Available online **www.ijpras.com**

International Journal of Pharmaceutical Research & Allied Sciences, 2025, 14(1):101-110
https://doi.org/10.51847/PZ7AUEMaSJ

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

From Molecular Correlation to Mechanistic Intervention: Reconstructing Causal Pharmacoinformatics for Drug Discovery, Safety Assessment, and Therapeutic Translation

Emma Wilson[1]*, Frederik De Jong[2], Lara Peters[2]

[1]Department of Causal Pharmacoinformatics and Mechanistic Inference, Faculty of Pharmacy, University of Edinburgh, Edinburgh, United Kingdom.
[2]Department of Drug Discovery and Therapeutic Translation, Faculty of Pharmaceutical Sciences, Utrecht University, Utrecht, Netherlands.Melbourne, Australia.

***Email:** emma.wilson@ed.ac.uk*

ABSTRACT

Pharmacoinformatics has become increasingly effective at identifying statistical relationships among molecular structure, biological measurements, and pharmaceutical outcomes, yet prediction alone does not establish what would happen if a compound, target, dose, pathway, or treatment policy were deliberately changed. This article develops an original, non-empirical causal-modelling framework for moving from molecular correlation toward mechanistic intervention in drug discovery, safety assessment, and therapeutic translation. The framework begins with explicit specification of variables, feasible interventions, comparators, outcomes, target systems, time horizons, and causal estimands. It then links causal-graph construction, molecular representation, biological context, identification analysis, counterfactual estimation, mechanistic interpretation, and claim-bounded validation. QSAR, network models, omics representations, and mechanistic models are positioned as complementary analytical modules rather than substitutes for causal identification. Perturbational experiments, external validation, temporal and chemical-space stress tests, negative controls, mediation analysis, and transport assessment are treated as distinct evidentiary requirements aligned with the intended claim. The proposed architecture also separates predictive uncertainty from uncertainty about causal structure, unmeasured confounding, intervention consistency, positivity, measurement error, and model transportability. Its central contribution is a gated causal-pharmacoinformatics logic in which no single predictive, explanatory, mechanistic, or validation component is sufficient by itself. The framework is intended to organize hypothesis formation, model design, evidence integration, and decision boundaries. It does not constitute an empirically validated model, clinical recommendation, regulatory standard, or deployment-ready system and requires domain-specific computational, experimental, translational, and implementation evaluation.

Key words: *Pharmacoinformatics, QSAR, Chemical space, External validation, Applicability domain, Calibration*

Received: 09 January 2025; **Revised:** 09 March 2025; **Accepted:** 11 March 2025

INTRODUCTION

Machine learning now supports molecular-property prediction, virtual screening, target prioritization, toxicity modelling, and other stages of drug discovery and development. These applications have expanded the capacity to rank candidates, detect complex associations, and integrate heterogeneous data, but their usual outputs remain predictions about observed labels or future observations under the data-generating conditions represented in training data [1]. Such outputs can be scientifically useful without answering the intervention question that

ultimately matters in pharmaceutical research: what outcome would follow if a molecular, biological, or treatment variable were deliberately changed?

The distinction begins at the level of molecular representation. Fingerprints, descriptors, graphs, strings, three-dimensional structures, and learned embeddings encode different aspects of chemical information and impose different invariances, resolutions, and inductive biases [2]. A model may therefore achieve strong prediction while relying on representation-specific regularities that are not stable under scaffold change, assay transfer, target perturbation, or altered biological context. Representation quality for prediction is consequently not equivalent to representation adequacy for causal reasoning.

Causal representation learning offers a more demanding objective: to organize learned variables around the data-generating mechanisms that remain meaningful under intervention and environmental change [3]. In pharmacoinformatics, this objective cannot be met by adding causal terminology to an otherwise associational pipeline. It requires an explicit intervention, a target system, a causal graph, identification assumptions, a counterfactual estimand, and validation evidence aligned with the claimed mechanism or pharmaceutical consequence.

This article presents an Original Causal-Modelling Framework Article rather than an empirical study or evidence review. It proposes a causal-pharmacoinformatics architecture that connects molecular interventions to biological mediators and pharmaceutical outcomes while separating established methodological concepts from newly proposed integration rules. The contribution is intended to clarify when QSAR, network, omics, interpretable machine-learning, and mechanistic models can support intervention-oriented reasoning, what additional evidence is required, and where claims must stop when causal assumptions or validation remain unresolved.

Why predictive association is insufficient for intervention

Predictive association and causal effect estimation address different questions. A predictive model estimates an outcome conditional on observed features, whereas an intervention-oriented analysis compares potential outcomes under specified alternatives. Counterfactual prediction therefore requires more than accurate forecasting: it requires a defined action, an admissible comparator, and assumptions linking observed data to unobserved outcomes under that action [4]. A molecular model that predicts toxicity for a candidate compound does not, by that fact alone, estimate the effect of changing a substructure, inhibiting a target, modifying exposure, or substituting a treatment.

The difference becomes consequential under dataset shift. Associational predictors may exploit pathways that are stable within a development dataset but unstable across laboratories, assays, chemical series, populations, or treatment policies. Causal structure can help distinguish features connected to the outcome-generating process from features whose predictive value depends on a particular environment, although causal formulation does not guarantee transportability [5]. For pharmaceutical modelling, external validity must therefore be linked to the intended intervention and deployment transition rather than inferred from performance on a single held-out sample.

Decision use intensifies this requirement. Causal machine learning can alter reasoning when several observed patterns are compatible with different underlying processes, because an intervention should act through the true data-generating structure rather than through a convenient correlate [6]. The methodological lesson is transferable to drug discovery: a biomarker associated with response may be a mediator, a proxy, a consequence, or a selection artifact. Acting on each role implies a different intervention and may produce a different outcome.

This article names the resulting mismatch the proposed intervention gap: the distance between what a predictive model estimates and what a pharmaceutical decision requires. The gap cannot be closed by discrimination, calibration, explanation, or mechanistic simulation alone. It narrows only when the intervention is explicit, causal roles are justified, the effect is identifiable, alternative explanations are addressed, and validation targets the claimed pathway. This interpretation is proposed as an organizing principle and does not imply that every pharmaceutical prediction should be reformulated as a causal effect.

Defining variables, interventions, outcomes, and causal estimands

A causal-pharmacoinformatics analysis must begin by stating the causal question without euphemism. Terms such as influence, driver, determinant, mechanism, or impact do not create causal meaning unless the corresponding intervention and contrast are defined [7]. Variables should therefore be classified by their role relative to the target question: exposure or treatment, outcome, confounder, mediator, effect modifier, collider, selection variable, proxy, instrument, or measurement. The same variable may occupy different roles under different estimands.

The target system must also be explicit. A compound effect may refer to a biochemical assay, a cell line, an organoid, an animal model, a patient subgroup, or a treatment policy, and validity in one system does not automatically establish validity in another. Target validity emphasizes that causal interpretation depends on whether the study design, population, setting, intervention, and outcome correspond to the intended decision [8]. In pharmaceutical modelling, this requires separating assay transport, chemical-space transport, biological transport, and clinical transport.

The proposed Pharmaceutical Causal Estimand Specification contains seven elements: target system, unit of intervention, feasible intervention, comparator, outcome, time horizon, and effect scale. Estimand specification removes ambiguity by stating exactly which treatment or intervention effect is sought and how deviations, competing events, or alternative intervention versions are handled [9]. For example, "the effect of target inhibition" is incomplete until the degree, timing, modality, biological context, comparator, and measured consequence are specified.

An estimand is a target of inquiry, not proof that an effect can be recovered. After specification, the analysis must examine consistency, exchangeability, positivity, temporal ordering, interference, measurement validity, and the plausibility of the causal graph. Molecular interventions require an additional feasibility condition: changes to descriptors or latent dimensions must correspond to chemically, biologically, or therapeutically realizable manipulations. The framework therefore rejects counterfactuals that are mathematically convenient but physically undefined.

Proposed causal pharmacoinformatics framework

Causal inference is increasingly relevant to drug discovery because target selection, compound optimization, safety assessment, and treatment translation are intervention problems rather than prediction problems alone [10]. The framework proposed here organizes these problems as a gated sequence: decision question, causal estimand, target-system boundary, causal graph, representation and measurement layer, identification analysis, counterfactual estimation, mechanistic triangulation, external transport assessment, and bounded claim formation. Failure at an early gate cannot be repaired solely by later predictive accuracy.

The causal graph is the framework's organizing object. It represents hypothesized relations among molecular structure, exposure, biological context, target engagement, pathway activity, cellular state, pharmacokinetics, pharmacodynamics, efficacy, and toxicity. Graphical causal analysis can make directional assumptions and candidate adjustment structures explicit before estimation [11]. However, the graph is not treated as a discovered universal mechanism; it is a revisable, target-specific model whose edges may be established, hypothesized, context-dependent, or unidentified.

Interventional data provide a stronger basis for orienting causal relations than observational co-variation alone. Large-scale perturbational studies demonstrate how interventions can constrain gene-network structure and expose discrepancies between observational and interventional dependencies [12]. Within the proposed framework, perturbations are used not merely to test predictive accuracy but to challenge graph edges, mediator assignments, effect directions, and transport assumptions. Their evidentiary value remains conditional on perturbation specificity, measurement quality, timing, and system relevance.

A causal pharmacoinformatics workflow therefore requires a pharmaceutical intervention question, an explicit graphical account of the data-generating structure, and revision against interventional evidence [10–12]. The integration of these elements into a non-substitutable architecture is the article's original contribution. The framework proposes six governing propositions: estimand primacy, prediction–identification non-substitutability, graph–representation compatibility, context-conditioned mechanism, validation triangulation, and claim-bounded translation.

Operationally, each study would produce a proposed Causal-Readiness Record documenting the estimand, graph assumptions, data source, representation choices, variable roles, identification strategy, overlap diagnostics, estimator, uncertainty analysis, perturbational evidence, external validation, and residual limitations. This record is not a score or certification. It is a structured disclosure mechanism intended to prevent predictive performance, interpretability, mechanistic plausibility, and decision readiness from being collapsed into a single claim.

Molecular structure, biological context, and treatment mechanism

Molecular representations are not neutral containers of causal information. Learned and engineered representations expose different structural regularities, and their usefulness depends on the property, endpoint, and chemical domain being modelled [13]. In the proposed framework, a representation is causally relevant only

when it preserves information needed to distinguish the intervention, contextual modifiers, mediators, and outcome-generating pathways specified by the estimand.

Interpretability does not resolve this requirement. Feature attribution, attention, saliency, and local explanations can reveal how a model generated an output, but they do not establish that the highlighted feature is a biological cause or therapeutically actionable mechanism [14]. Mechanistic interpretation therefore requires a causal role, temporal plausibility, intervention consistency, and evidence that the attributed feature participates in the proposed pathway.

Prospective prediction can nevertheless support causal investigation when it is followed by appropriate experimental testing. Computational antibiotic discovery has shown how molecular prediction can prioritize compounds for biological evaluation, but the resulting therapeutic and mechanistic claims depend on subsequent assays and validation rather than the model output alone [15]. The same principle applies to efficacy, toxicity, target engagement, and resistance predictions.

The proposed graph–representation compatibility criterion asks whether the chosen molecular and biological variables are sufficient for the target intervention question. Structure, exposure, cellular state, target abundance, pathway activity, and treatment context should be represented at resolutions compatible with the claimed mechanism. A failure at this stage restricts the permissible claim to association or prediction, even when the model remains practically useful.

Confounding, selection bias, mediation, and collider structures

Confounding arises when a common cause influences both the intervention or exposure and the outcome. Variable adjustment should therefore follow a justified causal graph rather than automated selection based only on predictive importance [16]. In pharmacoinformatics, potential confounders include assay protocol, compound series, target class, disease severity, concomitant therapy, laboratory practice, and exposure conditions.

Selection can create additional bias when inclusion in an assay, screen, cohort, or analysis depends on causes of both the model input and outcome. Conditioning on a collider may induce an association that is absent in the target population or reverse the apparent direction of an existing relationship [17]. Candidate filtering and complete-case analysis must consequently be treated as causal design choices.

Mediation analysis asks whether an intervention acts through an intermediate process, such as target engagement, pathway modulation, exposure, or cellular-state change. Direct and indirect effects are distinct estimands with different assumptions, and neither can be inferred merely by adding the suspected mediator to a predictive model [18]. Temporal ordering, mediator–outcome confounding, and intervention-induced confounding must be considered explicitly.

The framework therefore proposes a causal-role audit for every variable entering model development or interpretation. Each variable should be labelled relative to the estimand, and the consequences of conditioning, excluding, proxying, or intervening on it should be stated. These labels remain hypotheses about the target system and require revision when perturbational or external evidence contradicts the assumed graph.

Counterfactual prediction for compound and target interventions

Counterfactual prediction estimates differences between outcomes that would occur under alternative interventions. It therefore requires identification assumptions and effect-estimation methods that differ from ordinary outcome prediction [19]. In pharmaceutical settings, relevant interventions may include compound administration, structural modification, dose adjustment, target inhibition, pathway perturbation, formulation change, or selection of a treatment policy.

Individualized or context-specific effects create further challenges. Real-world effect estimation is sensitive to confounding, treatment-policy bias, informative missingness, limited overlap, and differences between the data source and target population [20]. Comparable problems arise when molecular series, assay conditions, or biological systems determine which interventions are observed and where reliable comparisons are available.

Meta-learners can combine flexible outcome models to estimate heterogeneous treatment effects across covariates [21]. Their value in pharmacoinformatics lies in separating nuisance prediction from effect estimation and allowing context-dependent responses. They do not, however, solve intervention ambiguity, unmeasured confounding, positivity failure, or incorrect causal structure.

The framework proposes an intervention hierarchy comprising structural, exposure, target, pathway, and treatment-policy interventions. Each level changes a different part of the causal system and may require a different comparator, representation, validation design, and decision boundary. Estimated individual effects should be

interpreted as model-dependent counterfactual quantities, not directly observed personal outcomes or guaranteed treatment responses.

Integration with QSAR, networks, omics, and mechanistic models

QSAR and drug-target prediction models can provide high-value modules for estimating molecular properties, activities, and interaction probabilities. Large-scale comparisons show their capacity to learn useful predictive structure across pharmacological tasks [22]. Within the proposed architecture, these outputs become inputs to causal reasoning only after their intervention meaning, applicability domain, and measurement role are specified. Network models add relational information about targets, pathways, drugs, diseases, and safety outcomes. Graph-based modelling of polypharmacy can represent complex patterns of adverse-effect risk across drug combinations [23]. Such edges remain predictive unless a causal direction, intervention, and identification strategy are justified, and network connectivity should not be interpreted automatically as mechanistic transmission.

Omics and foundation representations can encode regulatory state and transfer information across biological contexts. Pretrained network models have demonstrated predictive transfer to gene-regulatory and disease-relevant tasks [24]. For causal use, these representations must be connected to explicit perturbations, effect modifiers, temporal states, and outcomes, with validation in the target biological system.

QSAR, safety-network models, and transferable biological representations can serve as complementary modules within the causal architecture, although none independently identifies an intervention effect [22–24]. Mechanistic simulators may further test whether proposed pathways are dynamically coherent, but simulation fidelity does not replace identification from data. Integration is therefore modular and evidentiary rather than a merger into one undifferentiated model.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for integrated construct, components, evidence requirements, failure modes, and decision boundaries for from molecular correlation to mechanistic intervention..

Table 1. Integrated construct, components, evidence requirements, failure modes, and decision boundaries for From Molecular Correlation to Mechanistic Intervention.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Estimand specification	Ambiguous causal question	Target system, intervention, comparator, outcome, horizon	Proposed Pharmaceutical Causal Estimand Specification, informed by explicit causal questions [7] and estimand definition [9]	Aligns modelling with a decision-relevant effect	Feasible intervention and justified target contrast	Undefined or inconsistent intervention	Specification does not establish identifiability
Molecular representation	Loss or distortion of intervention-relevant information	Structure, descriptors, graphs, embeddings, chemical domain	Proposed graph–representation compatibility criterion; representation dependence is established [13]	Clarifies what molecular information supports the graph	Sensitivity analyses and chemical-domain evaluation	Shortcut learning, activity cliffs, representation loss	Predictive usefulness does not prove causal adequacy

Biological-context graph	Correlation mistaken for mechanism	Cell state, target, pathway, exposure, time, system context	Proposed target-specific causal graph informed by biological causal discovery [11]	Makes directional assumptions inspectable	Perturbation, temporal evidence, domain knowledge	Incorrect edges or omitted context	The graph is revisable, not universal
Bias and identification	Distorted intervention–outcome relation	Confounders, mediators, colliders, selection variables	Proposed causal-role audit informed by DAG guidance [16]	Determines defensible adjustment and identification	Graph justification, overlap, negative controls	Residual confounding, collider bias, positivity failure	Adjustment cannot compensate for an unidentified effect
Counterfactual estimation	Prediction substituted for effect estimation	Intervention, comparator, covariates, potential outcomes	Proposed intervention hierarchy informed by causal effect modelling [19]	Estimates average or context-specific contrasts	Identification assumptions and estimator diagnostics	Extrapolation and unstable heterogeneous effects	Estimated effects remain assumption-dependent
Integrated analytical modules	Fragmented QSAR, network, and omics evidence	Molecular, relational, omic, and mechanistic models	Proposed modular integration; predictive capabilities are established [22]	Combines complementary evidence without causal relabelling	Cross-module consistency and target-system validation	Error propagation and incompatible scales	Integration does not create causality
Validation architecture	Internal accuracy mistaken for external validity	Benchmarks, perturbations, external systems, temporal splits	Proposed validation-triangulation gate informed by molecular benchmarking [25]	Aligns tests with the claimed relation	Leakage-resistant splits, perturbation, replication, transport analysis	Memorization and context mismatch	No single test validates the complete framework
Translation boundary	Model output overstated as pharmaceutical consequence	Evidence strength, uncertainty, intended decision	Proposed Causal Translation Boundary; uncertainty must be evaluated empirically [26]	Restricts claims to their weakest unresolved gate	Calibrated prediction and explicit causal limitations	Unmeasured structure and miscalibration	No clinical, regulatory, or deployment claim follows automatically

Validation using perturbational and external evidence

Validation must be matched to the intended claim. Molecular benchmarks support reproducible comparisons only when tasks, endpoints, data splits, and evaluation procedures reflect the problem being tested [25]. Random hold-outs may assess interpolation, whereas scaffold, temporal, laboratory, assay, or population splits address different forms of external generalization. None alone establishes an intervention effect.

Benchmark design can obscure failure when structurally related observations cross the training–test boundary. Ligand-based evaluations may reward memorization rather than prospective generalization [27]. Causal-pharmacoinformatics validation should therefore test leakage, chemical-space coverage, intervention overlap, representation sensitivity, calibration, and performance across the transition corresponding to intended use.

Perturbational validation asks whether manipulating the proposed cause changes the predicted mediator or outcome in the expected direction. It may involve compound, dose, target, pathway, or genetic perturbations, with temporal and exposure controls. Interventional network evidence can challenge assumed edge direction [12]. Negative controls and alternative mechanisms should be included where feasible.

The proposed validation ladder distinguishes internal predictive assessment, chemical-space external validation, temporal or laboratory transport, perturbational validation, mechanistic triangulation, and decision-context evaluation. Benchmark standards inform the predictive layer [25], but later layers require different evidence. A failed higher-level test restricts the claim without necessarily invalidating every lower-level predictive use.

Limitations and assumptions

Molecular response surfaces are not uniformly smooth. Activity cliffs show that structurally similar compounds can exhibit sharply different activities, exposing local failures that aggregate metrics may conceal [28]. This limits counterfactual reasoning based on small structural changes, especially when the intervention enters an underrepresented region of chemical or conformational space.

Predictive uncertainty must also be estimated and evaluated rather than inferred from model complexity or ensemble construction. Molecular uncertainty methods differ in calibration and behaviour across tasks [26]. Even well-calibrated predictive uncertainty does not quantify uncertainty from an incorrect causal graph, unmeasured confounding, ambiguous intervention versions, or unsupported transport assumptions.

The framework additionally assumes that interventions can be defined consistently, relevant alternatives have sufficient overlap, outcomes are measured comparably, and interference is either absent or modelled. Feedback, adaptive treatment, metabolism, polypharmacy, compensatory pathways, and evolving disease states can violate simple acyclic or static representations. Measurement error may also change apparent variable roles.

These limitations prevent universal causal-readiness classification. The framework is a reasoning architecture rather than a validated scoring instrument, and its components may need modification across modalities and development stages. Failure to identify an effect should be reported as a boundary, not converted into a causal conclusion through interpretability, simulation, or predictive performance.

Research agenda for causal pharmaceutical modelling

Molecular foundation models create opportunities to learn transferable representations, but scale and pretraining do not themselves establish causal variables or mechanisms. Future work should connect representation learning with interventions, invariance, causal abstraction, and experimentally testable molecular relations [29]. The resulting models must distinguish reconstruction of observed regularities from prediction under manipulated chemical or biological states.

Distribution shift should be specified according to the actual transition of interest. Real-world molecular out-of-distribution settings differ by scaffold, property, assay, time, laboratory, and discovery programme [30]. Research should determine which shifts are addressable through representation, reweighting, causal transport, new experimentation, or explicit refusal to extrapolate.

The research agenda must jointly stress-test activity-cliff sensitivity, uncertainty calibration, and the causal adequacy of molecular foundation representations [26, 28, 29]. Priority studies should compare predictive and causal objectives under controlled perturbations, evaluate graph revision after contradictory evidence, and report how uncertainty changes when causal assumptions rather than model parameters are varied.

Figure 1 presents the original conceptual synthesis for causal graph connecting molecular interventions to pharmaceutical outcomes, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

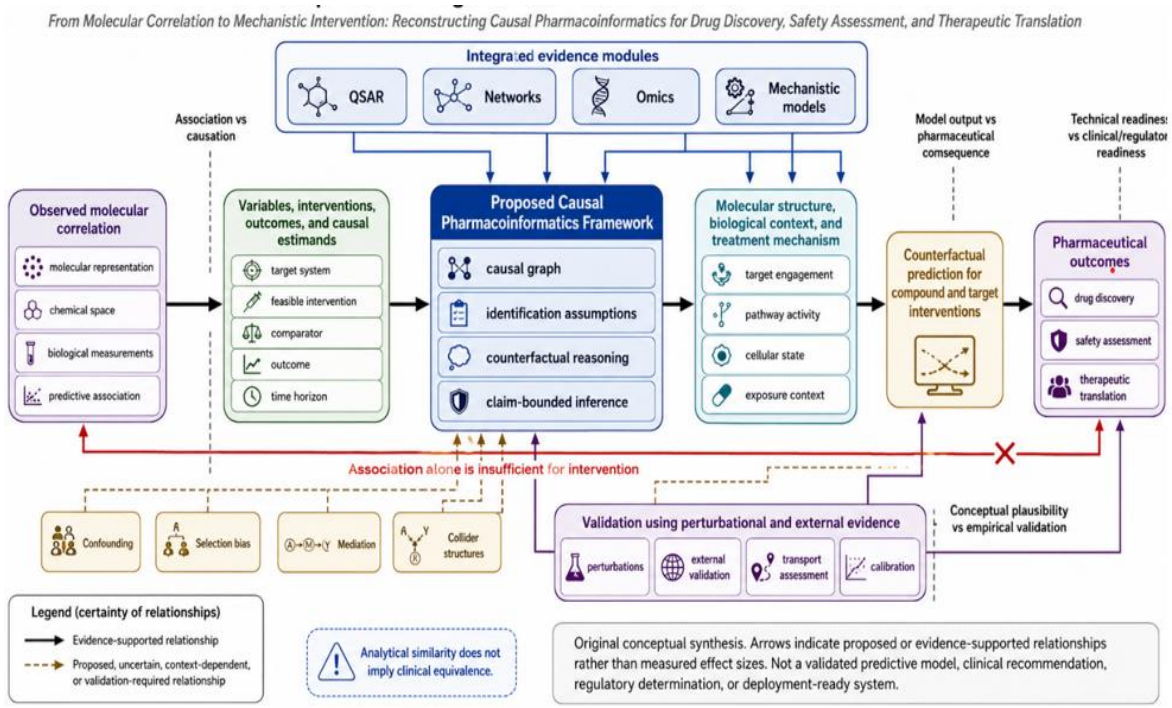

Figure 1. Causal Graph Connecting Molecular Interventions to Pharmaceutical Outcomes

The figure is an original conceptual synthesis developed for "From Molecular Correlation to Mechanistic Intervention: Reconstructing Causal Pharmacoinformatics for Drug Discovery, Safety Assessment, and Therapeutic Translation." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

The framework should next be tested through preregistered causal questions, perturbational datasets, cross-laboratory replication, prospective chemical-series transitions, and comparisons between predictive and intervention-oriented model selection. Reporting standards should require estimands, graph assumptions, overlap diagnostics, negative controls, calibration, transport targets, and residual uncertainty. Governance research should examine how such records support review without turning the framework into a certification system.

CONCLUSION

This article proposes a causal-pharmacoinformatics architecture for connecting molecular interventions to pharmaceutical outcomes through explicit estimands, target-system boundaries, causal graphs, representation choices, identification analysis, counterfactual estimation, mechanistic triangulation, and claim-aligned validation. Its central principle is non-substitutability: prediction, explanation, simulation, perturbation, and external validation answer different questions and cannot independently establish intervention validity. The framework may organize more transparent discovery, safety, and translation research, but its graphs, criteria, and decision boundaries remain proposed constructs requiring domain-specific empirical evaluation. It does not provide a clinical recommendation, regulatory determination, universal causal model, or deployment-ready system.

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. David L, Thakkar A, Mercado R, Engkvist O. Molecular representations in AI-driven drug discovery: A review and practical guide. J Cheminform. 2020;12(1):56. doi:10.1186/s13321-020-00460-5
3. Schölkopf B, Locatello F, Bauer S, Ke N, Kalchbrenner N, Goyal A, et al. Toward causal representation learning. Proc IEEE. 2021;109(5):612-34. doi:10.1109/JPROC.2021.3058954
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. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
6. 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
7. Hernán MA. The C-word: Scientific euphemisms do not improve causal inference from observational data. Am J Public Health. 2018;108(5):616-19. doi:10.2105/AJPH.2018.304337
8. Westreich D, Edwards JK, Lesko CR, Cole SR, Stuart EA. Target validity and the hierarchy of study designs. Am J Epidemiol. 2019;188(2):438-43. doi:10.1093/aje/kwy228
9. Kahan BC, Cro S, Li F, Harhay MO. Eliminating ambiguous treatment effects using estimands. Am J Epidemiol. 2023;192(6):987-94. doi:10.1093/aje/kwad036
10. Michoel T, Zhang JD. Causal inference in drug discovery and development. Drug Discov Today. 2023;28(10):103737. doi:10.1016/j.drudis.2023.103737
11. Amar D, Sinnott-Armstrong N, Ashley EA, Rivas MA. Graphical analysis for phenome-wide causal discovery in genotyped population-scale biobanks. Nat Commun. 2021;12(1):350. doi:10.1038/s41467-020-20516-2
12. Brown BC, Tokolyi A, Morris JA, Lappalainen T, Knowles DA. Large-scale causal discovery using interventional data sheds light on gene network structure in K562 cells. Nat Commun. 2025;16(1):9628. doi:10.1038/s41467-025-64353-7
13. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370-88. doi:10.1021/acs.jcim.9b00237
14. 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
15. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, et al. A deep learning approach to antibiotic discovery. Cell. 2020;180(4):688-702.e13. doi:10.1016/j.cell.2020.01.021
16. Tennant PWG, Murray EJ, Arnold KF, Berrie L, Fox MP, Gadd SC, et al. Use of directed acyclic graphs to identify confounders in applied health research: Review and recommendations. Int J Epidemiol. 2021;50(2):620-32. doi:10.1093/ije/dyaa213
17. Munafò MR, Tilling K, Taylor AE, Evans DM, Davey Smith G. Collider scope: When selection bias can substantially influence observed associations. Int J Epidemiol. 2018;47(1):226-35. doi:10.1093/ije/dyx206
18. Nguyen TQ, Schmid I, Stuart EA. Clarifying causal mediation analysis for the applied researcher: Defining effects based on what we want to learn. Psychol Methods. 2021;26(2):255-71. doi:10.1037/met0000299
19. 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
20. Bica I, Alaa AM, Lambert C, van der Schaar M. From real-world patient data to individualized treatment effects using machine learning: Current and future methods to address underlying challenges. Clin Pharmacol Ther. 2021;109(1):87-100. doi:10.1002/cpt.1907
21. Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. Proc Natl Acad Sci U S A. 2019;116(10):4156-65. doi:10.1073/pnas.1804597116
22. Mayr A, Klambauer G, Unterthiner T, Steijaert M, Wegner JK, Ceulemans H, et al. Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. Chem Sci. 2018;9(24):5441-51. doi:10.1039/C8SC00148K

23. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics. 2018;34(13). doi:10.1093/bioinformatics/bty294
24. Theodoris CV, Xiao L, Chopra A, Chaffin MD, Al Sayed ZR, Hill MC, et al. Transfer learning enables predictions in network biology. Nature. 2023;618(7965):616-24. doi:10.1038/s41586-023-06139-9
25. 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
26. Scalia G, Grambow CA, Pernici B, Li YP, Green WH. Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction. J Chem Inf Model. 2020;60(6):2697-717. doi:10.1021/acs.jcim.9b00975
27. 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
28. van Tilborg D, Alenicheva A, Grisoni F. Exposing the limitations of molecular machine learning with activity cliffs. J Chem Inf Model. 2022;62(23):5938-51. doi:10.1021/acs.jcim.2c01073
29. Lobentanzer S, Rodriguez-Mier P, Bauer S, Saez-Rodriguez J. Molecular causality in the advent of foundation models. Mol Syst Biol. 2024;20(8):848-58. doi:10.1038/s44320-024-00041-w
30. Tossou P, Wognum C, Craig M, Mary H, Noutahi E. Real-world molecular out-of-distribution: Specification and investigation. J Chem Inf Model. 2024;64(3):697-711. doi:10.1021/acs.jcim.3c01774