stronger mechanistic or intervention-oriented interpretations. It also introduces translational confidence criteria covering product identity, temporal ordering, exposure relevance, mechanistic plausibility, bias sensitivity, safety, replication, external validity, transportability, expert review, and transparent claim boundaries. Causal AI is positioned as a method for assumption mapping, evidence classification, counterfactual reasoning, validation planning, and uncertainty management rather than as an autonomous means of proving biological causation, clinical efficacy, or safety. The principal contribution is a structured pathway through which heterogeneous natural product data can be converted into testable causal hypotheses and validation priorities while preserving clear distinctions among association, prediction, causal inference, mechanistic evidence, experimental confirmation, and clinical effectiveness.

Key Words: *Causal inference, Causal AI, Natural product therapeutics, Phytopharmacology, Mechanism, intervention logic*

eIJPPR 2026; 16(2):53-69

HOW TO CITE THIS ARTICLE: Dupuis C, Perrin H, Morel E. Causal Inference for Natural Product Therapeutics: Moving Beyond Association Toward Mechanism, Intervention Logic, and Translational Confidence. Int J Pharm Phytopharmacol Res. 2026;16(2):53-69. https://doi.org/10.51847/6e7dEnHWlL

INTRODUCTION

Natural product therapeutics occupy an unusually heterogeneous evidentiary space in which botanical identity, phytochemical composition, biological activity, historical use, molecular prediction, experimental pharmacology, observational exposure, and clinical outcomes may all contribute to a single therapeutic narrative. This diversity creates discovery opportunities, but it also increases the risk that evidence generated for one purpose will be interpreted as supporting a stronger claim than it can justify. A compound–target prediction may be presented as target engagement, a pathway-enrichment result may be described as mechanism, or an observational exposure–outcome association may be interpreted as an intervention effect. Causal inference provides a disciplined alternative by requiring investigators to define the intervention, comparator, outcome, population, timing, and

Corresponding author: Charlotte Dupuis
Address: Department of AI for Anti-aging and Skin-whitening Phytocosmetics, Faculty of Pharmacy, University of Lille, Lille, France.
E-mail: ✉ charlotte.dupuis@univ-lille.fr
Received: 10 January 2026; **Revised:** 20 March 2026; **Accepted:** 28 March 2026

causal contrast before estimating or interpreting an effect. Target-trial emulation exemplifies this design-first principle by translating an observational effectiveness question into the protocol components of a hypothetical randomized trial, thereby making the intended intervention comparison and major design assumptions more transparent [1]. For natural product research, this orientation is especially important because the intervention is often chemically complex, preparation-dependent, dose-sensitive, and variable across sources, batches, formulations, and use contexts.

Natural products remain important sources of pharmacologically active compounds and therapeutic hypotheses, yet progression from discovery to credible translational evidence depends on more than chemical novelty or initial bioactivity. Reliable evaluation requires botanical authentication, reproducible preparation, phytochemical characterization, biologically relevant exposure, target and pathway validation, toxicity assessment, and evidence that the proposed intervention can be defined consistently across experimental and clinical settings [2]. These requirements complicate causal interpretation because a natural product name may refer to multiple non-equivalent preparations, each containing different constituent profiles and producing different exposures. Traditional-use information may help prioritize indications or identify historically relevant preparation contexts, but it does not establish comparative effectiveness or safety for a standardized intervention. Similarly, in vitro activity may identify a potentially responsive biological system while remaining disconnected from achievable concentrations, tissue distribution, metabolism, or human outcomes. The causal problem is therefore not limited to statistical confounding; it begins with the ontological question of what intervention is being evaluated and whether the experimental, observational, and clinical evidence refer to sufficiently comparable versions of that intervention.

Causal inference differs from ordinary prediction because it asks what would happen under alternative, explicitly defined interventions rather than merely estimating an outcome from observed features. Predictive models may accurately identify individuals with a high probability of an outcome while providing no valid estimate of how that outcome would change if a natural product were administered, withheld, reformulated, or compared with another therapy. Observational causal estimation consequently requires a defined estimand and assumptions concerning consistency, exchangeability, positivity, measurement, temporal ordering, and interference, together with methods appropriate to the treatment, outcome, and data structure [3]. In natural product therapeutics, these assumptions may be threatened by self-selection into supplement use, health-seeking behavior, indication severity, diet, concurrent medications, socioeconomic factors, product misclassification, adherence variation, batch heterogeneity, and differential outcome ascertainment. A causal analysis must therefore specify which variables are pre-exposure confounders, which may be mediators, which arise after treatment, and which may induce selection or collider bias if conditioned upon. Without this structure, increasing model complexity may improve prediction while leaving the intervention question unresolved.

Causal AI extends this reasoning by combining formal causal concepts with computational methods for representing assumptions, integrating heterogeneous evidence, estimating counterfactual contrasts, exploring effect heterogeneity, and prioritizing validation. Its value lies not in replacing experimental pharmacology or clinical evaluation, but in making the inferential pathway from association to action more explicit. Counterfactual and causal machine-learning approaches can help distinguish questions about prognosis from questions about intervention effects, thereby supporting more actionable analysis when identification assumptions are defensible [4]. For natural product therapeutics, an appropriately bounded causal AI system could classify evidence by inferential strength, construct and revise causal diagrams, identify potential confounding and selection mechanisms, map candidate mediators, compare alternative intervention definitions, and highlight where experimental perturbation or target-trial emulation is needed. The objective of this article is therefore to propose a structured causal AI framework that links natural product identity, phytochemical evidence, bioactivity, target and pathway hypotheses, omics, observational data, experimental validation, safety information, and translational criteria without treating computational outputs as independent proof of biological causation or therapeutic efficacy.

Causal thinking in natural product therapeutics

Causal thinking begins by representing a therapeutic question as a relation among interventions, mechanisms, outcomes, populations, and assumptions rather than as a search for statistically associated variables. Within a structural causal perspective, biological and clinical variables are understood as components of a data-generating system in which interventions alter selected variables and propagate effects through directed relations. Causal representation learning seeks representations that reflect underlying generative factors and remain informative under interventions or distributional change, rather than encoding only associations observed in a particular dataset [5]. Applied to natural product therapeutics, this distinction means that a model should not

merely learn that a phytochemical profile co-occurs with a gene-expression signature or phenotype. It should instead clarify whether the profile defines an intervention, whether exposure precedes the response, which variables plausibly transmit the effect, which common causes require control, and whether the proposed relations would remain meaningful under a deliberate change in dose, formulation, target activity, or treatment assignment. Such representations can strengthen hypothesis organization, but they remain conditional on data quality, causal assumptions, biological knowledge, and independent validation.

Potential-outcomes reasoning provides a complementary language for defining causal effects. For an individual or experimental unit, the causal question compares outcomes that would occur under alternative intervention strategies, although only one outcome is ordinarily observed. The estimand may concern an average effect, a conditional effect in a biologically defined subgroup, a direct effect not operating through a particular mediator, or another population-level contrast. Causal machine learning can support estimation of heterogeneous treatment effects and flexible adjustment in high-dimensional healthcare data, but it does not remove the need for exchangeability, positivity, consistency, appropriate temporal alignment, and reliable outcome measurement [6]. In natural product research, treatment versions must be specified with particular care. A crude extract, purified constituent, standardized fraction, dietary exposure, encapsulated supplement, and metabolite may not represent interchangeable interventions. Comparator definitions are equally consequential: nonuse, placebo, standard care, an active comparator, and an alternative dose answer different questions and may involve different confounding structures. The outcome must also be temporally and biologically aligned with the proposed mechanism, whether it is target engagement, a pathway response, a functional phenotype, an adverse event, or a clinical endpoint.

Causal AI should consequently be understood as an evidence-reasoning architecture rather than a single algorithm. It may incorporate graphical models, structural causal models, potential-outcomes estimators, causal discovery, mediation analysis, effect-modification assessment, target-trial emulation, sensitivity analysis, and transportability methods, but no component can guarantee causal validity in isolation. The broader future of causal inference depends on stronger integration of design, substantive knowledge, data quality, transparent assumptions, and methods capable of addressing complex longitudinal and heterogeneous systems [7]. This is directly relevant to natural product therapeutics, where data may span chemical assays, cell systems, animal models, omics platforms, electronic records, spontaneous safety reports, and clinical investigations. A causal AI framework must preserve the distinctions among these evidence-generating environments rather than pooling them as though they estimate the same quantity. Experimental systems may support perturbational mechanism evidence yet have limited external validity; observational data may represent real-world use while remaining vulnerable to confounding; and clinical trials may estimate intervention effects for one standardized preparation without validating every constituent-level mechanism.

Directed acyclic graphs offer a practical means of expressing these distinctions before analysis. A graph can encode hypothesized relations among product selection, phytochemical exposure, baseline disease severity, concurrent treatment, target engagement, pathway response, mediators, outcomes, measurement processes, and study participation. Its primary function is not graphical decoration but identification of causal assumptions, candidate adjustment sets, mediators that should not be controlled when estimating total effects, and colliders whose conditioning may introduce bias. Biological model selection should follow the causal question and domain knowledge rather than relying solely on automated variable selection, predictive fit, or statistical significance [8]. Structural causal models can extend the graph by specifying functional relations and intervention operations, whereas potential-outcomes notation can define the estimand associated with a particular intervention contrast. Together, these approaches provide a language for asking whether evidence supports association, causal identification, mechanism, or transportability. They do not, however, make assumptions true. The validity of the resulting analysis still depends on the defensibility of the graph, the quality and timing of measurements, the comparability of treatment versions, the availability of relevant confounders, and confirmation through appropriately designed experiments or clinical studies.

Limits of associational discovery

Molecular docking and computational target prediction are valuable early-stage tools because they reduce a large chemical and biological search space to a manageable set of hypotheses. Docking can propose binding orientations, approximate complementarity, and relative rankings under a selected receptor and scoring model, but its outputs remain sensitive to receptor flexibility, protonation, solvent representation, sampling, and scoring limitations. A docking pose or favorable score therefore does not establish direct target engagement, functional modulation, pathway mediation, organism-level benefit, or safety [9].

Target-prediction platforms similarly infer candidate proteins from ligand similarity, chemical features, or known interaction patterns. These predictions can guide assay selection and experimental prioritization, but their reliability depends on the composition of reference datasets, chemical-domain coverage, target annotation, and the similarity between the queried natural product constituent and previously characterized ligands [10]. Even when docking and target prediction converge on the same protein, the convergence may reflect shared database structure or modeling assumptions rather than independent causal evidence. Orthogonal binding assays, cellular target-engagement methods, functional perturbation, selectivity assessment, and rescue experiments are needed before a predicted interaction can contribute to a mechanistic causal claim.

Network pharmacology and pathway enrichment broaden the analysis from individual targets to connected biological systems, which is attractive for multicomponent natural products and polypharmacological hypotheses. However, a compound–target edge, target–disease link, network-proximity score, or highly connected node is ordinarily an associational representation assembled from databases, literature, predicted interactions, and curated annotations. Network pharmacology becomes more causally informative only when relationships are directionally specified, temporally plausible, linked to a defined intervention, and tested through perturbation or other mechanism-sensitive evidence [11]. Pathway-enrichment analysis poses a related challenge. An enriched gene set can indicate that genes associated with an exposure or phenotype are overrepresented in a defined biological process, but it cannot determine whether the pathway is activated or inhibited, whether it lies upstream or downstream of the observed phenotype, or whether it mediates an intervention effect [12]. Results may vary with the background universe, database version, gene-set overlap, thresholding, multiple-testing procedure, and composition of the input list. Accordingly, pathway and network results should be used to prioritize causal hypotheses and validation experiments rather than to label a mechanism as established.

Transcriptomic, proteomic, metabolomic, and microbiome evidence can reveal exposure-responsive signatures and connect multiple levels of biological organization, yet cross-sectional or correlational integration does not establish causal direction. Multi-omics approaches may strengthen systems-level interpretation by relating molecular layers, but they remain affected by batch variation, cellular composition, tissue selection, temporal mismatch, missingness, platform-specific measurement error, and incomplete coverage [13]. A transcript may change because it mediates an effect, compensates for an effect, reflects toxicity, or merely correlates with another causal process. Protein abundance may not represent activity; metabolite concentrations may be shaped by diet, circadian timing, organ function, or microbial metabolism; and microbiome associations may reflect compositional constraints, geographic differences, concurrent medication, or host behavior. Reverse causation is particularly problematic when the disease state alters product use, metabolism, diet, or microbial composition. Causal interpretation therefore requires longitudinal ordering, relevant exposure levels, replicated signatures, mechanistically justified mediators, and perturbation capable of distinguishing competing explanations. Omics evidence can reinforce a causal hypothesis when aligned with target engagement, temporal response, and functional validation, but it cannot substitute for those elements.

Observational exposure–outcome studies, traditional-use records, bioactivity screens, adverse-event reports, literature co-occurrence, and knowledge-graph links contribute additional forms of association. Observational studies may capture real-world product use but can be distorted by healthy-user behavior, confounding by indication, dietary patterns, socioeconomic differences, concurrent therapies, exposure misclassification, selective participation, missing follow-up, and outcome-detection differences. Traditional use can provide cultural and historical context for hypothesis generation, but heterogeneous preparations and selective documentation prevent it from establishing standardized efficacy or safety. Bioactivity screening demonstrates activity only under the tested assay conditions and may be affected by nonspecific cytotoxicity, aggregation, fluorescence interference, or concentrations that are not achievable in vivo. Natural-product pharmacology therefore requires authenticated source material, chemical characterization, relevant dosing, appropriate controls, reproducible assays, and cautious interpretation of experimental findings [14]. Adverse-event signals warrant investigation but are not causal risk estimates, while literature co-occurrence and knowledge-graph proximity may reproduce publication frequency, annotation bias, or circular evidence. **Table 1** summarises common associational evidence types in natural product therapeutics and the causal claims they cannot establish without additional validation.

Table 1. Limits of Associational Discovery in Natural Product Therapeutics: Docking, Target Prediction, Network Pharmacology, Pathway Enrichment, Omics Associations, Observational Signals, Traditional Use, Bioactivity Screening, Literature Co-Occurrence, and Causal Interpretation Boundaries

Associational evidence type	Typical natural product use	What it can support	What it cannot establish	Causal threat	Validation need	Interpretation boundary	Framework role
Docking association	Virtual screening, binding-pose generation and constituent prioritization	Plausible ligand orientation and comparative ranking under a specified model	Direct target engagement, functional modulation, pathway mediation, safety or therapeutic efficacy	Scoring error, receptor rigidity, protonation assumptions, solvent simplification and sampling limitations	Orthogonal binding, cellular engagement, counterscreening and functional perturbation	A docking result is a computational binding hypothesis	Computational target-hypothesis input
Target prediction	Nomination of proteins potentially interacting with a phytochemical	A prioritized candidate-target set for testing	Direct binding, target necessity, target sufficiency or intervention benefit	Training-data bias, ligand-similarity bias, incomplete target coverage and chemical-domain shift	Biophysical confirmation, engagement assays, genetic or pharmacological perturbation and rescue	Predicted targets remain provisional	Target-evidence classification and validation planning
Network pharmacology edge	Construction of compound–target–pathway–disease networks	Systems-level hypotheses and possible polypharmacological relations	Directionality, causal mediation, intervention effects or validated mechanism	Database incompleteness, annotation bias, degree bias, publication bias and circular evidence	Directed causal modeling, temporal evidence and perturbation of candidate nodes or edges	Network connectivity is not causal influence	Mechanism-hypothesis mapping
Pathway enrichment	Interpretation of genes or proteins associated with exposure or phenotype	Prioritization of overrepresented biological processes	Pathway activation, pathway mediation, direction of effect or therapeutic benefit	Multiple testing, overlapping gene sets, background selection and annotation bias	Targeted pathway assays, time-course evaluation, perturbation and rescue	Enrichment denotes statistical overrepresentation	Pathway-hypothesis input
Transcriptomic association	Differential-expression profiling after exposure	Exposure-responsive signatures and candidate regulatory processes	Necessity, sufficiency, causal mediation or clinical effect	Batch effects, cell-composition differences, temporal ambiguity and reverse causation	Replicated dose–time studies, perturbation, rescue and orthogonal assays	Expression change may be cause, consequence or correlate	Temporally qualified omics evidence
Proteomic association	Detection of abundance or modification changes	Candidate proteins, response signatures and pathway hypotheses	Functional activity, direct target engagement or	Limited coverage, abundance–activity mismatch, degradation	Activity assays, engagement testing, knockdown, overexpressi	Altered abundance is not equivalent to causal function	Target and pathway evidence

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

			therapeutic mechanism	and sample heterogeneity	on and replication		
Metabolomic association	Identification of exposure-related metabolic signatures	Candidate biomarkers, metabolic responses and exposure indicators	Metabolic mediation, pathway necessity or efficacy	Diet, circadian timing, organ function, sample handling and reverse causation	Longitudinal sampling, isotope tracing, enzyme perturbation and external replication	A metabolite association is not automatically a mediator	Exposure and response evidence
Microbiome association	Relating natural-product use to microbial taxa, functions or metabolites	Host–microbe hypotheses and candidate modifiers or mediators	Microbiome-mediated therapeutic effect or organism-specific causation	Compositionality, diet, geography, medication use, antibiotics and batch effects	Longitudinal studies, functional profiling, metabolite validation and controlled perturbation	Taxonomic correlation does not establish mediation	Contextual mediator or effect-modifier evidence
Observational exposure–outcome association	Evaluation of supplements, herbs or dietary exposures against health outcomes	Real-world association under observed conditions	Intervention benefit without identification assumptions	Confounding by indication, healthy-user bias, selection, exposure misclassification and reverse causation	Explicit causal design, appropriate comparator, confounding control, sensitivity analysis and replication	Causal interpretation remains conditional on design and assumptions	Observational causal-assessment input
Traditional use association	Prioritization of indications, preparation practices or cultural contexts	Historical context, plausibility and candidate-use selection	Standardized efficacy, comparative effectiveness, dose response or safety	Survivorship, selective documentation, heterogeneous preparations and reporting bias	Authentication, characterization, pharmacology, toxicology and clinical evaluation	Traditional use is contextual evidence rather than therapeutic proof	Evidence-foundation context
Bioactivity screening association	Identification of active extracts, fractions or compounds	Activity under a defined assay condition and concentration	In vivo exposure, selectivity, target mechanism, safety or clinical efficacy	Assay interference, aggregation, nonspecific cytotoxicity and concentration artifacts	Counterscreens, potency–selectivity profiling, exposure assessment and mechanism testing	An assay hit is conditional bioactivity evidence	Bioactivity-evidence layer
Adverse-event association	Signal detection from reports, records or claims	Identification of a potential safety concern requiring investigation	Causal attribution, incidence or comparative risk	Underreporting, stimulated reporting, notoriety bias, channeling and missing denominators	Comparator-based studies, temporal analysis, specialist review, sensitivity testing and replication	A safety signal is not a causal risk estimate	Safety and uncertainty review

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

Literature co-occurrence	Text mining of compound, target, pathway and disease terms	Rapid evidence retrieval and hypothesis generation	Biological interaction, mechanism or interventio n effect	Publication bias, popularity bias, duplicated claims and semantic ambiguity	Primary-source verification and independent experimental testing	Semantic proximity is not biological causality	Discovery and retrieval support
Knowledge-graph association	Link prediction, path ranking and candidate prioritizatio n	Integrative hypotheses across heterogeneous sources	Directional causal effects, validated mechanism , safety or efficacy	Inherited source errors, data leakage, circularity and uncertain provenance	Source-level audit, temporal validation, causal modeling and experimental confirmation	A graph score is not a causal estimate	Traceable hypothesis-generation layer

Mechanism and intervention logic

Mechanism evaluation should begin with a precisely formulated causal question rather than with an unconstrained search for molecular associations. The intervention must identify the natural product preparation, relevant constituent profile, dose, route, frequency, duration, adherence conditions, and exposure context, while the comparator must represent a meaningful alternative such as nonuse, placebo, standard care, an active comparator, or another formulation. Outcomes must be defined at an appropriate biological or clinical level and measured within a time window consistent with the proposed mechanism. A directed acyclic graph can then represent hypothesized relationships among baseline determinants, treatment selection, phytochemical exposure, target engagement, pathway activity, mediators, safety events, and outcomes. DAGs can improve confounder selection and reveal inappropriate adjustment for mediators or colliders, but their utility depends on the accuracy and completeness of the causal assumptions encoded by investigators [15]. They should therefore be treated as transparent models of current knowledge rather than as automatically valid descriptions of biological reality.

Causal discovery methods may supplement expert-designed diagrams by identifying conditional-independence patterns, candidate directions, or alternative structures in high-dimensional biological and clinical data. However, graphical discovery algorithms depend on assumptions concerning acyclicity, faithfulness, causal sufficiency, measurement quality, and the form of the data-generating process. Distinct causal structures may also be observationally equivalent, while latent confounding and selection can create misleading edges. Causal discovery should consequently generate testable structural hypotheses rather than confirm a biological mechanism independently [16]. Mediation analysis addresses a different question: whether and to what extent an intervention effect may operate through a specified intermediate variable, such as target engagement, pathway activity, a metabolite, or an inflammatory response. Valid mediation interpretation requires temporal ordering, explicit direct and indirect effect definitions, control of relevant exposure–outcome, exposure–mediator, and mediator–outcome confounding, and careful consideration of post-treatment confounders [17]. A biomarker that changes after exposure should not be designated a mediator merely because it is statistically associated with both the intervention and outcome.

Effect modification further distinguishes whether a causal contrast varies across biological, clinical, chemical, or environmental contexts. In natural product therapeutics, heterogeneity may arise from metabolic phenotype, microbiome composition, disease severity, genotype, concurrent treatment, formulation, baseline nutritional status, or differences in achievable exposure. Such variation must be defined relative to a particular causal estimand and measurement scale rather than inferred from unplanned subgroup comparisons [18]. Intervention logic must also preserve temporal alignment among eligibility, treatment assignment, comparator status, follow-up, and outcome measurement. When suitable longitudinal observational data are available, target-trial emulation can clarify these elements and reduce biases created by inconsistent time zero, immortal follow-up, treatment switching, or poorly defined eligibility [19]. Nevertheless, an emulation remains dependent on measured confounding, adequate overlap, valid treatment classification, and appropriate analytical assumptions. It cannot recreate randomization when major common causes of treatment and outcome are unavailable or poorly measured.

Mechanistic confidence increases when chemical plausibility, target evidence, temporal pathway response,

intervention specificity, and experimental perturbation converge. Computational target nomination may identify a candidate protein, but stronger evidence requires orthogonal binding or engagement assays, functional modulation, selectivity assessment, dose–response consistency, and experiments testing whether target inhibition, activation, deletion, or rescue changes the predicted phenotype. Natural-product target-identification strategies therefore gain credibility when complementary computational, biochemical, cellular, and functional methods converge rather than when a single method is treated as decisive [20]. Evidence triangulation should extend this logic across omics, network pharmacology, observational data, controlled experiments, safety evidence, and clinical studies where available. Convergence is most informative when the contributing approaches have different bias structures and independent data provenance. Contradictions should be preserved as evidence gaps rather than removed through selective synthesis. **Figure 1** illustrates how causal thinking can separate association, mechanism hypotheses, intervention logic, confounding, bias, and translational interpretation in natural product therapeutics.

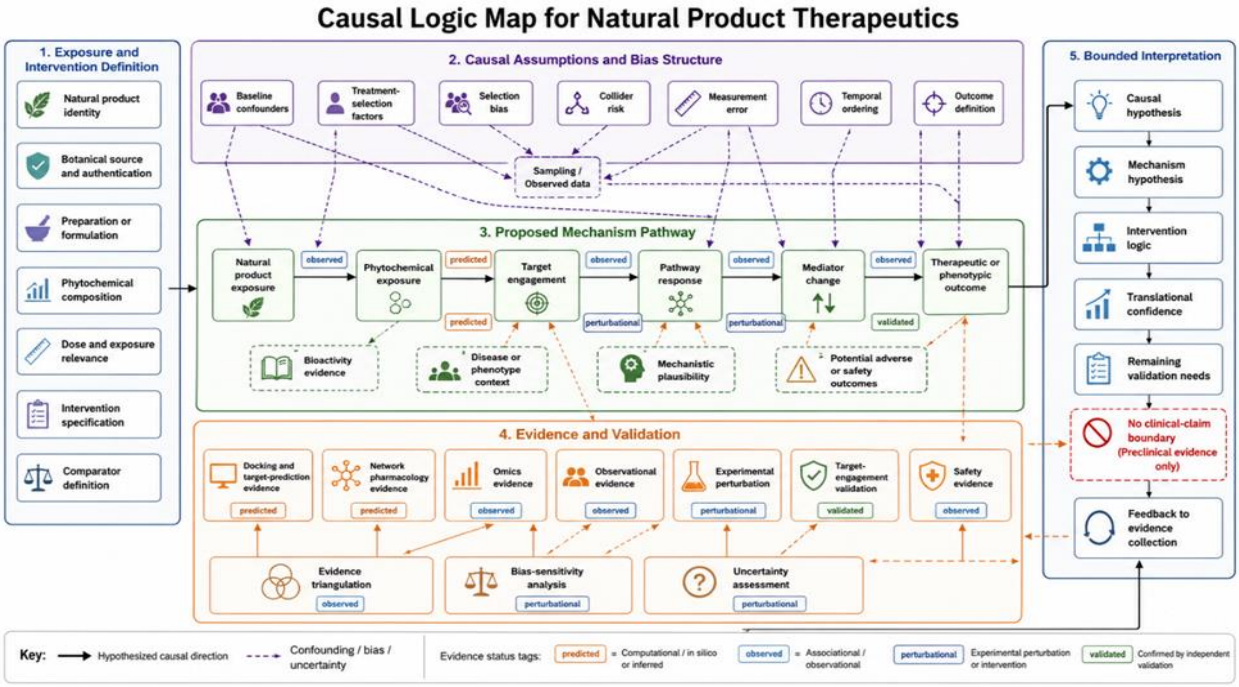

Figure 1. Causal Logic Map for Natural Product Therapeutics

Alt-text description

A conceptual causal logic map showing natural product exposure connected to phytochemical composition, target engagement, pathway response, mediators, confounders, safety signals, therapeutic outcomes, validation evidence, and causal interpretation boundaries.

Translational confidence criteria

Translational confidence should be assigned to a bounded causal claim rather than to a natural product, computational model, or evidence collection in the abstract. The first criterion is a clear causal estimand specifying the population, intervention, comparator, outcome, follow-up period, and effect measure under consideration. Confidence also depends on explicit assumptions about consistency, treatment versions, positivity, exchangeability, interference, missingness, and measurement. When observational estimates contribute to the claim, the analysis should describe the measured confounders, justify the adjustment set, and evaluate whether plausible unmeasured confounding could explain the observed association. Sensitivity analysis can quantify how strong an unmeasured confounder would need to be to alter a result, but it does not eliminate confounding or verify that the assumed bias structure is correct [21]. Robustness should therefore be interpreted alongside the causal diagram, study design, negative controls where feasible, alternative specifications, and the plausibility of omitted common causes.

External validity and transportability require separate evaluation because internally credible evidence may not apply to a different population, preparation, dose, healthcare context, or disease stage. A natural product tested as a standardized extract in a selected research population may not be equivalent to commercially available preparations or traditional formulations.

Differences in metabolism, concurrent medication, nutritional status, microbiome composition, genetic background, baseline risk, treatment access, and outcome ascertainment can modify intervention effects or alter the meaning of exposure. Transportability methods make the source population, target population, overlap conditions, measured effect modifiers, and required assumptions explicit [22]. Where formal transport estimation is infeasible, a structured qualitative assessment should still identify which features are expected to remain invariant and which require new evidence. Claims should remain limited to the populations, preparations, doses, and settings supported by the available evidence rather than being generalized from chemical similarity or broad botanical naming alone.

Safety evidence must be integrated into causal confidence rather than appended after efficacy-oriented reasoning. Spontaneous reports, electronic health records, claims databases, literature cases, toxicology studies, and interaction screens can identify possible harms, but safety signals are often vulnerable to underreporting, stimulated reporting, channeling, notoriety bias, exposure uncertainty, and missing denominators. Causal machine-learning approaches may support pharmacovigilance by improving adjustment, pattern detection, and heterogeneous-risk evaluation, yet detected associations still require clinically informed review and validation [23]. Translational confidence should be reduced when real-world studies use avoidable design choices that introduce time-related bias, inappropriate comparators, uncontrolled treatment switching, conditioning on post-exposure variables, or poorly justified covariate selection [24]. Safety interpretation should therefore include exposure timing, dechallenge and rechallenge information where available, plausible mechanisms, alternative causes, dose relationships, interaction potential, vulnerable populations, and replication across independent sources. A therapeutic hypothesis cannot be considered translationally mature when major safety uncertainties remain unresolved.

Longitudinal evidence creates additional challenges when treatment, confounders, mediators, and outcomes change over time. Standard regression may be inadequate when previous exposure affects a time-varying confounder that subsequently influences later exposure and outcome. G-formula and marginal structural approaches can address specified forms of treatment–confounder feedback under assumptions of consistency, positivity, exchangeability, and correct model specification [25]. More broadly, observational evidence becomes more causally informative when investigators begin with a hypothetical intervention protocol and align eligibility, treatment assignment, comparator status, follow-up, and analysis accordingly [26]. Translational confidence should therefore integrate temporal ordering, intervention specificity, biologically plausible exposure, replicated evidence, experimental perturbation, bias sensitivity, external validity, transportability, safety, and clinical evidence where available. Expert review is required to challenge causal assumptions, evaluate product comparability, assess toxicological plausibility, and calibrate the permissible claim. **Table 2** presents translational confidence criteria for evaluating causal claims in natural product therapeutics.

Table 2. Translational Confidence Criteria for Causal Claims in Natural Product Therapeutics: Causal Estimands, Explicit Assumptions, Confounding Control, Temporal Ordering, Mechanistic Plausibility, Intervention Specificity, Experimental Perturbation, Safety Evidence, External Validity, and Claim Boundaries

Confidence criterion	Core translational question	Evidence needed	Causal interpretation role	Failure risk	Required validation	Expert review need	Claim boundary
Clear causal estimand	What effect is being estimated, for whom, under which intervention contrast and over what period?	Population, intervention, comparator, outcome, timing and effect-measure definitions	Distinguishes a causal effect from association or prediction	Different analyses may answer incompatible questions	Estimand–design–analysis consistency check	Essential	No causal-effect claim without a defined estimand
Explicit assumptions	Which consistency, exchangeability, positivity, interference	Written assumptions and causal structure	Makes identification conditions transparent	Hidden assumptions can create unjustifi	Assumption audit and alternative causal models	Essential	Conclusions remain conditional on stated assumptions

	and measurement assumptions are required?			ed certainty			
Confounding control	Have common causes of intervention and outcome been identified and addressed?	DAG-informed baseline variables, design restrictions and diagnostics	Supports conditional exchangeability	Residual or unmeasured confounding	Balance checks, negative controls where feasible and sensitivity analysis	Essential	Inadequate control limits interpretation to association
Temporal ordering	Does exposure precede the mediator and outcome within a plausible interval?	Time-stamped intervention, mediator and outcome measurements	Supports directional interpretation	Reverse causation and time-window bias	Longitudinal design, lag analysis and prespecified time zero	High	Cross-sectional evidence cannot establish intervention direction
Mechanistic plausibility	Is the proposed mechanism compatible with chemistry, exposure, target engagement and biological function?	Chemical, pharmacokinetic, target, pathway and phenotypic evidence	Strengthens coherence of the causal hypothesis	Plausibility may substitute for direct testing	Orthogonal assays and perturbational confirmation	Essential	Plausibility alone is not therapeutic proof
Intervention specificity	Is the tested preparation sufficiently characterized and reproducible?	Authentication, phytochemical fingerprint, formulation, batch and dose data	Supports consistency and replication	Product heterogeneity, substitution or adulteration	Independent analytical characterization and batch control	Essential	Uncharacterized products cannot support product-specific claims
Dose or exposure relevance where supported	Are tested concentrations compatible with achievable biological exposure?	Dose–response, bioavailability, metabolism and tissue-exposure evidence	Connects experimental effects with feasible intervention conditions	Supra-physiological concentrations and exposure mismatch	Pharmacokinetic and exposure–response assessment	Essential for translational claims	Implausible exposure supports only a limited preclinical hypothesis
Experimental perturbation	Does deliberate manipulation of the intervention, target or pathway alter	Controlled exposure, target engagement, genetic or pharmacological perturbation and rescue	Tests functional necessity, sufficiency and pathway response	Off-target effects and model dependence	Orthogonal perturbations, appropriate controls and replication	Essential	Association-only evidence cannot establish mechanism

	the predicted outcome?						
Replicated evidence	Are findings reproduced across datasets, laboratories, assays or populations?	Independent replication with comparable definitions	Reduces chance findings and setting-specific artifacts	Overlapping data may be mistaken for independent confirmation	Provenance audit and external replication	High	Single-study evidence receives limited confidence
Safety evidence	Are adverse effects, interactions, vulnerable groups and exposure limits adequately assessed?	Toxicology, interaction studies, pharmacovigilance and human safety evidence where available	Balances possible benefit against harm	Underreporting, short follow-up and missing interaction data	Comparative safety studies and causal signal assessment	Essential	No clinical recommendation without adequate safety evidence
Bias sensitivity analysis	Would plausible unmeasured or measurement bias materially alter the conclusion?	Quantitative or structured qualitative bias parameters	Assesses robustness of causal interpretation	Sensitivity analyses may be omitted or overinterpreted	Multiple plausible scenarios and transparent assumptions	High	Fragile estimates receive low causal confidence
External validity	Does the study context resemble the intended use context?	Population, formulation, care setting and implementation information	Assesses applicability beyond the original study	Highly selected samples and artificial conditions	External cohort or pragmatic validation	High	Findings remain bounded to studied conditions
Transportability	Can evidence be extended across populations, preparations, doses or systems?	Effect modifiers, overlap, product comparability and source–target differences	Formalizes movement from source to target context	Unmeasured differences and positivity failure	Transport analysis or structured applicability assessment	Essential when extending claims	No extrapolation beyond supported settings
Clinical evidence where available	Do human intervention data support the stated therapeutic question?	Well-designed trials or credible causal observational studies	Provides the most direct evidence of human effectiveness	Small, biased or surrogate-only studies	Independent confirmation and clinically relevant outcomes	Essential	Preclinical evidence is not described as clinical efficacy
No clinical-claim boundary	What level of language is justified by the integrated evidence?	Claim-specific evidence grade, uncertainty and safety assessment	Prevents category errors across discovery stages	Premature efficacy or safety language	Final multidisciplinary claim audit	Mandatory	Framework outputs do not authorize diagnosis or treatment

Proposed causal AI framework

The proposed framework begins with a natural product evidence foundation that treats identity and provenance as causal prerequisites rather than descriptive metadata. Inputs include botanical source, taxonomic authentication, preparation method, phytochemical composition, formulation, batch information, bioactivity evidence, target predictions, network associations, omics findings, observational data, experimental studies, safety evidence, and traditional-use context where adequately documented. The framework classifies each input as association, prediction, causal hypothesis, causal estimate, mechanistic evidence, experimental validation, safety evidence, or clinical evidence. This classification prevents a high-dimensional evidence graph from flattening sources with different inferential meanings. Causal AI may assist by extracting structured relations, identifying conflicting intervention definitions, tracing evidence provenance, and detecting where a claim depends repeatedly on the same underlying database or study. Causal models can sometimes improve robustness by encoding relationships that remain meaningful under changing conditions, but their accuracy depends on the validity of the assumed structure and the context in which it is applied [27]. Accordingly, the framework does not infer that a graph representation is biologically correct simply because it improves prediction.

The second stage defines the causal question and estimand. Investigators specify the target population, intervention version, comparator, outcome, timing, adherence conditions, and effect measure, together with the intended use of the result. A causal structure module then represents assumptions through a DAG and, where appropriate, a structural causal model. Variables are labeled as baseline causes, treatment determinants, mediators, outcomes, selection variables, measurement processes, effect modifiers, or contextual factors. The framework identifies candidate confounders, warns against conditioning on colliders or post-treatment consequences, and records alternative plausible diagrams when knowledge is incomplete. For suitable longitudinal data, causal representation or treatment-effect models may help estimate potential outcomes or conditional effects, but these estimates remain vulnerable to unmeasured confounding, insufficient overlap, treatment misclassification, model misspecification, and inadequate benchmarking [28]. Model output is therefore accompanied by an identification statement that distinguishes what is estimated statistically from what is identified causally under the stated assumptions.

The third stage evaluates mechanism and intervention logic. Candidate paths from natural product exposure to target engagement, pathway response, mediator change, phenotype, outcome, and adverse effects are ranked according to evidence type, temporal coherence, exposure relevance, and perturbational support. Causal discovery may propose alternative structures, but these are labeled as hypotheses requiring biological review. Mediation modules assess whether a proposed intermediate can be evaluated under appropriate timing and confounding assumptions, while effect-modification modules examine whether causal contrasts may differ across product formulations, metabolic profiles, disease subtypes, or other prespecified contexts. Target-trial logic is activated when the question concerns comparative effectiveness or safety in observational data. The overall workflow follows a staged causal roadmap connecting question formulation, estimand definition, causal assumptions, identification, estimation, diagnostics, sensitivity analysis, and bounded interpretation [29]. Each stage records unresolved assumptions and directs them toward experimental, analytical, or data-collection priorities rather than concealing them within a composite score.

The fourth stage triangulates evidence and grades causal confidence. Computational evidence, omics data, network models, observational estimates, experimental perturbations, target-engagement studies, safety data, and clinical evidence are compared according to their independence, quality, temporal relevance, product comparability, and bias structure. A confidence grade is assigned to a specific claim, such as "compound X may engage target Y under the tested conditions" or "preparation Z is associated with outcome O in the studied population," rather than to the product as a whole. Causal language about an intervention is permitted only when the question, contrast, assumptions, and limitations are explicitly reported [30]. The confidence record describes supporting evidence, contradictory findings, untested alternatives, potential confounding, selection mechanisms, measurement uncertainty, sensitivity results, external-validity concerns, and remaining validation needs. Grades are qualitative or rule-based summaries of evidence maturity; they are not causal scores, probabilities of efficacy, or substitutes for experimental or clinical decisions.

The final stage integrates safety review, expert oversight, validation planning, and feedback updating. Natural product chemists assess identity and intervention equivalence; pharmacologists evaluate exposure, target engagement, and functional plausibility; toxicologists review adverse effects and interactions; causal methodologists examine identification and analysis; and clinical experts evaluate endpoint relevance and applicability. Causal machine learning may support heterogeneous-effect assessment and counterfactual prediction, but reliable use requires overlap, validation,

uncertainty estimation, and careful separation of treatment-effect modeling from ordinary prognostic prediction [31]. The framework converts unresolved uncertainty into a prioritized validation plan, which may include chemical recharacterization, binding confirmation, dose–response experiments, target perturbation, mediation testing, longitudinal data collection, comparator redesign, bias sensitivity analysis, external replication, or clinical evaluation. New evidence updates the causal diagram, evidence classifications, confidence judgment, and claim boundary. The framework remains explicitly non-clinical unless sufficient clinical evidence exists for the particular standardized intervention and bounded outcome. **Table 3** presents the proposed causal AI framework for natural product therapeutics.

Table 3. Proposed Causal AI Framework for Natural Product Therapeutics: Evidence Foundation, Causal Question Definition, DAG Construction, Bias Mapping, Mechanism Hypothesis, Intervention Logic, Evidence Triangulation, Causal Confidence Grading, Safety Review, Validation Planning, Feedback, and Claim Boundaries

Framework stage	Core purpose	Required input	Causal function	Potential output	Validation need	Failure risk	Feedback mechanism	Decision boundary
Natural product evidence foundation	Establish the identity, preparation and evidence under evaluation	Botanical identity, provenance, phytochemical profile, formulation, dose, bioactivity, targets, omics, observational, experimental and safety evidence	Defines the intervention and evidence provenance	Versioned natural product evidence dossier	Authentication, analytical characterization and provenance audit	Misidentification, contamination, batch heterogeneity or non-equivalent preparations	Update when source, batch, formulation or assay changes	No product-specific claim without adequate identity
Causal question definition	Convert a broad therapeutic claim into an answerable causal question	Population, intervention, comparator, outcome, timing and decision context	Defines the causal contrast and estimand	Structured causal-question record	Pharmacological, clinical and methodological review	Vague questions produce uninterpretable results	Revise after feasibility and stakeholder review	No causal analysis before the question is explicit
Evidence classification	Separate association, prediction, causal hypothesis, causal estimate and validation evidence	Retrieved evidence with study-design and provenance metadata	Prevents inferential category errors	Evidence-classification map	Independent source verification	Computational signals may be mislabeled as causal findings	Reclassify when stronger evidence becomes available	Evidence retains the weakest justified interpretation
DAG construction	Encode assumed causal relations	Temporal ordering, domain knowledge, empirical findings and candidate variables	Maps confounders, mediators, colliders, selection and outcomes	Versioned DAG with stated rationale	Expert review and alternative-DAG analysis	Missing variables, incorrect arrows or false certainty	Update when new evidence changes assumptions	A DAG is an assumption model, not causal proof
Confounder and bias mapping	Identify threats to causal identification	DAG, recruitment, measurement, missingness, treatment and outcome processes	Determines adjustment and sensitivity requirements	Bias register and proposed adjustment set	Balance diagnostics, negative controls where feasible and sensitivity analysis	Residual confounding, overadjustment, collider bias or selection bias	Revise when diagnostics reveal incompatibility	Major unresolved bias lowers or blocks causal grading
Mechanism hypothesis evaluation	Map compounds, targets,	Docking, target prediction, binding,	Constructs testable	Directed mechanism map with	Engagement assays, temporal	Associational edges may be	Promote or demote paths	Mechanism remains hypothetical

	pathways, mediators and outcomes	bioactivity, omics, network and experimental evidence	mechanism hypotheses	evidence labels	testing, perturbation and rescue	mistaken for mechanisms	after validation	without functional confirmation
Intervention logic specification	Define what is manipulated and what it is compared with	Preparation, dose, route, duration, adherence, comparator and outcome	Connects mechanism hypotheses to counterfactual strategies	Intervention protocol or target-trial specification	Exposure, feasibility and comparator review	Multiple treatment versions or inappropriate comparators	Update when formulation or exposure knowledge changes	No intervention effect for an undefined treatment
Evidence triangulation	Compare evidence streams with different assumptions and biases	Computational, molecular, omics, observational, experimental, safety and clinical evidence	Evaluates convergence and contradiction	Triangulation matrix and contradiction log	Independence, quality and provenance assessment	Correlated evidence may mimic independent replication	Reassess when evidence streams conflict or replicate	Evidence volume cannot replace causal quality
Causal confidence grading	Calibrate confidence in each bounded claim	Estimand, assumptions, bias profile, mechanism, perturbation, replication and transportability evidence	Converts synthesis into a transparent judgment	Claim-specific confidence grade and rationale	Reproducible grading rules and independent review	Composite grades may imply false precision	Revise after replication, validation or contradiction	A grade summarizes evidence and does not prove efficacy
Safety and uncertainty review	Integrate harm evidence and inferential fragility	Toxicology, interactions, pharmacovigilance, exposure and sensitivity evidence	Evaluates causal safety hypotheses and uncertainty	Safety register and uncertainty profile	Signal validation, bias analysis and specialist review	Underreported harms or interactions	Update with every material safety signal	Serious unresolved safety concerns stop escalation
Expert review	Challenge assumptions, classifications and claims	Evidence dossier, DAGs, analyses, confidence rationale and uncertainty	Adds domain knowledge and detects category errors	Documented review decision and revisions	Multidisciplinary review and conflict disclosure	Authority may substitute for evidence	Record disagreement and trigger reanalysis	Expert opinion cannot replace empirical validation
Validation planning	Convert uncertainty into prioritized studies	Evidence gaps, mechanism hypotheses, bias register and feasibility constraints	Selects studies that discriminate among causal explanations	Ranked validation plan	Prespecified endpoints, controls and decision criteria	Confirmatory bias or low-information experiments	Feed all validation results upstream	No advancement without meeting defined evidence needs
Feedback updating	Revise models and claims as evidence changes	New assays, datasets, replications, safety findings and clinical evidence	Updates causal structures, estimates and confidence	Versioned model and evidence history	Change logs, provenance and reproducibility checks	Silent model drift or selective updating	Mandatory re-evaluation after material evidence	Superseded claims must be withdrawn or revised
No clinical-claim boundary	Prevent framework output from being used as therapeutic authorization	Final evidence grade, uncertainty, safety status and clinical evidence	Calibrates permissible language	Explicit claim category and limitation statement	Independent claim-language review	Framework outputs may be marketed as efficacy or safety proof	Reassess only after suitable clinical validation	No diagnosis, prescribing or treatment recommendation follows from the framework alone

Figure 2 presents the proposed causal AI framework for natural product therapeutics, linking evidence

classification, causal question formulation, mechanism logic, intervention reasoning, translational confidence, validation planning, and feedback updating.

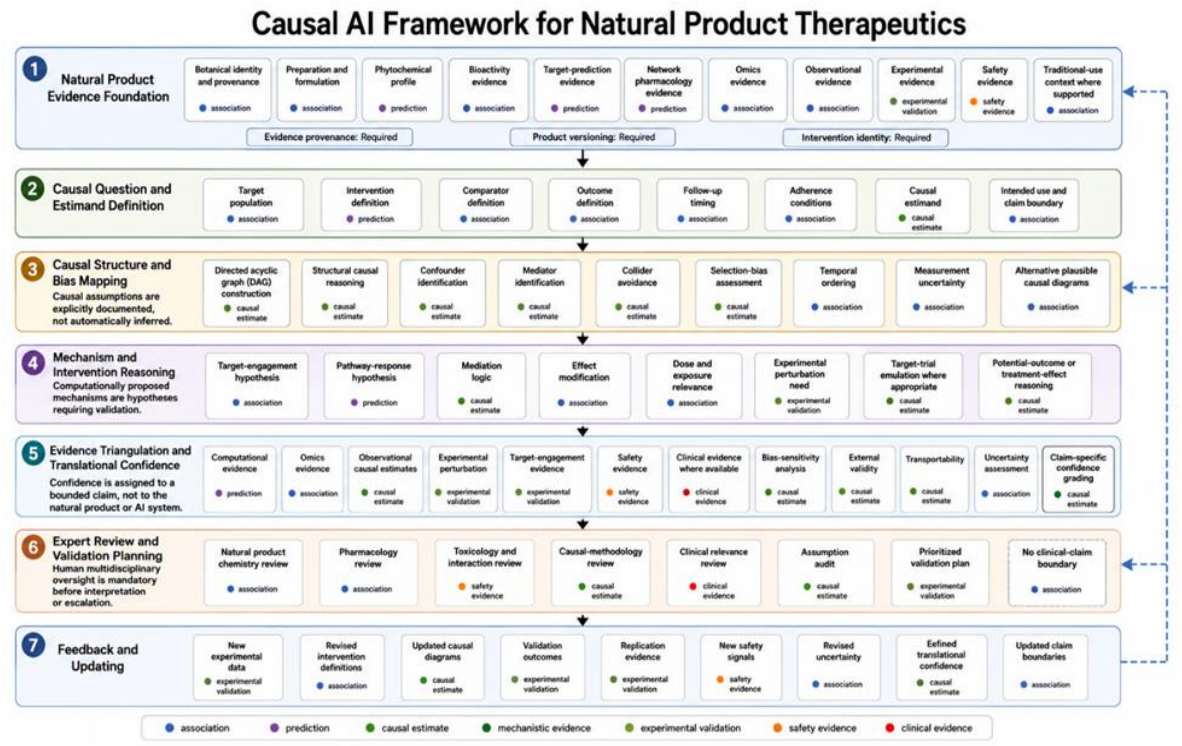

Figure 2. Causal AI Framework for Natural Product Therapeutics

Alt-text description

A conceptual causal AI framework showing natural product evidence flowing through causal question definition, DAG construction, bias mapping, mechanism assessment, intervention logic, evidence triangulation, confidence grading, safety review, validation planning, expert review, and feedback loops.

CONCLUSION

This article proposes a causal AI framework for natural product therapeutics that moves beyond associational discovery by integrating intervention definition, causal estimands, directed causal structures, confounding and bias assessment, mechanism hypotheses, mediation, effect modification, target-trial logic, experimental perturbation, evidence triangulation, safety evaluation, transportability, expert review, validation planning, and explicit claim boundaries. The framework treats docking, target prediction, network pharmacology, pathway enrichment, omics, observational signals, traditional use, and computational evidence as valuable but inferentially bounded inputs whose meaning depends on product identity, temporal ordering, exposure relevance, assumptions, and validation. Causal AI can strengthen evidence organization, counterfactual reasoning, assumption transparency, uncertainty assessment, and prioritization of informative experiments, but it cannot independently establish biological causation, therapeutic efficacy, clinical safety, or regulatory acceptability. Translational confidence should therefore be assigned to specific, clearly worded claims and revised as new experimental, observational, safety, or clinical evidence becomes available. By preserving distinctions among association, prediction, causal hypothesis, causal estimation, mechanistic confirmation, and clinical effectiveness, the proposed framework offers a disciplined foundation for more transparent and intervention-oriented natural product research.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

[1] Hernán MA, Wang W, Leaf DE. Target trial emulation: A framework for causal inference from

observational data. JAMA. 2022;328(24):2446-7. doi:10.1001/jama.2022.21383

[2] Atanasov AG, Zotchev SB, Dirsch VM, International Natural Product Sciences Taskforce, Supuran CT. Natural products in drug discovery: Advances and opportunities. Nat Rev Drug Discov. 2021;20(3):200-16. doi:10.1038/s41573-020-00114-z

[3] Shi J, Norgeot B. Learning causal effects from observational data in healthcare: A review and summary. Front Med (Lausanne). 2022;9:864882. doi:10.3389/fmed.2022.864882

[4] 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

[5] Schölkopf B, Locatello F, Bauer S, Ke NR, Kalchbrenner N, Goyal A, et al. Toward causal representation learning. Proc IEEE. 2021;109(5):612-34. doi:10.1109/JPROC.2021.3058954

[6] Sanchez P, Voisey JP, Xia T, Watson HI, O'Neil AQ, Tsaftaris SA. Causal machine learning for healthcare and precision medicine. R Soc Open Sci. 2022;9(8):220638. doi:10.1098/rsos.220638

[7] Mitra N, Roy J, Small D. The future of causal inference. Am J Epidemiol. 2022;191(10):1671-6. doi:10.1093/aje/kwac108

[8] Laubach ZM, Murray EJ, Hoke KL, Safran RJ, Perng W. A biologist's guide to model selection and causal inference. Proc Biol Sci. 2021;288(1943):20202815. doi:10.1098/rspb.2020.2815

[9] Pinzi L, Rastelli G. Molecular docking: Shifting paradigms in drug discovery. Int J Mol Sci. 2019;20(18):4331. doi:10.3390/ijms20184331

[10] Daina A, Michielin O, Zoete V. SwissTargetPrediction: Updated data and new features for efficient prediction of protein targets of small molecules. Nucleic Acids Res. 2019;47(W1). doi:10.1093/nar/gkz382

[11] Nogales C, Mamdouh ZM, List M, Kiel C, Casas AI, Schmidt HHHW. Network pharmacology: Curing causal mechanisms instead of treating symptoms. Trends Pharmacol Sci. 2022;43(2):136-50. doi:10.1016/j.tips.2021.11.004

[12] Reimand J, Isserlin R, Voisin V, Kucera M, Tannus-Lopes C, Rostamianfar A, et al. Pathway enrichment analysis and visualization of omics data using g, GSEA, Cytoscape and EnrichmentMap. Nat Protoc. 2019;14(2):482-517. doi:10.1038/s41596-018-0103-9

[13] Hasin Y, Seldin M, Lusis A. Multi-omics approaches to disease. Genome Biol. 2017;18(1):83. doi:10.1186/s13059-017-1215-1

[14] 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

[15] Tennant PWG, Murray EJ, Arnold KF, Berrie L, Fox MP, Gadd SC, et al. Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: Review and recommendations. Int J Epidemiol. 2021;50(2):620-32. doi:10.1093/ije/dyaa213

[16] Glymour C, Zhang K, Spirtes P. Review of causal discovery methods based on graphical models. Front Genet. 2019;10:524. doi:10.3389/fgene.2019.00524

[17] Rijnhart JJM, Lamp SJ, Valente MJ, MacKinnon DP, Twisk JWR, Heymans MW. Mediation analysis methods used in observational research: A scoping review and recommendations. BMC Med Res Methodol. 2021;21(1):226. doi:10.1186/s12874-021-01426-3

[18] Bours MJL. Tutorial: A nontechnical explanation of the counterfactual definition of effect modification and interaction. J Clin Epidemiol. 2021;134:113-24. doi:10.1016/j.jclinepi.2021.01.022

[19] Fu EL. Target trial emulation to improve causal inference from observational data: What, why, and how? J Am Soc Nephrol. 2023;34(8):1305-14. doi:10.1681/ASN.0000000000000152

[20] Li G, Peng X, Guo Y, Gong S, Cao S, Qiu F. Currently available strategies for target identification of bioactive natural products. Front Chem. 2021;9:761609. doi:10.3389/fchem.2021.761609

[21] VanderWeele TJ, Ding P. Sensitivity analysis in observational research: Introducing the E-value. Ann Intern Med. 2017;167(4):268-74. doi:10.7326/M16-2607

[22] Degtiar I, Rose S. A review of generalizability and transportability. Annu Rev Stat Appl. 2023;10:501-24. doi:10.1146/annurev-statistics-042522-103837

[23] Zhao Y, Yu Y, Wang H, Li Y, Deng Y, Jiang G, et al. Machine learning in causal inference: Application in pharmacovigilance. Drug Saf. 2022;45(5):459-76. doi:10.1007/s40264-022-01155-6

[24] Bykov K, Patorno E, D'Andrea E, He M, Lee H, Graff JS, et al. Prevalence of avoidable and bias-inflicting methodological pitfalls in real-world studies of medication safety and effectiveness. Clin Pharmacol Ther. 2022;111(1):209-17. doi:10.1002/cpt.2364

[25] Kawahara T, Shiba K, Tsuchiya A. Application of causal inference methods in the analysis of observational neurosurgical data: G-formula and marginal structural model. World Neurosurg. 2022;161:310-15. doi:10.1016/j.wneu.2021.09.141

[26] Hernán MA. Methods of public health research—strengthening causal inference from observational data. N Engl J Med. 2021;385(15):1345-8. doi:10.1056/NEJMp2113319

[27] Richens JG, Lee CM, Johri S. Improving the accuracy of medical diagnosis with causal machine learning. Nat Commun. 2020;11(1):3923. doi:10.1038/s41467-020-17419-7

[28] Rao S, Mamouei M, Salimi-Khorshidi G, Li Y, Ramakrishnan R, Hassaine A, et al. Targeted-BEHRT: Deep learning for observational causal inference on longitudinal electronic health records. IEEE Trans Neural Netw Learn Syst. 2024;35(4):5027-38. doi:10.1109/TNNLS.2022.3183864

[29] Dang LE, Gruber S, Lee H, Dahabreh IJ, Stuart EA, Williamson BD, et al. A causal roadmap for generating high-quality real-world evidence. J Clin Transl Sci. 2023;7(1). doi:10.1017/cts.2023.635

[30] Dahabreh IJ, Bibbins-Domingo K. Causal inference about the effects of interventions from observational studies in medical journals. JAMA. 2024;331(21):1845-53. doi:10.1001/jama.2024.7741

[31] Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, et al. Causal machine learning for predicting treatment outcomes. Nat Med. 2024;30(4):958-68. doi:10.1038/s41591-024-02902-1