

PREDICTING SEVERE DRUG–DRUG INTERACTIONS USING POLYPHARMACY, PHARMACOKINETIC PATHWAYS, AND ADVERSE EVENT REPORTS

Hiroshi Nakamura^{1*}, Yuta Kato¹

1. Department of Intelligent Pharmaceutical Engineering, Faculty of Pharmaceutical Sciences, Nagoya University, Nagoya, Japan.

ARTICLE INFO

Received:

26 January 2025

Received in revised form:

25 March 2025

Accepted:

27 March 2025

Available online:

28 April 2025

Keywords: Drug–drug interaction prediction, Polypharmacy, Pharmacovigilance, Pharmacokinetics, FAERS, Machine learning

ABSTRACT

Severe drug–drug interactions remain a major source of preventable patient harm and healthcare burden. Conventional systems often rely on static pairwise interaction tables and therefore struggle to reflect the complexity of real-world polypharmacy. Existing prediction tools insufficiently integrate pharmacokinetic mechanisms, spontaneous adverse event evidence, and the patient's full medication context. This gap can contribute to both missed severe interactions and excessive low-priority alerts. The objective is to develop a conceptual machine learning model that predicts the severity of a potential drug–drug interaction for a given drug pair or multi-drug regimen. The model would use features derived from polypharmacy exposure, pharmacokinetic pathway overlap, and real-world adverse event reports. A gradient-boosted classification model would be constructed using structured predictors that represent concomitant drug burden, CYP450 and transporter pathway overlap, and pharmacovigilance signal measures from FAERS or Vigibase. The target label would represent serious clinical consequences such as hospitalization, death, or other medically significant outcomes. Conceptually, the model could identify severe interactions that are incompletely represented in standard compendia, particularly when risk emerges from multi-drug exposure rather than a single pairwise mechanism. It would also provide an interpretable decomposition of risk across pharmacokinetic, polypharmacy, and real-world evidence components. A severity-focused predictive model could support safer prescribing by prioritizing clinically urgent interactions for review. Such an approach could reduce alert fatigue by distinguishing high-concern interaction signals from lower-severity flags.

This is an open-access article distributed under the terms of the [Creative Commons Attribution-Non Commercial-Share Alike 4.0 License](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows others to remix, and build upon the work non commercially.

To Cite This Article: Nakamura H, Kato Y. Predicting Severe Drug–Drug Interactions Using Polypharmacy, Pharmacokinetic Pathways, and Adverse Event Reports. *Pharmacophore*. 2025;16(2):43-52. <https://doi.org/10.51847/DPN1uRuBbP>

Introduction

Severe drug–drug interactions are a persistent challenge in medication safety because patients with multimorbidity often receive several medications whose combined effects are difficult to anticipate. Predictive modeling has become increasingly relevant as computational studies have shown that interaction risk can be inferred from chemical, biological, phenotypic, and network-based drug features [1]. Polypharmacy side effects have also been modeled as graph-structured phenomena, suggesting that harmful outcomes may emerge from relationships among multiple medications rather than from isolated drug pairs alone [2]. A predictive model for severe interactions should therefore account for both the presence of specific interacting drugs and the broader medication network in which they occur.

Current interaction knowledge bases are useful for alerting clinicians, but they are often static, pairwise, and limited in their ability to represent patient-specific severity. Approaches using structural similarity and interaction networks have demonstrated that known pharmacokinetic and pharmacodynamic knowledge can enrich DDI prediction beyond simple lookup tables [3]. However, models based only on known pairwise relationships may miss clinically important risk when multiple drugs share toxicities, metabolic pathways, or transporter dependencies. This limitation motivates a severity-calibrated framework that evaluates a regimen as a contextual exposure rather than a set of independent pairs.

Large-scale pharmacovigilance databases and electronic health records provide real-world evidence on adverse outcomes that may not be fully captured during premarketing evaluation. Clinical side-effect profiles have been used as predictive features for DDI discovery, showing that adverse event patterns can encode clinically meaningful relationships between drugs [4]. Genetic interaction information and adverse drug interaction labels have also been used to improve prediction, indicating that

Corresponding Author: Hiroshi Nakamura; Department of Intelligent Pharmaceutical Engineering, Faculty of Pharmaceutical Sciences, Nagoya University, Nagoya, Japan. E-mail: hiroshi.nakamura@outlook.com

real-world and biological evidence can be complementary [5]. Combining these resources with mechanistic pharmacology could support earlier recognition of severe interaction signals.

This manuscript proposes a conceptual predictive model that synthesizes polypharmacy burden, pharmacokinetic pathway overlap, and spontaneous adverse event report signals to estimate the severity of potential DDIs. Deep learning models have improved prediction of drug–drug and drug–food interactions by learning latent relationships from heterogeneous drug information [6]. Matrix factorization and graph-based approaches have further shown that interaction patterns can be learned from relational data structures [7, 8]. The intended contribution of the proposed model is not to replace clinical judgment but to provide a more comprehensive, interpretable, and clinically actionable risk assessment.

Background

Mechanisms of Drug–Drug Interactions

Drug–drug interactions can arise through pharmacokinetic mechanisms, such as CYP450 inhibition, induction, or transporter competition, and through pharmacodynamic mechanisms, such as additive or synergistic toxicity. Mechanism-based inhibition of CYP enzymes is especially important because it can produce sustained exposure changes that persist beyond simple co-administration [9]. Pharmacodynamic interaction frameworks emphasize that combined drug effects may intensify toxicity even when pharmacokinetic exposure changes are modest [10]. These mechanistic principles suggest that many severe DDIs are at least partially predictable when drug metabolism, transport, and shared toxicity pathways are represented as model features.

Polypharmacy and Interaction Risk

Polypharmacy increases interaction risk because each additional drug can create new pairwise and higher-order combinations, especially when medications share metabolic enzymes or adverse effect profiles. Graph convolutional modeling of polypharmacy side effects has shown that multi-drug adverse outcomes can be represented as relational patterns across drugs and side effects [2]. Network and semantic similarity approaches further indicate that the topology of drug relationships can contribute useful information for predicting potential interactions [11]. A polypharmacy interaction network therefore offers a more realistic representation of risk than a list of independent drug pairs.

As summarized in **Table 1**, a polypharmacy interaction network can capture pairwise combinations, higher-order medication patterns, shared metabolic pathways, overlapping adverse effect profiles, and drug–side effect relationships, providing a more realistic representation of interaction risk than treating each drug pair as independent.

Table 1. Key features of a polypharmacy interaction network for DDI risk prediction

Network feature	Relevance to polypharmacy risk	Example interpretation
Pairwise drug combinations	Each added medication creates additional possible drug–drug interactions.	A 5-drug regimen produces more potential interaction pairs than a 2-drug regimen.
Higher-order combinations	Some adverse outcomes may arise from the combined effect of three or more drugs rather than a single pair.	Multiple CNS depressants may jointly increase sedation or fall risk.
Shared metabolic pathways	Drugs using the same enzymes may compete for metabolism or alter exposure levels.	Several drugs metabolized by CYP enzymes may increase toxicity risk.
Overlapping adverse effect profiles	Drugs with similar toxicities may produce additive or synergistic harm.	Two QT-prolonging drugs may increase arrhythmia risk.
Drug–side effect relationships	Graph-based models can represent links between drugs and observed adverse outcomes.	A model may connect multiple drugs to a shared side effect node.
Network topology and semantic similarity	Structural relationships among drugs can help identify potential interactions not obvious from pairwise lists.	Drugs close together in a similarity network may share mechanisms or adverse outcomes.
Relational risk patterns	Polypharmacy risk can be modeled as patterns across drugs, mechanisms, and side effects.	A cluster of medications may indicate elevated risk for a specific adverse event.

Adverse Event Report Databases and Signal Detection

Spontaneous adverse event reporting systems can reveal suspected interaction harms that occur in routine practice, including outcomes influenced by age, comorbidity, and concomitant medications. A clinically relevant reference set of adverse drug–drug interactions illustrates the importance of distinguishing signals with plausible clinical relevance from noisy associations in pharmacovigilance data [12]. Signal detection methods using spontaneous reports can prioritize possible safety concerns, but they remain vulnerable to under-reporting, confounding, duplicate reports, and reporting stimulation. A predictive model should therefore treat FAERS or Vigibase signals as informative evidence rather than definitive proof of causation.

Machine Learning in Drug Interaction Prediction

Machine learning methods for DDI prediction have progressed from similarity-based frameworks to deep neural networks, heterogeneous graph models, and multimodal fusion approaches. Neural network models using integrated similarity have

shown that diverse similarity measures can improve DDI prediction when combined within a learning framework [13]. Multimodal deep learning has extended this idea by integrating drug features from multiple sources to predict interaction events [14]. Nevertheless, many prior models focus on whether an interaction exists rather than whether the interaction is likely to be clinically severe, leaving an important gap for severity-oriented pharmacovigilance modeling.

Clinical Decision Support for DDIs

Clinical decision support systems can reduce preventable medication harm, but excessive low-specificity interaction alerts may contribute to alert fatigue and override behavior. Multitask deep learning for DDI prediction suggests that models can learn multiple interaction-related outcomes simultaneously, which could support more differentiated alert categories [15]. Automated and graph-based models also imply that risk stratification can be updated as new relationships are learned from expanding drug knowledge bases [16]. A clinically useful system should therefore prioritize severe, actionable interactions while presenting lower-concern signals in a less disruptive form.

Model Development Overview

High-Level Predictive Pipeline

For a given drug pair or medication regimen, the proposed pipeline would construct a feature vector from the patient's full medication list, the pharmacokinetic pathway profiles of the drugs, and historical adverse event signals for the combination. Semi-nonnegative matrix factorization has shown that comprehensive DDI prediction can benefit from representing interaction structure in a way that remains interpretable to domain experts [7]. The model would output a calibrated probability that the interaction could be severe, with severity defined conceptually by outcomes such as hospitalization, death, or an equivalent serious clinical event. This pipeline would support both pairwise predictions and contextual predictions that reflect the surrounding regimen.

Figure 1 presents the proposed severity-calibrated DDI prediction architecture, showing how medication exposure context, pharmacokinetic pathway overlap, adverse event signals, uncertainty encoding, machine learning classification, interpretability, and clinical decision support are connected in a single medication-safety workflow.

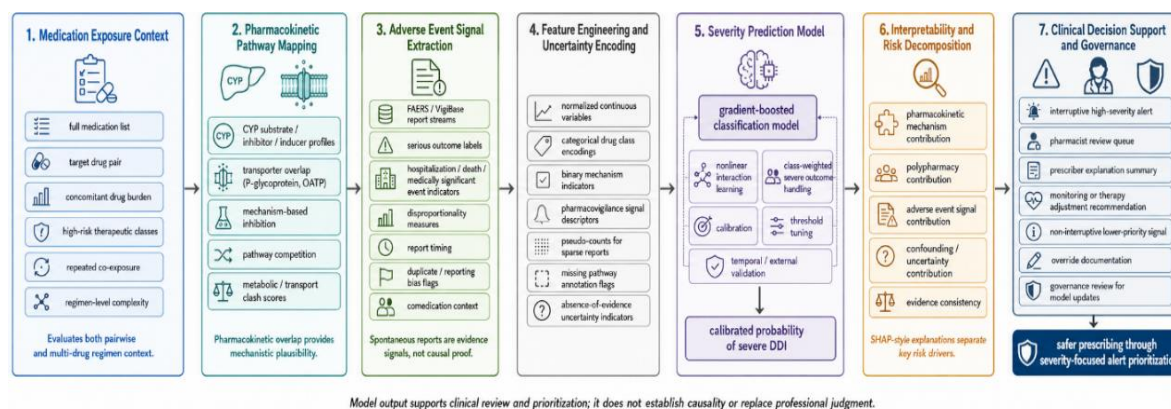


Figure 1. Severity-Calibrated Machine Learning Architecture for Predicting Severe Drug-Drug Interactions from Polypharmacy Context, Pharmacokinetic Pathways, and Adverse Event Signals

Core Input Feature Groups

The core input features would include polypharmacy indicators, pharmacokinetic mechanism descriptors, and pharmacovigilance signal variables. Substructure-based DDI prediction has shown that molecular components can help identify interaction patterns, while multi-scale feature fusion demonstrates the value of combining feature types across levels of representation [17, 18]. Pharmacokinetic features would encode CYP substrate, inhibitor, and inducer relationships, as well as transporter substrate and inhibitor overlap. Pharmacovigilance features would represent disproportionality signals and report characteristics while being adjusted conceptually for demographic and medication-context confounding.

Design Principles

The design principles are patient-contextual prediction, interpretability, periodic updating, and alert prioritization. Heterogeneous information network models illustrate how diverse drug information can be integrated into a unified predictive architecture [19]. Link-aware graph attention and pre-trained heterogeneous graph neural network approaches also show how relational structure can be exploited while preserving the ability to emphasize influential connections [20, 21]. In a clinical setting, the model should provide explanations that pharmacists and prescribers can understand, rather than functioning as an opaque warning engine.

Data Sources and Feature Engineering

Polypharmacy and Drug-Exposure Data

Polypharmacy and drug-exposure features would be derived from medication lists in electronic health records, pharmacy claims, or medication reconciliation systems. These features could include total concomitant drug burden, the presence of high-risk therapeutic classes, shared toxicity profiles, and indicators of repeated co-exposure across patient populations. Adverse interaction prediction using matrix tri-factorization demonstrates that adverse DDI labels and drug attributes can be integrated into a structured feature space [22]. Feature engineering should also capture the regimen context because the risk of a severe interaction may depend on multiple co-administered drugs rather than the target pair alone.

Pharmacokinetic Pathway Feature Construction

Pharmacokinetic pathway construction would compile drug-level annotations for CYP450 substrate, inhibitor, and inducer activity, together with transporter profiles such as P-glycoprotein and organic anion transporting polypeptide involvement. Prediction methods incorporating pharmacokinetic and pharmacodynamic knowledge indicate that mechanistic drug features can strengthen interaction inference [3]. Mechanism-based CYP inhibition literature further supports representing inhibitory strength, reversibility, and persistence as conceptually distinct feature dimensions [9]. The resulting pathway matrix would allow computation of overlap scores that quantify metabolic or transport “clash” potential for each drug pair or regimen.

Adverse Event Signal Extraction from FAERS/VigiBase

Adverse event signal extraction would query spontaneous reporting systems for reports involving the drug pair or regimen and clinically severe outcomes. Timing-aware pharmacovigilance methods show that the temporal relationship between co-exposure and adverse event reporting can support prioritization of interaction signals [23]. Clinical discovery work combining spontaneous reporting systems and electronic health records also suggests that known pharmacokinetic mechanisms can help contextualize signals found in real-world data [24]. The model would use disproportionality-derived features as evidence of reported harm while explicitly accounting for confounding by comedications, indication, reporter type, and patient characteristics.

Table 2 defines the feature architecture required to convert polypharmacy exposure, pharmacokinetic mechanisms, adverse event evidence, patient-context proxies, and uncertainty indicators into a severity-calibrated DDI prediction model.

Table 2. Feature Architecture for Severity-Calibrated Drug–Drug Interaction Prediction

Feature domain	Representative variables	Analytical purpose	Severe-DDI relevance	Interpretation value for clinicians
Polypharmacy exposure context	Total concomitant medication count; target drug pair; number of high-risk therapeutic classes; repeated co-exposure indicators; regimen complexity index	Represents the medication environment in which the interaction occurs rather than treating the pair in isolation	Severe harm may emerge when several drugs contribute overlapping metabolic, toxic, or physiologic burdens	Helps clinicians see whether the alert is driven by the full regimen rather than a single drug pair
Pairwise interaction structure	Known DDI label presence; drug-pair frequency; historical co-prescribing patterns; similarity to known severe pairs	Establishes baseline interaction plausibility and prior evidence	Allows comparison between established interaction knowledge and newly predicted severe risk	Supports recognition of whether the model is reinforcing known risk or identifying an underrepresented signal
CYP450 pathway overlap	Shared CYP substrate status; inhibitor-substrate combinations; inducer-substrate combinations; inhibitory strength; mechanism-based inhibition flag	Encodes metabolic competition, exposure elevation, or exposure reduction mechanisms	Strong CYP inhibition or induction can produce clinically serious concentration changes	Provides mechanistic explanation for why drug exposure may become unsafe
Transporter pathway overlap	P-glycoprotein substrate/inhibitor overlap; OATP involvement; renal transporter competition; transporter uncertainty flag	Captures non-CYP pharmacokinetic pathways that can alter tissue distribution or clearance	Transporter-mediated effects may intensify toxicity or reduce therapeutic control	Makes non-obvious transporter-mediated interaction pathways visible to pharmacists
Shared pharmacodynamic toxicity	QT prolongation liability; bleeding risk; CNS depression; nephrotoxicity; hepatotoxicity; serotonergic burden	Represents additive or synergistic toxicity independent of pharmacokinetic change	Severe outcomes may occur even when drug levels are not substantially altered	Helps clinicians distinguish exposure-mediated risk from shared toxicity risk
Pharmacovigilance disproportionality	Reporting odds ratio; proportional reporting ratio; information component; severe outcome report	Converts spontaneous report evidence into structured signal features	Serious adverse event clustering can indicate clinically important interaction concern	Shows whether the alert is supported by observed harm in real-world reporting systems

	count; outcome seriousness category			
Report-context and bias indicators	Reporter type; geographic source; drug age; duplicate-report flag; stimulated reporting period; report completeness score	Prevents overinterpretation of noisy spontaneous reporting data	Apparent severity may reflect reporting behavior rather than true interaction risk	Encourages cautious interpretation of uncertain or potentially biased signals
Patient and exposure context proxies	Age group; renal/hepatic impairment proxy; comorbidity burden; indication class; exposure timing; medication start sequence	Reduces confounding by indication and baseline disease severity	Complex patients may have high baseline risk unrelated to the DDI itself	Helps clinicians judge whether the prediction reflects interaction risk or underlying patient vulnerability
Sparse-data uncertainty features	No-report indicator; missing pathway annotation flag; pseudo-count features; low-exposure combination indicator	Separates absence of evidence from evidence of low risk	Rare or new drug combinations may be dangerous despite limited reports	Prevents false reassurance when data are incomplete
Outcome label structure	Hospitalization; death; life-threatening event; medically significant event; severe intervention requirement	Defines the target as severity, not merely interaction existence	Aligns prediction with clinically urgent harm rather than minor interaction flags	Supports alert prioritization according to clinical consequence

Predictive Model Architecture

Model Choice and Rationale

A gradient-boosted tree model would be an appropriate primary architecture because it can handle mixed feature types, nonlinear interactions, missingness, and clinically interpretable feature contributions. Ensemble deep neural network approaches have also been applied to DDI prediction, showing that model combinations can capture complex interaction patterns from heterogeneous drug data [25]. However, a gradient-boosted framework may be preferable for an initial clinical prediction model because it supports transparent feature inspection and flexible calibration without requiring an opaque representation. Deep or graph-based components could later be added as auxiliary feature generators if they improve conceptual interpretability and clinical usefulness.

Input Feature Vector and Preprocessing

The input vector would combine normalized continuous variables, categorical drug class encodings, binary mechanism indicators, and pharmacovigilance signal descriptors. Convolutional and automated graph neural network approaches show that learned representations can extract useful drug-interaction features from molecular or network data, but structured preprocessing remains important when integrating clinical and pharmacovigilance variables [16, 26]. Drug pairs with no spontaneous reports should not be treated as risk-free; instead, the model could use priors, pseudo-counts, or uncertainty indicators to distinguish absence of evidence from evidence of absence. Missing pathway annotations would similarly be encoded as uncertainty rather than silently imputed as non-interaction.

Output: Severe Interaction Probability

The model output would be a calibrated probability that a drug pair or regimen could produce a serious adverse event under the patient's current medication context. Similarity-based safety signal frameworks demonstrate that predictions can be framed as prioritization tools rather than deterministic causal claims [27]. Reviews of DDI databases, web servers, and computational models emphasize the importance of comparing computational predictions with established knowledge resources while recognizing that knowledge bases may be incomplete [28]. The classification threshold would therefore be tuned conceptually according to clinical tolerance for false positives, alert burden, and the consequences of missing a severe interaction.

Handling Data Imbalance, Confounding, And Polypharmacy Complexity

Addressing Severe Interaction Rarity

Severe DDI outcomes are expected to be uncommon relative to the large number of possible drug combinations, so the model should be designed to avoid learning only the majority non-severe pattern. Methods such as class-weighted learning, focal loss, or clinically guided sampling could help the model attend to severe interaction signals without overstating weak associations. PASS-based interaction prediction demonstrates that rare but biologically plausible interaction mechanisms can still be captured when pharmacological knowledge is represented systematically [29]. The model should also preserve uncertainty for sparse combinations so that rare severe signals are prioritized for review rather than converted into unsupported certainty.

Adjusting for Confounding by Indication and Polypharmacy

Confounding by indication is a central challenge because patients receiving complex regimens may already have severe disease, organ dysfunction, or high baseline adverse event risk. A DDI prediction model should therefore include patient-level context such as comorbidity indicators, medication class proxies, and exposure timing to reduce the chance that disease severity is mistaken for interaction severity. Domain adaptation approaches for DDI prediction suggest that models can be structured to generalize across data sources while accounting for differences in observed populations and reporting environments [30]. In the proposed framework, confounding adjustment would support more cautious interpretation of pharmacovigilance signals and prevent the model from over-prioritizing drug combinations merely because they are used in medically complex patients. As shown in **Figure 2**, incorporating patient-level context, medication proxies, and exposure timing can help distinguish disease-related adverse event risk from true interaction-related risk, reducing the likelihood that drug combinations used in medically complex patients are incorrectly prioritized as high-risk DDIs.




Figure 2. Conceptual framework for confounding adjustment in DDI prediction

Extending from Pairwise to Regimen-Level Risk

A regimen-level risk score would aggregate pairwise predictions while accounting for overlapping mechanisms, shared toxicities, and repeated pathway conflicts across the full medication list. Graph convolutional approaches to DDI prediction show that drug relationships can be modeled as connected structures rather than isolated comparisons, which supports extension from pairwise risk to network-level risk [31]. Size-adaptive molecular substructure modeling also indicates that interaction-relevant information may vary across drugs and mechanisms, making rigid pairwise scoring insufficient for complex regimens [32]. The proposed aggregation strategy should therefore avoid double-counting similar mechanisms while still recognizing that multiple modest risks may combine into a clinically serious interaction concern.

Model Interpretability and Clinical Acceptability

SHAP-Based Explanation of Interaction Severity

For each flagged interaction, the model should provide an explanation that identifies the major drivers of predicted severity, such as strong CYP inhibition, transporter overlap, high polypharmacy burden, or a consistent adverse event signal. Deep multimodal fusion models demonstrate the value of integrating heterogeneous evidence streams, but clinical acceptance depends on making those streams visible and interpretable to users [33]. A SHAP-style explanation could translate complex feature interactions into a concise risk decomposition that indicates why the model elevated the alert. This explanation should distinguish mechanistic plausibility from real-world reporting evidence so that clinicians can judge whether the alert is actionable.

Embedding Explanations into Clinical Workflow

Explanations should be embedded into the prescribing or pharmacy review workflow as a concise “why this alert?” summary rather than a separate technical report. Self-supervised molecular representation learning can support predictive discovery, but

bedside implementation requires translating learned signals into clinically meaningful categories [34]. The interface should show the suspected mechanism, the seriousness of the potential outcome, and the evidence source supporting the alert. This design would help clinicians decide whether to avoid the combination, adjust therapy, increase monitoring, or document a justified override. As shown in **Figure 3**, DDI explanations should be integrated directly into prescribing or pharmacy review workflows as a concise “why this alert?” summary that presents the suspected mechanism, the seriousness of the potential outcome, and the supporting evidence source, thereby enabling clinicians to choose whether to avoid the combination, adjust therapy, increase monitoring, or document a justified override.

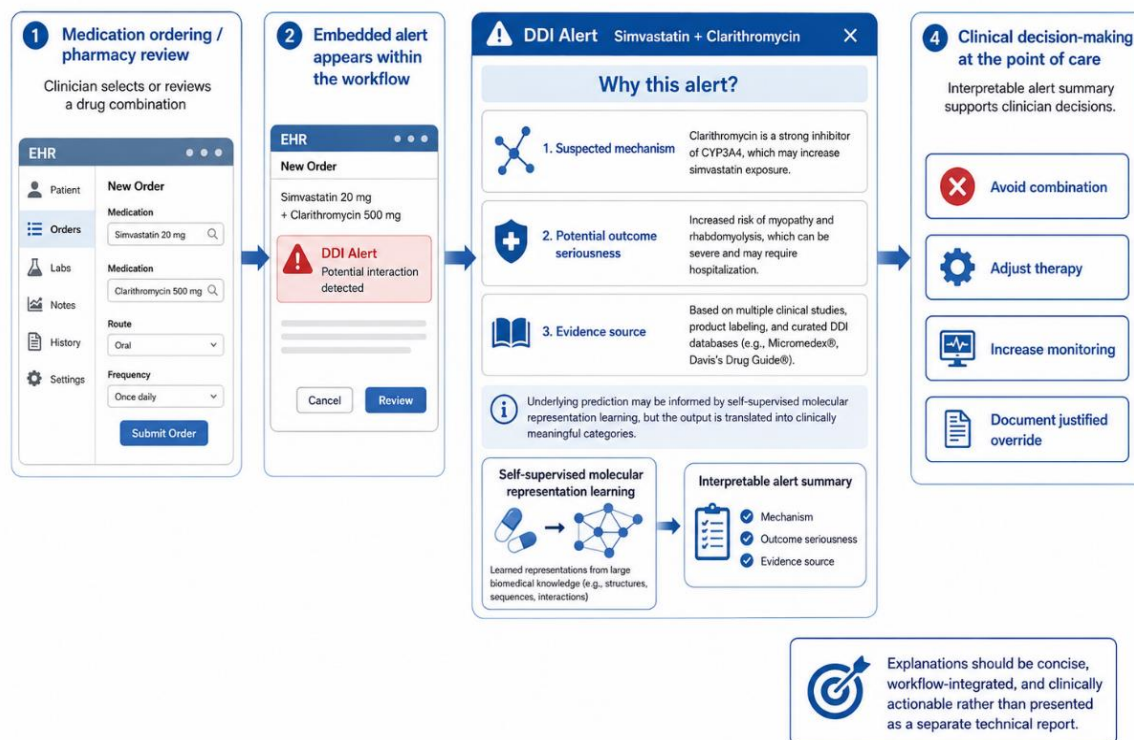


Figure 3. Workflow-integrated explainable DDI alert design for clinical decision support

Integration Into Clinical Decision Support

Prioritising Alerts in CPOE and Pharmacy Review

The model's severity probability could be used to tier interaction alerts within computerized provider order entry and pharmacy verification systems. High-concern interactions would be presented as interruptive alerts when the predicted harm is clinically serious and the evidence is mechanistically plausible, while lower-concern interactions would be displayed less disruptively. Multimodal and graph-based DDI prediction methods support the idea that heterogeneous evidence can distinguish interaction types rather than merely identify interaction presence [14, 20]. This prioritization would be intended to reduce alert fatigue by reserving workflow interruption for interactions that most require immediate clinical attention.

Continuous Learning from New Adverse Event Reports

The model should be updated periodically as spontaneous reports, drug labels, and real-world medication exposure patterns change. Timing-aware FAERS-based prioritization demonstrates that the temporal relationship between co-exposure and reported harm can add useful context when reassessing interaction signals over time [23]. Continuous learning should include monitoring for feature drift, reporting shifts, and emerging pharmacokinetic mechanisms in newly marketed drugs. Model updates would need governance review before deployment so that new signals improve safety without introducing unstable or poorly validated alerts.

Evaluation Strategy

Predictive Performance Metrics

Evaluation should examine discrimination, precision-recall behavior for severe interactions, and calibration without relying on a single summary statistic. A clinically useful model should also be compared conceptually with simpler baselines, including static DDI compendia, pharmacovigilance-only models, and mechanism-only models. Comprehensive reviews of DDI computational resources emphasize that model evaluation should consider both prediction quality and alignment with

established interaction knowledge bases [28]. Sensitivity analyses should focus on clinically acceptable alert burden and the trade-off between missed severe interactions and excessive warnings.

Temporal and External Validation

Temporal validation would train the model on earlier evidence and evaluate it against later interaction signals, helping assess whether predictions remain useful as prescribing patterns and reporting behavior change. External validation across healthcare systems or pharmacovigilance sources would test whether learned relationships generalize beyond a single reporting environment. Clinical discovery using spontaneous reporting systems and electronic health records supports the value of combining independent real-world data streams when assessing suspected DDIs [24]. The model should also be compared across settings with different patient complexity profiles to identify where recalibration may be needed.

Clinical Utility Assessment

Clinical utility should be evaluated by simulating how the model would change alert prioritization, pharmacist review workload, and recognition of previously under-emphasized severe interactions. A reference set of clinically relevant adverse DDIs can support expert review of whether model-prioritized alerts correspond to plausible and actionable safety concerns [12]. The assessment should examine whether explanations improve clinician confidence and whether lower-priority alerts can be safely moved to non-interruptive displays. Because the model is intended as a decision-support tool, evaluation should emphasize workflow fit, interpretability, and safe triage rather than claiming autonomous diagnostic authority.

Limitations

Reporting Bias and Incompleteness of Spontaneous Reports

Spontaneous adverse event reports are incomplete, selectively reported, and influenced by publicity, drug age, reporter behavior, and clinical suspicion. A model using FAERS or Vigibase evidence must therefore avoid treating absence of reports as absence of risk, particularly for newer drugs or rarely used combinations. Pharmacovigilance-derived DDI discovery can identify useful signals, but those signals require careful interpretation because reporting databases were not designed as controlled epidemiologic cohorts [24]. The proposed model should explicitly flag data-sparse predictions as uncertain and route them for cautious clinical review rather than definitive action.

Causality Attribution and Real-World Confounding

A predicted severe interaction risk does not establish causality because observed associations may reflect confounding by disease severity, co-medication patterns, or reporting artifacts. Even when pharmacokinetic mechanisms are plausible, the model's output should be interpreted as a risk signal that requires clinical reasoning and, where appropriate, confirmatory evaluation. Pharmacodynamic and pharmacokinetic interaction frameworks show that biological plausibility can strengthen interpretation, but plausibility alone cannot prove that a specific outcome was caused by the interaction [9, 10]. The safest role for the model is therefore to prioritize review, monitoring, or alternative therapy decisions rather than to make final clinical determinations.

Conclusion

A machine learning model for predicting severe drug–drug interactions could integrate polypharmacy burden, pharmacokinetic pathway overlap, and spontaneous adverse event evidence into a unified risk assessment. Such a model would move beyond static pairwise tables by considering the patient's full medication context. Its output would be a severity-oriented estimate intended to support clinical review rather than replace professional judgment.

The main strength of this approach is its holistic integration of mechanistic and real-world evidence. Pharmacokinetic pathway features would provide biological plausibility, polypharmacy features would represent regimen complexity, and adverse event signals would reflect harms observed in routine care. Interpretable outputs could help clinicians understand why an alert was generated and whether it deserves immediate action.

Important challenges remain, including incomplete reporting, missing pathway annotations, confounding by indication, and differences between healthcare systems. The model would require careful validation, transparent governance, and prospective clinical evaluation before routine deployment. Regulatory and institutional acceptance would also depend on demonstrating that the tool improves prioritization without introducing unsafe automation bias.

Future work should focus on collaborative benchmark datasets linking pharmacovigilance reports, pharmacokinetic pathway knowledge, medication exposure histories, and clinically adjudicated severity labels. Shared standards would make it easier to compare models, refine feature definitions, and evaluate clinical utility across settings. Pilot implementation in clinical decision support systems could then test whether severity-calibrated DDI prediction improves medication safety while reducing unnecessary alert burden.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Zhang W, Chen Y, Liu F, Luo F, Tian G, Li X. Predicting potential drug-drug interactions by integrating chemical, biological, phenotypic and network data. *BMC Bioinformatics*. 2017;18(1):18.
2. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*. 2018;34(13):i457-66.
3. Takeda T, Hao M, Cheng T, Bryant SH, Wang Y. Predicting drug-drug interactions through drug structural similarities and interaction networks incorporating pharmacokinetics and pharmacodynamics knowledge. *J Cheminform*. 2017;9(1):16.
4. Liu R, AbdulHameed MD, Kumar K, Yu X, Wallqvist A, Reifman J. Data-driven prediction of adverse drug reactions induced by drug-drug interactions. *BMC Pharmacol Toxicol*. 2017;18(1):44.
5. Qian S, Liang S, Yu H. Leveraging genetic interactions for adverse drug-drug interaction prediction. *PLoS Comput Biol*. 2019;15(5):e1007068.
6. 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):E4304-11.
7. Yu H, Mao KT, Shi JY, Huang H, Chen Z, Dong K, et al. Predicting and understanding comprehensive drug-drug interactions via semi-nonnegative matrix factorization. *BMC Syst Biol*. 2018;12(Suppl 1):14.
8. Zhang W, Chen Y, Li D, Yue X. Manifold regularized matrix factorization for drug-drug interaction prediction. *J Biomed Inform*. 2018;88:90-7.
9. Deodhar M, Al Rihani SB, Arwood MJ, Darakjian L, Dow P, Turgeon J, et al. Mechanisms of CYP450 inhibition: understanding drug-drug interactions due to mechanism-based inhibition in clinical practice. *Pharmaceutics*. 2020;12(9):846.
10. Niu J, Straubinger RM, Mager DE. Pharmacodynamic drug-drug interactions. *Clin Pharmacol Ther*. 2019;105(6):1395-406.
11. Kastrin A, Ferk P, Leskošek B. Predicting potential drug-drug interactions on topological and semantic similarity features using statistical learning. *PLoS One*. 2018;13(5):e0196865.
12. Kontsioti E, Maskell S, Dutta B, Pirmohamed M. A reference set of clinically relevant adverse drug-drug interactions. *Sci Data*. 2022;9(1):72.
13. Rohani N, Eslahchi C. Drug-drug interaction predicting by neural network using integrated similarity. *Sci Rep*. 2019;9(1):13645.
14. Deng Y, Xu X, Qiu Y, Xia J, Zhang W, Liu S. A multimodal deep learning framework for predicting drug-drug interaction events. *Bioinformatics*. 2020;36(15):4316-22.
15. Zhou D, Miao L, He Y. Position-aware deep multi-task learning for drug-drug interaction extraction. *Artif Intell Med*. 2018;87:1-8.
16. Gao J, Wu Z, Al-Sabri R, Oloulade BM, Chen J. AutoDDI: drug-drug interaction prediction with automated graph neural network. *IEEE J Biomed Health Inform*. 2024;28(3):1773-84.
17. Chen Y, Ma T, Yang X, Wang J, Song B, Zeng X. MUFFIN: multi-scale feature fusion for drug-drug interaction prediction. *Bioinformatics*. 2021;37(17):2651-8.
18. Nyamabo AK, Yu H, Shi JY. SSI-DDI: substructure-substructure interactions for drug-drug interaction prediction. *Brief Bioinform*. 2021;22(6):bbab133.
19. Abbas K, Hao C, Yong X, Hasan MK, Islam S, Rahman AH. Graph neural network-based drug-drug interaction prediction. *Sci Rep*. 2025;15(1):30340.
20. Hong Y, Luo P, Jin S, Liu X. LaGAT: link-aware graph attention network for drug-drug interaction prediction. *Bioinformatics*. 2022;38(24):5406-12.
21. Li Z, Tu X, Chen Y, Lin W. HetDDI: a pre-trained heterogeneous graph neural network model for drug-drug interaction prediction. *Brief Bioinform*. 2023;24(6):bbad385.
22. Zhu J, Liu Y, Zhang Y, Li D. Attribute supervised probabilistic dependent matrix tri-factorization model for the prediction of adverse drug-drug interaction. *IEEE J Biomed Health Inform*. 2021;25(7):2820-32.
23. Battini V, Cocco M, Barbieri MA, Powell G, Carnovale C, Clementi E, et al. Timing matters: A machine learning method for the prioritization of drug-drug interactions through signal detection in the FDA adverse event reporting system and their relationship with time of co-exposure. *Drug Saf*. 2024;47(9):895-907.
24. Jeong E, Su Y, Li L, Chen Y. Discovering clinical drug-drug interactions with known pharmacokinetics mechanisms using spontaneous reporting systems and electronic health records. *J Biomed Inform*. 2024;153:104639.
25. Vo TH, Nguyen NT, Le NQ. Improved prediction of drug-drug interactions using ensemble deep neural networks. *Med Drug Discov*. 2023;17:100149.

26. Lee G, Park C, Ahn J. Novel deep learning model for more accurate prediction of drug-drug interaction effects. *BMC Bioinformatics*. 2019;20(1):415.
27. Ibrahim H, El Kerdawy AM, Abdo A, Eldin AS. Similarity-based machine learning framework for predicting safety signals of adverse drug-drug interactions. *Inform Med Unlocked*. 2021;26:100699.
28. Zhao Y, Yin J, Zhang L, Zhang Y, Chen X. Drug-drug interaction prediction: databases, web servers and computational models. *Brief Bioinform*. 2024;25(1):bbad445.
29. Dmitriev AV, Filimonov DA, Rudik AV, Pogodin PV, Karasev DA, Lagunin AA, et al. Drug-drug interaction prediction using PASS. *SAR QSAR Environ Res*. 2019;30(9):655-64.
30. Abdul Raheem AK, Dhannoon BN. A novel deep learning model for drug-drug interactions. *Curr Comput Aided Drug Des*. 2024;20(5):666-72.
31. Wang F, Lei X, Liao B, Wu FX. Predicting drug-drug interactions by graph convolutional network with multi-kernel. *Brief Bioinform*. 2022;23(1):bbab511.
32. Nyamabo AK, Yu H, Liu Z, Shi JY. Drug-drug interaction prediction with learnable size-adaptive molecular substructures. *Brief Bioinform*. 2022;23(1):bbab441.
33. Gan Y, Liu W, Xu G, Yan C, Zou G. DMFDDI: deep multimodal fusion for drug-drug interaction prediction. *Brief Bioinform*. 2023;24(6):bbad397.
34. Kpanou R, Dallaire P, Rousseau E, Corbeil J. Learning self-supervised molecular representations for drug-drug interaction prediction. *BMC Bioinformatics*. 2024;25(1):47.