Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

A CAUSAL GRAMMAR FOR DRUG–DRUG INTERACTIONS FROM MOLECULAR INITIATION TO EXPOSURE CHANGE AND CLINICAL CONSEQUENCE

Meenal Joshi[1]*, Rohan Patil[2], Arjun Nair[3], Pooja Reddy[4]

1. ***Department of Causal Drug–Drug Interaction Modeling, Faculty of Pharmacy, Indian Institute of Science, Bengaluru, India.***
2. ***Department of Molecular Initiation and Exposure Change, Faculty of Pharmacy, IIT Bombay, Mumbai, India.***
3. ***Department of Clinical Consequence and Interaction Grammar, Faculty of Pharmaceutical Sciences, University of Delhi, Delhi, India.***
4. ***Department of Pharmacological Causal Reasoning, Faculty of Pharmacy, University of Agricultural Sciences Dharwad, Dharwad, India.***

ARTICLE INFO

Received:
05 February 2025
Received in revised form:
14 April 2025
Accepted:
16 April 2025
Available online:
28 April 2025

Keywords: Drug–drug interaction, Causal ontology, Pharmacokinetics, Drug metabolism, Membrane transporters, Knowledge graph

ABSTRACT

Drug–drug interaction knowledge is distributed across molecular assays, enzyme and transporter studies, pharmacokinetic models, clinical investigations, computational predictions, knowledge bases, and prescribing alerts. These resources frequently describe different portions of an interaction while using incompatible entities, relation types, temporal assumptions, evidence labels, and levels of causal interpretation. Consequently, a molecular perturbation may be represented as equivalent to an exposure change, a predicted association may appear indistinguishable from an observed interaction, and a pairwise alert may omit the dose, timing, patient state, or physiological process that determines its relevance. This article proposes a causal grammar for representing drug–drug interactions as evidence-bearing chains extending from a context-specific molecular initiation event through enzymes, transporters, disposition processes, systemic or tissue exposure changes, and bounded clinical consequences. The grammar distinguishes interacting-drug roles, molecular mechanisms, biological mediators, pharmacokinetic transitions, contextual qualifiers, evidence assertions, contradictions, uncertainty, and applicability boundaries. It is designed to preserve direction, magnitude, temporal sequence, dose and regimen conditions, anatomical location, genotype, current metabolic phenotype, organ function, inflammation, and other patient-specific modifiers without reducing them to a single interaction score. The proposed construct also separates molecular plausibility, computational prediction, mechanistic explanation, quantitative simulation, clinical observation, and decision relevance as different evidence states requiring different forms of validation. Its purpose is not to generate dosing recommendations or establish clinical causality autonomously, but to provide a coherent semantic architecture for evidence integration, mechanistic reasoning, model documentation, and context-aware decision support. The grammar remains conceptual and requires formal ontology testing, pharmacological adjudication, quantitative evaluation, interoperability assessment, and prospective human-centered validation before operational use.

*This is an **open-access** article distributed under the terms of the Creative Commons Attribution-Non Commercial-Share Alike 4.0 License, which allows others to remix, and build upon the work non commercially.*

To Cite This Article: Joshi M, Patil R, Nair A, Reddy P. A Causal Grammar for Drug–Drug Interactions from Molecular Initiation to Exposure Change and Clinical Consequence. Pharmacophore. 2025;16(2):86-97. https://doi.org/10.51847/XhzZht5y4r

Introduction

Drug–drug interactions arise through diverse processes that include reversible or time-dependent enzyme inhibition, enzyme induction, substrate competition, modulation of uptake or efflux transporters, and combinations of these mechanisms. Their consequences may occur during absorption, intestinal extraction, hepatic uptake, metabolism, biliary excretion, renal secretion, or tissue distribution. Computational pharmaceutical science has added further representations based on molecular features, literature extraction, graph structures, and predictive models. These developments have expanded the available evidence but

Corresponding Author: Meenal Joshi; Department of Causal Drug–Drug Interaction Modeling, Faculty of Pharmacy, Indian Institute of Science, Bengaluru, India. E-mail: meenal.joshi@iisc.ac.in.

have not produced a shared semantic account of what an interaction statement means, which causal level it describes, or how it should be connected to subsequent exposure and clinical consequences [1–3].

This unresolved problem is especially visible in computational drug–drug interaction prediction. Models may learn from molecular similarity, reported adverse effects, known pairwise relations, biomedical knowledge graphs, or combinations of these sources. Performance on a held-out benchmark can indicate successful pattern recovery under the benchmark's assumptions, but it does not establish whether the model has identified the responsible enzyme or transporter, whether the inferred relation is directional, whether it applies to an unseen compound, or whether it remains valid under a different dose, regimen, population, or disease state. Cold-start behavior and uneven representation of drugs and relation types further limit the interpretation of aggregate predictive performance [4]. Thus, the central difficulty cannot be resolved by optimizing one classification score, link-prediction measure, confidence value, or alert-severity category.

The deeper problem is an absence of causal granularity. A statement that "drug A interacts with drug B" can refer to a molecular binding event, inhibition of a metabolic enzyme, altered transporter activity, a change in clearance, a difference in systemic exposure, a tissue-specific concentration change, an adverse-effect association, or an electronic prescribing warning. These statements are not semantically interchangeable. Enzyme-mediated interaction evidence, for example, may progress from an in vitro inhibition observation to a mechanistic interpretation and then to clinical pharmacokinetic evaluation, with uncertainty introduced at each transition [1]. When these levels are collapsed, a plausible initiation mechanism may be overstated as an established clinical consequence, whereas clinically observed exposure changes may be disconnected from the biological processes needed to explain or generalize them.

This article therefore proposes a causal grammar for drug–drug interactions. The term *grammar* denotes a controlled set of entities, typed relations, qualifiers, evidence states, and composition rules through which an interaction can be expressed as a conditional causal chain. The proposed grammar assigns context-specific roles to interacting drugs; distinguishes molecular initiation events from enzyme, transporter, and disposition processes; represents systemic and tissue exposure changes explicitly; and permits clinical consequences only when supported by an appropriate evidence path. Direction, magnitude, timing, dose, administration sequence, patient context, provenance, contradiction, uncertainty, and applicability are represented as integral components rather than optional narrative annotations. The contribution is conceptual and ontological: it is intended to organize evidence and constrain reasoning, not to claim empirical validation, universal coverage, autonomous prediction, regulatory acceptance, or readiness for clinical deployment.

Why Drug–Drug Interaction Knowledge Remains Semantically Fragmented

Drug–drug interaction evidence is first fragmented because it is generated for different scientific purposes. Molecular assays characterize binding, inhibition, induction, or transport under experimentally defined conditions. Pharmacokinetic investigations measure concentration–time changes. Modeling studies connect mechanistic parameters to simulated exposure. Clinical reports describe treatment responses, toxicities, or monitoring decisions. Literature-mining systems extract relations from text, while interaction databases and prescribing systems compress evidence into pairwise records or alert categories. An ontology-grounded knowledge graph can connect some of these sources and preserve supporting or conflicting assertions, but its construction also reveals that extracted relations may be incomplete, contradictory, or dependent on how entities and evidence statements were normalized [5]. Fragmentation therefore concerns not only where information is stored but also what type of assertion each source contains.

A second source of fragmentation is the inconsistent definition of entities and relations. Knowledge graphs can integrate drugs, proteins, pathways, diseases, phenotypes, and studies, yet integration is meaningful only when identifiers and predicates retain their scientific interpretation [6]. One resource may use *inhibits* to denote a biochemical effect on an enzyme, another may use it for a reduction in drug clearance, and another may apply the same label to an observed decrease in treatment effect. Similarly, *interacts with* may be symmetric even when the pharmacological relationship is directional. Recent analyses of drug-discovery knowledge graphs continue to identify inconsistent source modeling, limited standardization, incomplete provenance, and variable interpretability as barriers to reliable reuse [7]. A causal grammar must therefore define not only the participating objects but also the level at which each relation operates.

A third problem is that biomedical datasets differ in scope, curation standards, identifiers, evidence density, and relation semantics. Some resources emphasize molecular mechanisms, whereas others prioritize clinical associations or adverse events. Some preserve negative or uncertain findings, while others contain only reported positive relations. These differences create uneven graph connectivity and can make frequently studied drugs appear more causally informative than poorly represented drugs [8]. When datasets are merged without an explicit evidence model, repeated assertions may be mistaken for independent confirmation, missing relations may be interpreted as negative evidence, and provenance may be lost. The proposed grammar addresses this problem by treating every evidence statement as a versioned assertion linked to its source, method, context, and interpretive status rather than as an unqualified edge.

A fourth source of fragmentation arises during the movement from language to structured knowledge. Heterogeneous knowledge graphs can improve extraction of drug–drug interaction relations from biomedical literature by supplying additional entity and relational context [9]. Nevertheless, an extracted sentence-level relation does not automatically establish the direction of the pharmacokinetic effect, the relevant enzyme or transporter, the affected exposure variable, or the clinical importance of the interaction. Extraction, normalization, causal interpretation, and clinical adjudication are separate operations. The causal grammar therefore requires explicit separation between a reported assertion, a computationally inferred relation, a

mechanistically supported transition, and a clinically supported consequence. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop why drug-drug interaction knowledge remains semantically fragmented within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for why drug-drug interaction knowledge remains semantically fragmented

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Molecular initiation evidence	What physical or biochemical event begins the proposed interaction?	Binding, inhibition, induction, competition, or transporter-modulation evidence generated under specified experimental conditions	Defines the earliest mechanistic element of the causal chain	Treating an in vitro molecular event as an established in vivo interaction	Molecular initiation indicates mechanistic possibility but does not alone establish exposure change
Enzyme and transporter evidence	Which biological mediator is affected, in which organ or cellular location?	Evidence concerning metabolic enzymes, uptake transporters, efflux transporters, regulatory receptors, or coupled mediator systems	Connects molecular initiation to a drug-disposition process	Assigning one mediator when several processes jointly determine disposition	Mediator attribution is conditional on the experimental and physiological context
Pharmacokinetic evidence	Which disposition variable or concentration–time profile changes?	Clinical pharmacokinetic observations or appropriately qualified mechanistic models	Defines the transition from biological mechanism to systemic or tissue exposure	Reducing exposure change to a binary interaction label	Exposure must specify the affected drug, compartment, direction, timing, and evidentiary basis
Clinical-consequence evidence	Is the exposure change associated with altered efficacy, toxicity, treatment failure, or monitoring need?	Pharmacodynamic, clinical-outcome, or adjudicated decision evidence	Establishes the possible downstream relevance of the interaction	Inferring clinical harm or benefit solely from molecular or exposure evidence	A pharmacokinetic interaction does not automatically establish a clinically important consequence
Computational prediction	Has a model inferred a relation from molecular, textual, or graph-derived information?	Statistical, machine-learning, or link-prediction output evaluated under a defined benchmark	Supports hypothesis generation and identification of missing relations	Treating predictive accuracy as mechanistic confirmation or clinical usefulness	Prediction is a distinct evidence state and requires independent mechanistic and clinical evaluation
Literature-derived assertion	What interaction claim is explicitly reported in a publication?	Text extraction, manual curation, or structured evidence abstraction	Preserves the original reported statement and its source	Conflating extraction confidence with scientific truth	A reported assertion must remain linked to its wording, design, context, and limitations
Identity and terminology mapping	Do different resources refer to the same drug, metabolite, enzyme, transporter, process, or outcome?	Canonical identifiers, terminology mappings, synonym control, and versioning	Enables integration without losing entity identity	Identifier collision, false equivalence, or silent merging of distinct concepts	Technical mapping must not erase scientifically important distinctions
Directionality	Which entity acts on which mediator or drug, and what variable increases or decreases?	Subject–predicate–object structure with an explicit affected property	Prevents symmetric pairwise labels from replacing causal relations	Reversal of perpetrator and object roles or unspecified effect polarity	Direction must be assertion-specific and tied to a named variable
Temporal and regimen context	Under what dose, route, sequence, duration, and time state does the relation hold?	Study protocol, dosing history, onset, persistence, reversibility, and washout information	Defines the conditions under which the assertion is interpretable	Treating a time-dependent effect as immediate, constant, or permanent	A timeless interaction statement cannot support time-specific inference
Patient-context evidence	Which physiological, genetic, disease, or organ-function state modifies the interaction?	Genotype, phenotype, age, organ function, inflammation, comorbidity, and population evidence	Restricts the applicability of mechanistic and clinical assertions	Generalizing from one population or phenotype to all patients	Context variables are mandatory when they materially change the interaction
Provenance and evidence independence	Where did an assertion originate, and are apparently repeated findings independent?	Source, version, extraction method, study design, and transformation history	Enables audit, updating, and weighting of evidence	Counting duplicated database records as independent confirmation	Provenance completeness supports traceability but does not guarantee scientific validity
Contradiction and uncertainty	Do sources disagree in mechanism, direction, magnitude, timing, or clinical interpretation?	Explicit comparison of evidence assertions and their contexts	Preserves unresolved disagreement instead of forcing premature consensus	Silent overwriting, averaging, or removal of conflicting evidence	Contradictory evidence should remain visible until contextual or empirical resolution is justified

Decision-support representation	What bounded action or interpretation may the evidence support?	Context-aware rules, clinical workflow evidence, and human adjudication	Connects the ontology to possible downstream use	Converting an evidence representation into an autonomous recommendation	The causal grammar may organize decision evidence but does not itself prescribe treatment

From Molecular Initiation to Enzymes, Transporters, Exposure, and Outcomes

A causal representation of drug–drug interactions must begin with a molecular initiation event but cannot end there. A perpetrator drug may inhibit or induce an enzyme, compete for a transporter, modify transporter expression, or participate in a coupled process affecting uptake, metabolism, and efflux. Physiologically based pharmacokinetic models provide one means of connecting transporter perturbations to organ-level disposition and concentration–time changes, while clinical pharmacology establishes that transporter-mediated interactions may influence systemic exposure, tissue exposure, or both. Transporter and enzyme activities can also be interdependent, particularly when hepatic uptake controls intracellular substrate availability for metabolism or efflux [10–12]. The proposed grammar therefore represents molecular initiation, biological mediation, disposition-process change, and exposure alteration as distinct but composable causal layers.

The first layer specifies the context-dependent roles of the interacting compounds and the type of initiation event. A drug may act as a perpetrator with respect to one pathway while simultaneously functioning as an object of another mechanism. The predicate must therefore identify the affected enzyme, transporter, or regulatory process and distinguish reversible inhibition, time-dependent inhibition, induction, competition, or another supported mechanism. Quantitative prediction of CYP3A4 inhibition illustrates why this distinction matters: estimates depend on the selected method, input parameters, experimental data, and evaluation assumptions [13]. The grammar consequently attaches method provenance, evidence status, uncertainty, and applicability conditions to the initiation relation rather than encoding inhibition as an unconditional drug property.

The second and third layers represent the mediator and the affected disposition process. An enzyme-centered pathway may alter intestinal or hepatic metabolism, whereas transporter modulation may affect absorption, hepatic uptake, biliary excretion, renal secretion, or tissue distribution. Disease state can modify these pathways before any perpetrator drug is introduced. Inflammation, for example, may change cytochrome P450 activity and thereby alter the expected metabolism of a substrate [14]. The grammar treats such patient-state variables as causal modifiers of the mediator rather than as detached demographic annotations. This structure permits a represented interaction to distinguish a direct drug-mediated perturbation from a disease-mediated change, a gene-associated baseline, or a combined context in which several determinants shape the observed phenotype.

The fourth and fifth layers concern exposure and clinical consequence. An exposure-change event must identify the affected drug or active metabolite, the relevant compartment, direction, temporal profile, and evidentiary basis. Increased systemic exposure, decreased systemic exposure, altered tissue exposure, and indeterminate change are separate outcomes. A clinical consequence is then represented only when evidence supports an additional transition from exposure to efficacy, toxicity, treatment failure, tolerability, or monitoring need. This prevents a molecular observation from being promoted directly to a clinical claim and prevents a clinically observed association from being backfilled with an unsupported mechanism. The proposed chain is therefore not a deterministic pipeline. It is an evidence-constrained grammar in which each transition can be observed, modeled, predicted, inferred, contradicted, or left unestablished, and in which absence of support for one link limits the strength of the complete causal interpretation.

Core Vocabulary and Relations of the Causal Grammar

The proposed causal grammar begins with a controlled vocabulary that separates scientific objects from the assertions made about them. Its principal entity classes are drugs and active metabolites, enzymes, transporters, regulatory receptors, anatomical compartments, disposition processes, exposure variables, patient-context states, clinical consequences, evidence objects, and decision-use statements. These classes are deliberately narrower than generic biomedical categories because a useful drug–drug-interaction ontology must identify what is being altered, where the alteration occurs, and which downstream variable may consequently change. Knowledge-graph scholarship supports explicit entity and relation design as a prerequisite for interpretable integration rather than treating graph connectivity itself as meaningful evidence [15]. The proposed vocabulary therefore constrains how entities may participate in an interaction assertion and prevents a molecular target, physiological process, exposure measure, and clinical event from being represented as interchangeable nodes.

Drug roles are defined at the level of an individual assertion. A *perpetrator drug* initiates or contributes to a represented change, whereas an *object drug* or its active metabolite undergoes the affected disposition, exposure, or consequence. These roles are contextual rather than intrinsic: the same compound may be a perpetrator in one pathway and an object in another. The grammar also allows mutual interaction, competing substrates, and multiple perpetrators when supported by evidence. Relations among these entities are typed according to causal level. Examples include *reversibly inhibits*, *induces*, *competes for*, *reduces activity of*, *increases clearance of*, *decreases systemic exposure to*, and *is associated with increased toxicity of*. Preserving such heterogeneous relation types is essential because biomedical integration loses scientific meaning when distinct biological and clinical relations are collapsed into a generic connection [16].

The grammar distinguishes asserted relations from inferred relations. An asserted relation is traceable to a study, model, curated source, or explicit textual statement. An inferred relation is generated through a formally stated composition rule, such as combining an enzyme-inhibition assertion with evidence that the object drug is substantially cleared through that enzyme.

Both relation classes require provenance, but inferred relations additionally require a record of the rule, input assertions, graph version, and uncertainty. Evaluation must be relation-specific because a benchmark that mixes easy and difficult relation types may obscure failures in clinically consequential predicates. Biomedical link-prediction frameworks demonstrate the importance of standardized relation definitions, controlled evaluation splits, and leakage-aware assessment [17]. Nevertheless, relation-prediction performance remains an informatics measure and cannot alone establish that the inferred path is pharmacologically causal.

Composition rules specify which relations may be connected. A molecular-initiation event may connect to a mediator only when the affected enzyme, transporter, or regulatory system is identified or explicitly marked as unknown. A mediator may connect to a disposition process only when the anatomical and physiological relationship is supported. A disposition process may connect to an exposure-change event only when the affected drug, metabolite, or compartment is named. A clinical consequence requires a further evidence-bearing transition and cannot be produced solely from an interaction prediction. Pharmacological knowledge graphs illustrate how typed entities and relations can support reproducible computational reasoning [18]. Relation-aware DDI models likewise show the value of retaining interaction-specific structure rather than reducing all drug pairs to a single binary predicate [19]. The proposed grammar extends this principle by imposing causal-level constraints, evidence-state labels, and explicit prohibitions against unsupported semantic shortcuts. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop core vocabulary and relations of the causal grammar within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Core vocabulary and relations of the causal grammar

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Interaction assertion	Two or more identified drugs, an asserted relation, source, and context	Encapsulates a versioned DDI claim	Traceable interaction statement	Source may be incomplete, duplicated, or incorrect	Provenance audit and expert review
Context-specific drug roles	Perpetrator, object, mutual interactor, substrate, inhibitor, or inducer designation	Establishes causal orientation	Directionally interpretable role assignment	Roles may change across mechanisms, doses, or time points	Mechanism-specific competency testing
Molecular-initiation event	Binding, inhibition, induction, competition, or transporter-modulation evidence	Defines the earliest supported perturbation	Typed initiation relation	In vitro evidence may not translate in vivo	Experimental-context and translation assessment
Biological mediator	Enzyme, uptake transporter, efflux transporter, receptor, or coupled system	Connects initiation to physiology	Identified affected mediator	Several mediators may contribute simultaneously	Pharmacological expert adjudication
Anatomical location	Intestine, liver, kidney, blood–brain barrier, or other tissue	Localizes mediator and process	Site-specific causal relation	Location may be inferred rather than observed	Anatomical and physiological consistency review
Disposition process	Absorption, extraction, metabolism, secretion, excretion, or distribution evidence	Represents the affected handling process	Process-level transition	Process contribution may be model dependent	Comparison with mechanistic and clinical PK evidence
Exposure-change event	Observed or modeled systemic or tissue concentration information	Expresses direction and, when supported, magnitude of exposure alteration	Qualified exposure assertion	Magnitude, timing, or compartment may be uncertain	External PK evaluation and calibration
Clinical consequence	Pharmacodynamic, efficacy, toxicity, tolerability, or monitoring evidence	Represents bounded downstream relevance	Typed consequence assertion	Exposure change may not translate into an observable outcome	Clinical-outcome or adjudicated evidence
Direction qualifier	Increased, decreased, unchanged, bidirectional, or indeterminate status	Preserves polarity of the affected variable	Directionally explicit relation	Direction may differ by endpoint or time	Endpoint-specific review
Magnitude qualifier	Qualitative category or source-reported quantitative estimate	Characterizes interaction strength	Magnitude-bearing assertion	Measures may not be comparable across designs	Metric definition and uncertainty assessment
Temporal qualifier	Onset, sequence, duration, persistence, reversibility, washout, or steady state	Represents temporal causality	Time-bounded relation	Timing is often incompletely reported	Temporal competency questions and case validation
Dose and regimen context	Dose, route, formulation, frequency, duration, and administration sequence	Restricts the conditions under which an assertion holds	Regimen-bounded interaction statement	Extrapolation beyond studied regimens may be unjustified	Cross-regimen evaluation
Patient-context state	Genotype, phenotype, age, organ function, inflammation, comorbidity, pregnancy, or population	Represents modifiers of mediator activity or clinical susceptibility	Context-bounded applicability statement	Patient state can change over time	Longitudinal and population-specific validation

Evidence assertion	Study, model, report, extraction, or curated statement	Separates evidence from the causal relation it supports	Evidence-linked graph object	Evidence quality and independence vary	Source-level quality and duplication assessment
Provenance record	Source identifier, design, version, curator, extraction method, and transformation history	Enables reproducibility and updating	Auditable assertion history	Complete provenance does not guarantee correctness	Independent reconstruction audit
Contradiction relation	Two or more incompatible evidence assertions	Preserves disagreement in direction, magnitude, context, or mechanism	Explicit conflict representation	Apparent contradiction may reflect context differences	Context-sensitive adjudication
Applicability boundary	Population, dose, mechanism, tissue, or evidence limitation	Restricts valid inference	Bounded reasoning result	Missing boundaries may be misread as universality	External-validity assessment
Decision-use statement	Qualified evidence chain, intended user, and workflow context	Describes a limited permissible use	Human-readable bounded interpretation	May be mistaken for a recommendation	Human-factors and safety evaluation

Representing Direction, Magnitude, Timing, Dose, and Patient Context

Directionality is a foundational requirement because the phrase *drug A interacts with drug B* conceals several distinct orientations. The perpetrator may increase or decrease mediator activity; the resulting process may increase or decrease clearance; and the final exposure effect may move in the opposite direction from the process effect. For example, reduced metabolic activity may lower clearance while increasing systemic exposure. Direction must therefore be attached to a named property at each causal level rather than inherited from a generic interaction edge. Temporal induction further demonstrates that direction cannot be separated from onset and duration. Induction develops and resolves over time, and its interpretation depends on the relevant enzyme-turnover process, dosing history, model assumptions, and observation window [20]. The grammar consequently represents direction as a local relation property and does not assume that one polarity applies uniformly across the entire chain.

Magnitude is represented separately from direction. The grammar permits a qualitative description when numerical evidence is unavailable, but any quantitative assertion must retain the measured or modeled endpoint, units, uncertainty, study design, and source. It must also identify whether magnitude concerns molecular inhibition, enzyme activity, clearance, concentration–time exposure, or a clinical consequence. These values cannot be treated as interchangeable. Patient phenotype can change the magnitude of an interaction even when genotype remains constant. Controlled clinical evidence indicates that co-medication can shift CYP2C19 activity away from the phenotype expected from genotype alone [21]. The grammar therefore distinguishes inherited genotype, inferred baseline phenotype, currently observed phenotype, and drug-induced phenoconversion as separate but related states.

Population and physiological context are required because interaction magnitude is not stable across all patients. Age, pregnancy, renal or hepatic impairment, inflammatory disease, comorbidity, and altered protein binding may change baseline physiology, perpetrator exposure, mediator activity, object-drug clearance, or susceptibility to downstream effects. Evidence concerning special populations shows that interaction magnitude may differ from that observed in healthier or more narrowly defined study groups [22]. These differences should not be handled merely by appending a population label to a universal interaction record. Instead, the patient or population state must modify the relevant mediator, disposition, exposure, or consequence relation. An assertion may therefore be valid for one physiological state and unresolved, attenuated, or intensified in another.

Dose, route, administration sequence, duration, and co-medication history provide the remaining conditions under which a DDI assertion holds. Gene–drug–drug interaction modeling illustrates that genotype, inhibitor exposure, and substrate disposition may operate jointly rather than as independent modifiers [23]. Phenoconversion evidence also shows that current metabolic behavior can diverge from genotype-predicted activity because of interacting medications, disease, and environmental influences [24]. The proposed grammar represents these determinants compositionally: genotype may establish a baseline mediator state; disease or inflammation may alter that state; a perpetrator drug may produce an additional perturbation; and the administered regimen determines the temporal and quantitative expression of the interaction. This approach avoids treating patient context as a flat list of covariates and instead locates each contextual variable at the causal transition it modifies.

Reasoning Across Mechanistic and Clinical Evidence

Reasoning within the proposed grammar is evidence constrained. A computational system may identify a plausible path from a drug to an enzyme, from that enzyme to an object drug, and from the object drug to a reported adverse effect. Such a path may be useful for hypothesis generation, but its components can originate from evidence with different designs, populations, directions, and levels of certainty. Local knowledge-graph subgraphs can improve the efficiency and interpretability of typed interaction prediction by restricting attention to relations surrounding a candidate drug pair [25]. In the causal grammar, however, a local subgraph is classified as an explanatory computational object rather than a confirmed mechanistic chain. It must display the assertions on which it depends, identify inferred transitions, and expose missing causal levels.

Molecular-feature models occupy another evidence tier. Structural and molecular descriptors can support prediction of drug–drug or drug–food interactions, particularly when interaction records are incomplete [26]. Yet a successful prediction may reflect chemical similarity, shared targets, reporting patterns, or correlated features rather than the actual mechanism responsible for an interaction in a patient. The grammar therefore permits molecular predictions to support a *plausibility relation*, but not an unqualified *causes exposure change* or *causes toxicity* relation. Advancement to stronger causal status requires evidence concerning the mediator, affected disposition process, relevant exposure variable, and context in which the effect occurs.

Clinical and adverse-effect data provide essential downstream evidence but also require careful interpretation. Multi-relational graph models demonstrate that different drug combinations can be associated with different polypharmacy side-effect types rather than one undifferentiated interaction outcome [27]. Local mechanistic paths and global graph structure may also supply complementary forms of computational evidence [28]. The proposed grammar preserves these sources as separate evidence objects and allows them to converge only when their entities, direction, timing, and patient context are compatible. **Figure 1** presents the drug–drug interaction causal grammar, showing how the article's key components and boundaries are connected within reasoning across mechanistic and clinical evidence.

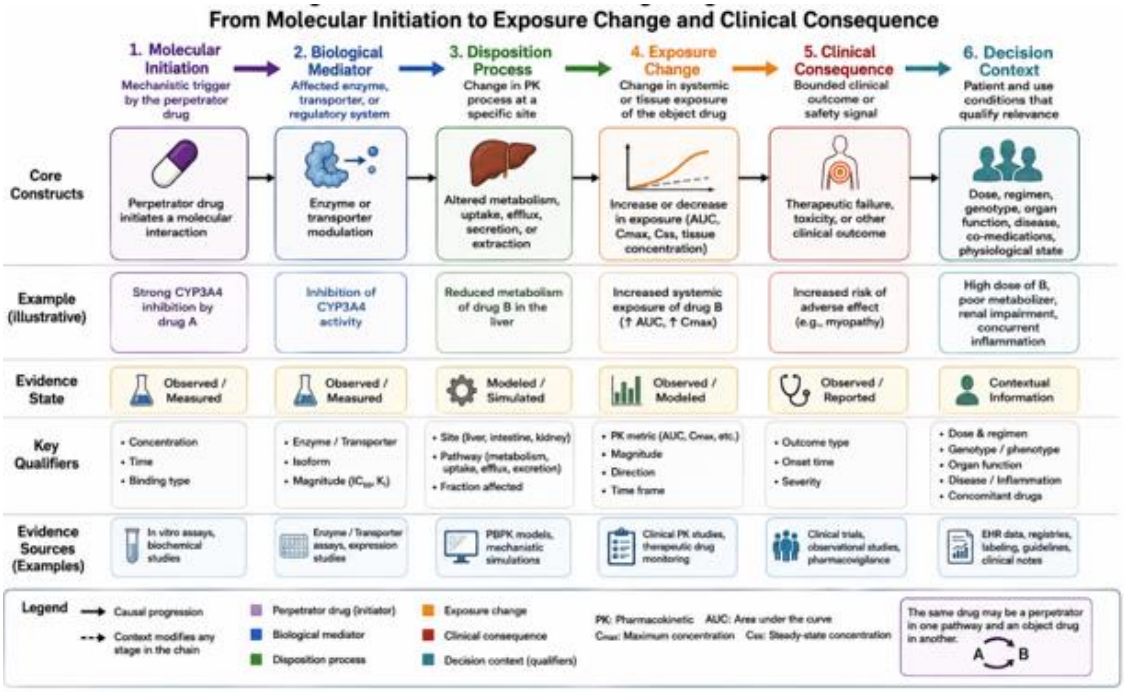

Figure 1. The causal grammar represents drug–drug interactions as evidence-bearing chains. A perpetrator drug initiates a molecular event that affects a mediator (enzyme, transportor, or regulatory system), which alters a disposition process, leading to a change in exposure of the object drug and potentially a clinical consequence. Patient and use context qualifies relevance. Each link carries context-specific qualifiers, evidence states, and sources.

Figure 1. Drug–Drug Interaction Causal Grammar.

The figure is an original conceptual synthesis that organizes the article's central contribution across drug-drug interaction knowledge remains semantically fragmented, molecular initiation to enzymes, transporters, exposure, and outcomes, molecular initiation to enzymes, core vocabulary and relations of the causal grammar, representing direction, magnitude, timing, dose, and patient context, representing direction. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Interpretable knowledge-subgraph methods can expose which graph relations contribute to a predicted interaction [29]. This is valuable for auditing and hypothesis formation, but the displayed subgraph should not be described as a causal explanation merely because it is understandable to a human reader. Within the grammar, reasoning results are assigned bounded evidence statuses: *association supported*, *computationally predicted*, *mechanistically plausible*, *experimentally supported*, *exposure supported*, *clinically associated*, or *clinically adjudicated*. Progression between these states requires evidence appropriate to the claimed level. Contradictory paths remain visible, and an unresolved transition prevents the complete chain from being represented as established.

Interoperability, Validation, and Decision-Support Applications

Interoperability requires more than mapping drug names to shared identifiers. A transferable DDI representation must preserve interacting-drug roles, affected mediators, anatomical locations, disposition processes, exposure variables, temporal conditions, patient states, evidence provenance, and uncertainty. Context-aware clinical decision-support research indicates that interaction screening can become more relevant when patient and medication circumstances are considered rather than

presenting every pairwise warning in the same manner [30]. The proposed grammar may support such contextualization by allowing systems to query not merely whether two drugs are linked, but whether the represented pathway is applicable to the current dose, sequence, organ function, phenotype, and evidence state. This potential use remains conditional on validated mappings, complete input data, and safeguards against suppressing clinically important information.

Validation must proceed in layers. Semantic validation asks whether terms and predicates are correctly defined and whether domain experts can use the grammar to answer pre-specified competency questions. Structural validation asks whether relation constraints prevent invalid causal compositions. Informatics evaluation tests extraction, mapping, contradiction detection, and reasoning under controlled conditions. Mechanistic evaluation determines whether represented paths agree with established pharmacology. Quantitative evaluation examines whether model-linked paths correspond to observed concentration–time changes. Clinical evaluation addresses whether the represented evidence improves interpretation or action in realistic workflows. Override studies demonstrate why alert generation cannot serve as a substitute for clinical usefulness: clinicians may dismiss alerts for reasons ranging from genuine irrelevance to workflow burden or inadequate presentation [31].

Human reasoning and decision provenance should also be represented. Structured categories for DDI-alert overrides can make clinician responses more interpretable and comparable across systems [32]. Within the proposed grammar, an override would not erase the original interaction assertion. Instead, it would be represented as a separate decision event linked to the patient context, alert content, clinician rationale, and subsequent action. This distinction supports auditability and permits examination of whether an alert was irrelevant, previously addressed, accepted with monitoring, or dismissed despite residual risk. The grammar does not assign correctness automatically to either the alert or the override; both remain evidence-bearing objects requiring contextual review.

Patient-specific applicability is particularly important when organ function changes drug disposition. Renal decision-support studies show that alert relevance and override behavior depend on the patient's physiological state and medication circumstances [33]. More broadly, evidence on electronic DDI alerts remains heterogeneous in intervention design, workflow integration, outcomes, and demonstrated effectiveness [34]. Ontology conformance should therefore never be treated as proof of clinical benefit. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop interoperability, validation, and decision-support applications within the article's central argument.

Table 3. Evaluation, implementation, and research priorities arising from A Causal Grammar for Drug–Drug Interactions from Molecular Initiation to Exposure Change

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Vocabulary validity	Whether entity classes and predicates are pharmacologically precise, non-redundant, and understandable	Structured expert review, ontology competency questions, terminology comparison	Ambiguous definitions or inconsistent expert interpretation	Agreement does not prove completeness or clinical value	Determines whether the grammar can support reliable curation
Drug-role validity	Whether perpetrator, object, and mutual roles are assigned correctly under changing mechanisms and regimens	Mechanism-specific curated cases	Reversed or permanently fixed role assignments	Correct role assignment does not establish effect magnitude	Necessary for directional reasoning
Relation-direction validity	Whether subject, predicate, object, and affected variable retain correct polarity	Gold-standard directional assertions and error analysis	Reversed causal arrows or unspecified affected variables	Direction may differ across endpoints and time points	Prevents misleading mechanistic and exposure statements
Causal-chain completeness	Whether initiation, mediator, process, exposure, and consequence levels are represented without unsupported jumps	Expert-adjudicated multi-layer DDI cases	Direct movement from prediction or assay result to clinical consequence	A complete path may still be biologically wrong	Supports bounded mechanistic interpretation
Temporal validity	Whether onset, duration, sequence, persistence, reversibility, and washout are represented accurately	Time-dependent inhibition and induction cases	Timeless representation of transient or delayed effects	Temporal evidence may remain sparse	Supports regimen-specific interpretation
Context adequacy	Whether dose, route, genotype, phenotype, organ function, disease, and population modifiers are captured	Context-rich clinical and modeling cases	Universal assertion despite known contextual dependence	Recorded context may still be incomplete	Determines patient or population applicability
Provenance fidelity	Whether every assertion can be traced to a source, version, method, and transformation	Reproducibility and reconstruction audits	Orphaned edges, lost source history, or duplicated evidence	Traceability does not prove scientific quality	Enables updating, auditing, and evidence review
Contradiction handling	Whether disagreements in mechanism, direction, magnitude, or context remain visible	Discordant evidence sets and adjudication exercises	Silent overwriting or unjustified averaging	Contradiction representation does not resolve truth	Prevents false consensus

Extraction validity	Whether literature-derived relations retain the meaning and limits of the source	Manually curated corpora and relation-level assessment	Correct entity pair but incorrect mechanism or direction	Corpus performance may not generalize	Determines reliability of automated evidence ingestion
Computational reasoning validity	Whether inferred relations obey grammar constraints and avoid shortcut learning or leakage	Relation-specific benchmarks, unseen-drug tests, temporal splits	High aggregate performance with invalid paths	Benchmark success is not mechanistic validation	Supports hypothesis generation only
Mechanistic validity	Whether represented mediators and processes agree with experimental pharmacology	In vitro, in vivo, and expert evidence	Plausible but unsupported or incorrect pathway	Mechanistic evidence may not predict exposure magnitude	Supports experimental prioritization
Exposure validity	Whether represented mechanisms correspond to observed systemic or tissue exposure changes	Clinical PK data and qualified mechanistic models	Correct mechanism but incorrect exposure direction, timing, or magnitude	Exposure agreement does not establish clinical utility	Supports model interpretation and study design
Clinical-consequence validity	Whether downstream efficacy, toxicity, or treatment effects are independently supported	Clinical outcomes, pharmacodynamic evidence, or adjudicated cases	Clinical claim based only on molecular or PK evidence	Confounding and sparse outcomes may limit causal inference	Determines whether consequence statements are defensible
Interoperability	Whether data can be exchanged without loss of direction, context, provenance, or evidence state	Round-trip mappings across terminologies and data models	Relation flattening, identifier collision, or missing qualifiers	Technical exchange does not guarantee shared scientific meaning	Supports multi-resource integration
Human interpretability	Whether domain users understand the represented path, uncertainty, and boundaries	Pharmacist, pharmacologist, modeler, and prescriber studies	Misreading inference as observation or certainty	Understanding does not prove improved decisions	Necessary for transparent decision support
Alert relevance	Whether contextualized representations improve useful alerting without hiding important interactions	Workflow simulation, prospective user evaluation, missed-alert review	Persistent alert fatigue or inappropriate suppression	Process measures may not reflect patient outcomes	Supports cautious CDS development
Governance and versioning	Whether updates, deprecations, mappings, and rule changes are controlled	Change-management audits and governance review	Untraceable semantic drift	Governance cannot compensate for weak evidence	Enables sustainable implementation
Deployment readiness	Whether semantic, technical, human-factors, safety, and clinical requirements have all been met	Staged prospective implementation and safety monitoring	Advancement based on one successful evaluation layer	No single metric establishes readiness	Defines a boundary against premature operational use

Limitations and Boundary Conditions

The principal limitation of the proposed grammar is that formal semantic coherence cannot guarantee biological correctness. A well-formed causal path may connect valid entity and relation types while relying on an incorrect molecular mechanism, an incomplete account of disposition, or evidence that does not apply to the relevant dose or population. Computational generalization is particularly vulnerable when drugs, interaction types, or mechanisms are poorly represented in development data [4]. The grammar can expose missing evidence, provenance, and applicability limits, but it cannot manufacture experimental support where none exists. Its evaluation must therefore include pharmacological and clinical adjudication rather than relying solely on logical consistency or benchmark performance.

A second limitation is the dynamic and context-dependent nature of drug disposition. Enzyme activity, transporter function, organ performance, inflammation, genotype–phenotype relationships, adherence, diet, and concurrent treatment can change over time. Disease-mediated changes in metabolic activity complicate the interpretation of an interaction even when the interacting drugs and nominal doses remain the same [14]. Evidence from special populations further indicates that magnitude and clinical relevance may differ across physiological contexts [22]. The grammar can represent these modifiers only when they are available, sufficiently defined, and linked to the appropriate causal transition. Missing patient information must therefore produce an explicit applicability limitation rather than an assumed default state.

A third limitation concerns intended use. Interpretable subgraphs, evidence paths, or ontology-derived explanations may improve traceability but can still encourage overconfidence if visual coherence is mistaken for causal confirmation [29]. Likewise, electronic alert evidence does not establish that a semantically richer representation will improve outcomes in every clinical environment [34]. The grammar is not a dosing engine, prescribing rule, regulatory classification, or autonomous clinical decision system. It cannot determine whether a medication combination should be initiated, continued, adjusted, monitored, or discontinued without compound-specific evidence, current patient information, qualified professional judgment, and setting-appropriate validation.

Research Agenda and Implementation Priorities

The first research priority is formal ontology development and semantic validation. A multidisciplinary group spanning clinical pharmacology, drug metabolism and pharmacokinetics, pharmacometrics, transporter science, pharmacogenomics, biomedical informatics, pharmacy, and clinical medicine should define competency questions that the grammar must answer. These questions should include whether a represented interaction identifies the perpetrator and object correctly, distinguishes molecular initiation from exposure change, preserves timing and dose, and prevents unsupported movement into a clinical consequence. Evaluation should use relation-specific and leakage-aware designs rather than one aggregate graph score [17]. Negative, contradictory, multi-mechanism, and context-dependent examples are particularly important because they test whether the grammar can represent uncertainty without forcing a single answer.

The second priority is evidence-linked mechanistic and quantitative validation. Curated case sets should span reversible inhibition, time-dependent inhibition, induction, transporter effects, coupled enzyme–transporter processes, gene–drug–drug interactions, disease-mediated phenoconversion, and tissue-specific exposure. For each case, independent experts should compare the represented path with available experimental, PBPK, clinical pharmacokinetic, and outcome evidence. Longitudinal representations are needed because phenotype can change as medication, disease, and physiological state change [24]. Future studies should test whether the grammar improves model documentation, identification of unsupported assumptions, external-validation planning, and recognition of contexts in which quantitative extrapolation is inappropriate.

The third priority is staged interoperability and human-centered implementation. Mappings to established drug, gene, protein, phenotype, anatomical, and clinical terminologies should be tested through round-trip exchange to determine whether direction, qualifiers, provenance, and contradictions survive translation. Prototype decision-support applications should begin in retrospective or simulated environments and should compare explanation quality, alert relevance, clinician comprehension, override rationale, and missed-risk detection. Context-aware alert research provides a useful foundation, but any implementation must evaluate whether increased specificity introduces new omissions or unsafe filtering [30]. Progression to prospective use should occur only after semantic, mechanistic, quantitative, workflow, human-factors, privacy, governance, and safety requirements have been assessed independently.

Conclusion

A drug–drug interaction is not one relation but a conditional sequence of evidence-bearing transitions. The proposed causal grammar makes those transitions explicit by separating context-specific drug roles, molecular initiation events, enzymes and transporters, disposition processes, systemic or tissue exposure changes, and bounded clinical consequences. It further treats direction, magnitude, timing, dose, administration sequence, patient state, evidence provenance, contradiction, uncertainty, and applicability as integral components of interpretation. This architecture is intended to prevent molecular plausibility, computational prediction, quantitative simulation, clinical association, and decision relevance from being collapsed into a single interaction label or performance measure. Its value lies in organizing heterogeneous evidence, exposing unsupported causal jumps, and providing a shared structure for future curation, modeling, validation, and context-aware decision support. The contribution remains a proposed ontology rather than a validated model or operational tool. Its scientific usefulness will depend on formal semantic testing, pharmacological adjudication, quantitative and clinical evaluation, interoperable implementation, transparent governance, and continued human oversight.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Yadav J, Maldonato BJ, Roesner JM, Vergara AG, Paragas EM, Aliwarga T, et al. Enzyme-mediated drug-drug interactions: A review of in vivo and in vitro methodologies, regulatory guidance, and translation to the clinic. Drug Metab Rev. 2025;57(3):211-43. doi:10.1080/03602532.2024.2381021
2. Russell LE, Yadav J, Maldonato BJ, Chien HC, Zou L, Vergara AG, et al. Transporter-mediated drug-drug interactions: Regulatory guidelines, in vitro and in vivo methodologies and translation, special populations, and the blood-brain barrier. Drug Metab Rev. 2025;57(3):244-71. doi:10.1080/03602532.2024.2364591
3. Luo H, Yin W, Wang J, Zhang G, Liang W, Luo J, et al. Drug-drug interactions prediction based on deep learning and knowledge graph: A review. iScience. 2024;27(3):109148. doi:10.1016/j.isci.2024.109148
4. Ahmad MM, Shah ZBA, Nies HW. A systematic review of molecular structures, knowledge graphs, and cold-start scenario in drug-drug interaction prediction. Comput Biol Med. 2025;190:110122. doi:10.1016/j.compbiomed.2025.110122

5. Taneja SB, Callahan TJ, Paine MF, Kane-Gill SL, Kilicoglu H, Joachimiak MP, et al. Developing a knowledge graph for pharmacokinetic natural product-drug interactions. J Biomed Inform. 2023;140:104341. doi:10.1016/j.jbi.2023.104341
6. MacLean F. Knowledge graphs and their applications in drug discovery. Expert Opin Drug Discov. 2021;16(9):1057-69. doi:10.1080/17460441.2021.1910673
7. Serra A, Fratello M, Federico A, Greco D. An update on knowledge graphs and their current and potential applications in drug discovery. Expert Opin Drug Discov. 2025;20(5):599-619. doi:10.1080/17460441.2025.2490253
8. Bonner S, Barrett IP, Ye C, Swiers R, Engkvist O, Bender A, et al. A review of biomedical datasets relating to drug discovery: A knowledge graph perspective. Brief Bioinform. 2022;23(6). doi:10.1093/bib/bbac404
9. Asada M, Miwa M, Sasaki Y. Integrating heterogeneous knowledge graphs into drug-drug interaction extraction from the literature. Bioinformatics. 2023;39(1). doi:10.1093/bioinformatics/btac754
10. Taskar KS, Pilla Reddy V, Burt H, Posada MM, Varma M, Zheng M, et al. Physiologically-based pharmacokinetic models for evaluating membrane transporter mediated drug-drug interactions: Current capabilities, case studies, future opportunities, and recommendations. Clin Pharmacol Ther. 2020;107(5):1082-115. doi:10.1002/cpt.1693
11. Gessner A, König J, Fromm MF. Clinical aspects of transporter-mediated drug-drug interactions. Clin Pharmacol Ther. 2019;105(6):1386-94. doi:10.1002/cpt.1360
12. Alluri RV, Li R, Varma MVS. Transporter-enzyme interplay and the hepatic drug clearance: What have we learned so far? Expert Opin Drug Metab Toxicol. 2020;16(5):387-401. doi:10.1080/17425255.2020.1749595
13. Ren HC, Sai Y, Chen T, Zhang C, Tang L, Yang CG. Predicting the drug-drug interaction mediated by CYP3A4 inhibition: Method development and performance evaluation. AAPS J. 2022;24(1):12. doi:10.1208/s12248-021-00659-w
14. Lenoir C, Rollason V, Desmeules JA, Samer CF. Influence of inflammation on cytochromes P450 activity in adults: A systematic review of the literature. Front Pharmacol. 2021;12:733935. doi:10.3389/fphar.2021.733935
15. Zeng X, Tu X, Liu Y, Fu X, Su Y. Toward better drug discovery with knowledge graph. Curr Opin Struct Biol. 2022;72:114-26. doi:10.1016/j.sbi.2021.09.003
16. Himmelstein DS, Lizee A, Hessler C, Brueggeman L, Chen SL, Hadley D, et al. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. eLife. 2017;6. doi:10.7554/eLife.26726
17. Breit A, Ott S, Agibetov A, Samwald M. OpenBioLink: A benchmarking framework for large-scale biomedical link prediction. Bioinformatics. 2020;36(13):4097-8. doi:10.1093/bioinformatics/btaa274
18. Zheng S, Rao J, Song Y, Zhang J, Xiao X, Fang EF, et al. PharmKG: A dedicated knowledge graph benchmark for biomedical data mining. Brief Bioinform. 2021;22(4). doi:10.1093/bib/bbaa344
19. Su X, Hu L, You ZH, Hu P, Zhao B. Attention-based knowledge graph representation learning for predicting drug-drug interactions. Brief Bioinform. 2022;23(3). doi:10.1093/bib/bbac140
20. Hariparsad N, Ramsden D, Taskar K, Badée J, Venkatakrishnan K, Reddy MB, et al. Current practices, gap analysis, and proposed workflows for PBPK modeling of cytochrome P450 induction: An industry perspective. Clin Pharmacol Ther. 2022;112(4):770-81. doi:10.1002/cpt.2503
21. Abouir K, Exquis N, Gloor Y, Daali Y, Samer CF. Phenoconversion due to drug-drug interactions in CYP2C19 genotyped healthy volunteers. Clin Pharmacol Ther. 2024;116(4):1121-9. doi:10.1002/cpt.3378
22. Bettonte S, Berton M, Marzolini C. Magnitude of drug-drug interactions in special populations. Pharmaceutics. 2022;14(4):789. doi:10.3390/pharmaceutics14040789
23. Storelli F, Desmeules J, Daali Y. Physiologically-based pharmacokinetic modeling for the prediction of CYP2D6-mediated gene-drug-drug interactions. CPT Pharmacometrics Syst Pharmacol. 2019;8(8):567-76. doi:10.1002/psp4.12411
24. Klomp SD, Manson ML, Guchelaar HJ, Swen JJ. Phenoconversion of cytochrome P450 metabolism: A systematic review. J Clin Med. 2020;9(9):2890. doi:10.3390/jcm9092890
25. Yu Y, Huang K, Zhang C, Glass LM, Sun J, Xiao C. SumGNN: Multi-typed drug interaction prediction via efficient knowledge graph summarization. Bioinformatics. 2021;37(18):2988-95. doi:10.1093/bioinformatics/btab207
26. Ryu JY, Kim HU, Lee SY. Deep learning improves prediction of drug-drug and drug-food interactions. Proc Natl Acad Sci U S A. 2018;115(18). doi:10.1073/pnas.1803294115
27. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics. 2018;34(13). doi:10.1093/bioinformatics/bty294
28. Ren ZH, You ZH, Yu CQ, Li LP, Guan YJ, Guo LX, et al. A biomedical knowledge graph-based method for drug-drug interactions prediction through combining local and global features with deep neural networks. Brief Bioinform. 2022;23(5). doi:10.1093/bib/bbac363
29. Wang Y, Yang Z, Yao Q. Accurate and interpretable drug-drug interaction prediction enabled by knowledge subgraph learning. Commun Med (Lond). 2024;4(1):59. doi:10.1038/s43856-024-00486-y
30. Muylle KM, Gentens K, Dupont AG, Cornu P. Evaluation of an optimized context-aware clinical decision support system for drug-drug interaction screening. Int J Med Inform. 2021;148:104393. doi:10.1016/j.ijmedinf.2021.104393
31. Nanji KC, Seger DL, Slight SP, Amato MG, Beeler PE, Her QL, et al. Medication-related clinical decision support alert overrides in inpatients. J Am Med Inform Assoc. 2018;25(5):476-81. doi:10.1093/jamia/ocx115

32. Wright A, McEvoy DS, Aaron S, McCoy AB, Amato MG, Kim H, et al. Structured override reasons for drug-drug interaction alerts in electronic health records. J Am Med Inform Assoc. 2019;26(10):934-42. doi:10.1093/jamia/ocz033
33. Shah SN, Amato MG, Garlo KG, Seger DL, Bates DW. Renal medication-related clinical decision support alerts and overrides in the inpatient setting. J Am Med Inform Assoc. 2021;28(6):1081-7. doi:10.1093/jamia/ocaa222
34. Holbrook AM, Matos Silva J, Faruque JAY, Deng J, Schneider T, Jaffer A. Electronic drug-drug interaction alerts in clinical practice: A systematic review. J Am Med Inform Assoc. 2025;32(10):1617-28. doi:10.1093/jamia/ocaf139