DrugInsight: Multi-Source Evidence Fusion for Explainable Drug-Drug Interaction Prediction Using Graph Neural Networks

Ayman Uzayr
Department of Computer Science and Engineering
Methodist College of Engineering and Technology
Hyderabad, India
160722733086@methodist.edu.in

Motassim Khan
Department of Computer Science and Engineering
Methodist College of Engineering and Technology
Hyderabad, India
160722733098@methodist.edu.in

Mustafa Meraj
Department of Computer Science and Engineering
Methodist College of Engineering and Technology
Hyderabad, India
160722733076@methodist.edu.in

Ms. R. Prathyusha
Department of Computer Science and Engineering
Methodist College of Engineering and Technology
Hyderabad, India
prathyusha@methodist.edu.in

***Abstract*—DrugInsight is an end-to-end, interpretable drug-drug interaction (DDI) prediction system. It evaluates interaction probability and severity while generating structured clinical explanations. The system utilizes RDKit and PyTorch Geometric to process drug SMILES strings into molecular graphs. These graphs are encoded by an AttentiveFP Graph Neural Network (GNN) into fixed-length structural embeddings. The embeddings are concatenated with a comprehensive 12-dimensional pharmacological feature vector. This vector encompasses shared biological entities from DrugBank and post-market safety signals (PRR) from the TWOSIDES database. A Multi-Layer Perceptron (MLP) binary classifier processes this combined representation. The final interaction probability is computed using a tiered fusion model. This model adaptively weights curated DrugBank rules, ML classifier scores, and TWOSIDES pharmacovigilance signals to generate a unified risk index and severity classification. The system achieves a validation AUC of 0.7065 under a rigorous cold-start drug-level split. This demonstrates robust generalization to entirely unseen compounds. DrugInsight ensures high transparency through a rule-based explainer detailing metabolic competition, pharmacodynamic overlap, and explicit clinical recommendations. The system is packaged for versatile deployment via a Streamlit web interface, a FastAPI REST endpoint, and a command-line interface (CLI).**

***Index Terms*—Drug-drug interactions, graph neural networks, AttentiveFP, evidence fusion, explainability, pharmacovigilance, DrugBank, TWOSIDES**

I. Introduction

Drug-drug interactions (DDIs) constitute a significant source of preventable adverse drug events (ADEs). They contribute to an estimated 74,000 emergency department visits and 195,000 hospitalizations annually in the United States alone [1]. Polypharmacy is increasingly prevalent, particularly among elderly and chronically ill patients. As a result, the combinatorial space of potential DDIs far exceeds the capacity of manual curation efforts. Commercial drug interaction databases such as DrugBank [2], Medscape, and Lexicomp provide robust curated interaction records. However, they are inherently limited by the scope of manually reviewed drug pairs and cannot generalize to newly approved compounds.

Recent advances in deep learning, particularly graph neural networks (GNNs), show promise in learning molecular representations directly from chemical structures [3], [4]. Unfortunately, most GNN-based systems operate as black-box classifiers. They provide binary predictions without clinically interpretable explanations, thereby limiting their utility in clinical decision support. Furthermore, purely structure-based approaches fail to incorporate critical pharmacological context. Clinicians rely heavily on shared metabolizing enzymes, pharmacological target overlaps, and post-market safety signals when evaluating interaction risks.

In this work, we present **DrugInsight**, an end-to-end DDI prediction system addressing these domain limitations through three key innovations:

1) **Multi-source evidence fusion**: DrugInsight combines GNN-derived molecular predictions with curated DrugBank pharmacological evidence and real-world TWOSIDES [9] pharmacovigilance signals. This is implemented through an adaptive tiered fusion layer. This approach mitigates gaps in any single data source and achieves broader coverage than individual database-based DDI checkers.
2) **Structured explainability**: The system generates mechanism-grounded clinical explanations detailing metabolic competition, pharmacodynamic overlap, and patient safety signals. It leverages a rule-based explainer built on CYP enzyme interaction knowledge and shared targets. This provides clinical transparency without the computational overhead or hallucination risks of large

language models.

3) **Cold-start generalization**: The model is rigorously evaluated under a zero-overlap drug-level split protocol. Validation pairs consist entirely of drugs unseen during training. This simulates the realistic scenario of evaluating newly approved or rarely studied compounds.

The remainder of this paper is organized as follows: Section II reviews related work, Section III details the system architecture and methodology, Section IV describes the experimental setup, Section V presents results and case studies, Section VI discusses limitations and future directions, and Section VII concludes the paper.

II. Related Work

A. GNN-Based DDI Prediction

Graph neural networks have emerged as a powerful paradigm for molecular property prediction. DeepDDI [3] pioneered the use of structural similarity profiles for DDI prediction but did not directly learn from molecular graphs. KGNN [4] incorporated knowledge graph embeddings alongside molecular features, while SSI-DDI [5] employed substructure-level interaction modeling. MR-GNN [6] introduced multi-resolution graph representations, and MHCADDI [7] proposed multi-head cross-attention mechanisms for capturing inter-molecular dependencies. AttentiveFP [8], which DrugInsight employs, uses graph attention with virtual super-nodes and multi-timestep readout for molecular graph encoding.

B. Knowledge-Based and Database Approaches

DrugBank [2] remains the gold-standard curated drug interaction database, providing mechanism-level annotations for known interactions. Commercial tools such as Medscape and Lexicomp aggregate similar curated data. While these resources provide high-precision interaction records, their coverage is inherently limited to manually reviewed pairs and cannot scale to the combinatorial explosion of potential drug combinations, particularly for newly approved compounds.

C. Pharmacovigilance Signal Detection

Post-market surveillance databases such as FAERS and TWOSIDES [9] offer real-world evidence of drug co-administration risks. The Proportional Reporting Ratio (PRR) is a standard disproportionality measure that quantifies the excess adverse event reporting for drug combinations relative to individual drugs. Prior work has used these signals primarily for retrospective analysis; DrugInsight integrates them as a prospective evidence source within the prediction pipeline.

D. Explainability in DDI Prediction

Explainability remains a critical gap in most DDI prediction systems. Attention-based methods [8] provide implicit feature importance scores but do not generate human-readable clinical explanations. Knowledge-grounded approaches [10] have explored ontology-based reasoning but typically require extensive manual knowledge engineering. DrugInsight's approach combines curated CYP enzyme interaction data, DrugBank mechanism annotations, and pharmacovigilance signals into a structured explanation framework without relying on a large language model (LLM).

III. Methodology

A. System Architecture Overview

Fig. 1 presents the end-to-end DrugInsight pipeline. Given two drug identifiers, the system: (1) resolves drug names against the DrugBank database, (2) converts SMILES representations into molecular graphs, (3) encodes each graph via an AttentiveFP GNN, (4) extracts pharmacological pair features, (5) classifies interaction probability via an MLP, (6) fuses the ML prediction with curated evidence and pharmacovigilance signals, and (7) generates a structured clinical explanation.

B. Data Sources and Preprocessing

1) DrugBank Data Extraction: DrugBank [2] exports are parsed into structured CSV tables: drug metadata, interaction pairs, enzymes, targets, transporters, carriers, pathways, and SMILES structures. Drugs are filtered to retain only those with valid, parseable SMILES strings (verified using RDKit `MolFromSmiles`), yielding approximately 3,803 drugs with molecular structure data.

2) TWOSIDES Pharmacovigilance Data: The TWOSIDES database [9] provides post-market adverse event reporting data with Proportional Reporting Ratios (PRR) for drug combinations. A RxNorm bridge file maps DrugBank IDs to RxNorm identifiers, enabling cross-database linkage.

3) Interaction Enrichment: Each interaction pair is enriched with computed pharmacological counts (shared enzymes, targets, transporters, carriers, pathways) and the maximum TWOSIDES PRR via bidirectional RxNorm join. PRR values are capped at the 99th percentile to control outliers, and duplicate pairs are deduplicated by retaining the highest-PRR record. The enriched dataset contains approximately 936,178 interaction pairs.

C. Molecular Graph Representation

Each drug's SMILES string is converted into a PyTorch Geometric `Data` object using RDKit. Hydrogen atoms are explicitly added before feature extraction. The molecular graph $G = (V, E)$ consists of atoms as nodes and chemical bonds as undirected edges.

Atom Features (8-dimensional): Table I details the per-node feature vector.

Bond Features (6-dimensional): Each bond is encoded as a 6-dimensional vector: single, double, triple, and aromatic bond type (one-hot), conjugation flag, and ring membership flag. All bonds are represented as bidirectional edges.

D. GNN Encoder: AttentiveFP

The molecular graph encoder employs AttentiveFP [8], a graph attention network architecture with virtual super-node readout that has demonstrated strong performance on molecular property prediction tasks. Table II shows the encoder configuration.

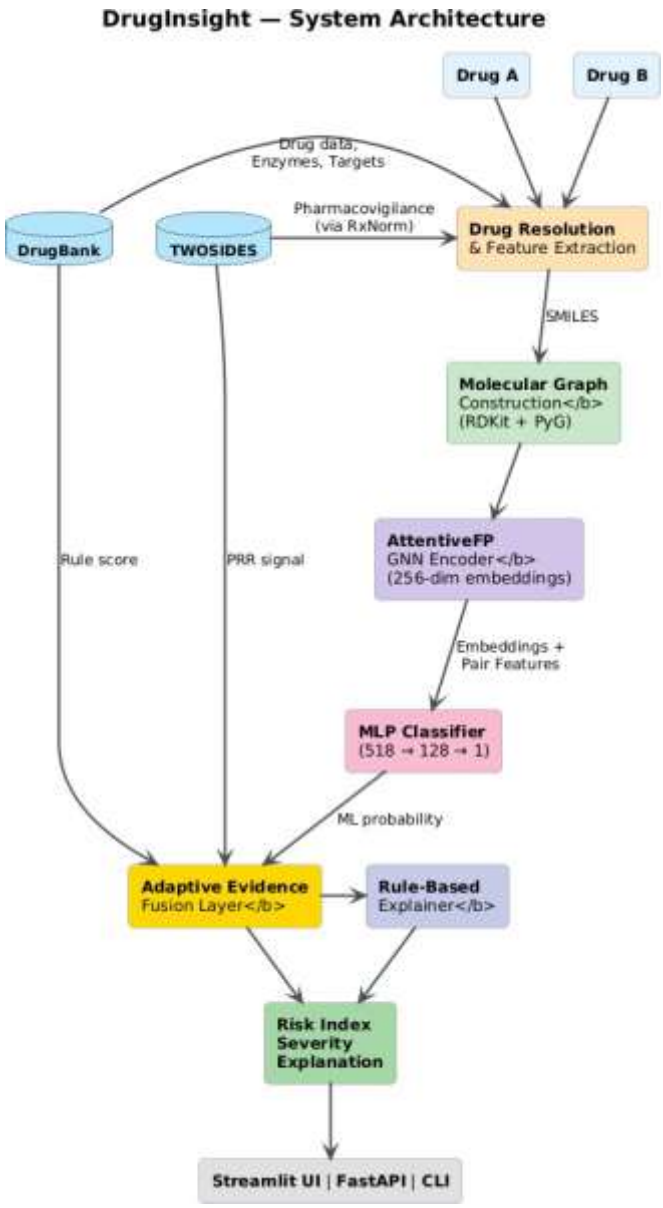

Fig. 1. End-to-end DrugInsight pipeline. Two drug identifiers are resolved, encoded as molecular graphs by AttentiveFP, classified by the MLP, and the prediction is fused with DrugBank rules and TWOSIDES signals before explanation generation.

TABLE I
ATOM-LEVEL NODE FEATURES (8-DIMENSIONAL)

Feature	Type	Description
Atomic number	int	Element identity
Degree	int	Number of bonds
Formal charge	int	Net charge
Hybridization	int	sp/sp2/sp3 as integer
Aromaticity	binary	In aromatic ring
H-count	int	Total attached hydrogens
Ring membership	binary	Part of any ring
Normalized mass	float	Atomic mass / 100

TABLE II
ATTENTIVEFP GNN ENCODER CONFIGURATION

Parameter	Value
in_channels	8
edge_dim	6
hidden_channels	128
out_channels	256
num_layers	4
num_timesteps	2
dropout	0.3

The GNN produces a 256-dimensional global graph embedding for each drug molecule. The attention mechanism allows the network to learn atom-level importance for molecular property prediction.

E. DDI Classifier: Multi-Layer Perceptron

The two 256-dimensional drug embeddings are concatenated with a 12-dimensional pharmacological feature vector (described in Section III-F), yielding a 524-dimensional pair representation. This is processed by a 3-layer MLP trunk with batch normalization, ReLU activation, and dropout regularization:

$$\mathbf{h} = \text{MLP}([\mathbf{e}_A \,|\, \mathbf{e}_B \,|\, \mathbf{f}_{pair}]) \quad (1)$$

where $\mathbf{e}_A, \mathbf{e}_B \in \mathbb{R}^{256}$ are drug embeddings and $\mathbf{f}_{pair} \in \mathbb{R}^{12}$ is the pharmacological feature vector. Table III details the layer-wise configuration of the trunk architecture.

TABLE III
MLP TRUNK ARCHITECTURE

Layer	Components	Input	Output
1	Linear → BatchNorm → ReLU → Dropout(0.5)	524	512
2	Linear → BatchNorm → ReLU → Dropout(0.5)	512	256
3	Linear → BatchNorm → ReLU → Dropout(0.5)	256	128

A **probability head** (Linear 128→1) produces an interaction logit passed through sigmoid. A **severity head** (Linear 128→3) for Minor/Moderate/Major classification is architecturally defined but reserved for future training when severity labels become available.

F. Pair-Level Pharmacological Features

To complement learned molecular representations with curated biological context, a 12-dimensional pharmacological feature vector is computed for each drug pair, as detailed in Table IV.

Normalization caps are set from empirical 99th-percentile data distribution analysis to prevent dominant features from overwhelming the model.

G. Adaptive Evidence Fusion Layer

A key differentiator of DrugInsight is its tiered evidence fusion layer, which combines three independent scoring sources under clinical safety heuristics. Unlike single-database

TABLE IV
PAIR-LEVEL PHARMACOLOGICAL FEATURES (12-DIMENSIONAL)

Feature	Source	Norm. Cap	Rationale
Shared enzyme count	DrugBank	/5	Metabolic competition
Shared target count	DrugBank	/2	PD synergy/antagonism
Shared transporter count	DrugBank	/2	Distribution interactions
Shared carrier count	DrugBank	/1	Distribution interactions
Shared pathway count	DrugBank	/1	Downstream cascades
Shared major CYPs	DrugBank	/3	Dominant metabolic routes
CYP3A4 shared	DrugBank	binary	Specific primary enzyme
CYP2D6 shared	DrugBank	binary	Specific primary enzyme
CYP2C9 shared	DrugBank	binary	Specific primary enzyme
Max PRR	TWOSIDES (via RxNorm bridge)	/245.05	Peak adverse safety signal
TWOSIDES signals count	TWOSIDES	/1882	Breadth of safety signals
TWOSIDES found flag	TWOSIDES	binary	Signal availability

TABLE V
DATASET STATISTICS

Source	Content	Scale
DrugBank (enriched)	Interaction pairs	~936,178
DrugBank (filtered)	Drugs with valid SMILES	~3,803
DrugBank enzymes	Metabolizing enzyme records	~760 KB
DrugBank targets	Pharmacological targets	~2.7 MB
DrugBank pathways	Biological pathways	~23 MB
TWOSIDES (filtered)	Pharmacovigilance signals	~87 MB
RxNorm bridge	DrugBank ↔ RxNorm mapping	Cross-link

DDI checkers, DrugInsight integrates complementary evidence streams to achieve broader coverage.

1) Evidence Sources: Three independent scores are computed:

- **Rule Score (DrugBank):** Additive scoring based on curated mechanisms and shared entities.
- **ML Score (GNN+MLP):** Sigmoid of classifier logit.
- **TWOSIDES Score:** Priority PRR / distinct signals count.

2) Adaptive Weighting Tiers: The fusion weights adapt based on algorithmic tiers determining the availability of specific evidence sets:

- **Tier 1 (Direct Hit):** $P_{fused} = 0.7 \times \text{rule} + 0.3 \times ML$ (if no ML, $1.0 \times$ rule).
- **Tier 2 (Evidence Fusion):** Logistic regression unification of signals (Rule + ML + TWOSIDES).
- **Tier 3 (ML Only):** $P_{fused} = 1.0 \times ML$.

This tiered system ensures graceful confidence degradation: when curated evidence exists, it safely constrains predictions. A risk index is derived as $fused\_prob \times 100$ (integer 0–100), and severity is classified as Major (≥ 70), Moderate (40–69), or Minor (< 40).

H. Rule-Based Explainability Module

DrugInsight generates structured clinical explanations without any LLM through a hierarchical rule-based reasoning pipeline:

1) **Metabolic mechanism**: Checks CYP inhibitor/inducer knowledge base → DrugBank enzyme actions → generic substrate competition.
2) **Pharmacodynamic mechanism**: Reports shared pharmacological targets.
3) **Pharmacovigilance note**: Translates PRR into textual signal strength (weak < 3, moderate 3–10, strong > 10).
4) **Priority**: If a curated DrugBank mechanism text exists for the pair, it is used as the primary explanation.
5) **Clinical recommendation**: Severity-dependent advice (avoid concurrent use / use with caution / standard monitoring).

IV. EXPERIMENTAL SETUP

A. Dataset

Table V summarizes the datasets and their roles in the system.

B. Training Configuration

Table VI summarizes all hyperparameters and training details.

TABLE VI
TRAINING HYPERPARAMETERS

Parameter	Value
Optimizer	Adam (two parameter groups)
GNN learning rate	3×10^{-5}
Classifier learning rate	1×10^{-4}
Weight decay	5×10^{-4}
LR scheduler	ReduceLROnPlateau (mode='max', factor=0.5, patience=3)
Batch size	64
Max epochs	20
Early stopping patience	6 epochs (based on val AUC)
Improvement threshold	1×10^{-4} (for AUC)
Loss function	BCEWithLogitsLoss
Dropout	0.5 (classifier), 0.3 (GNN)
Symmetry augmentation	50% random drug swap per batch
Gradient clipping	max_norm = 1.0
Seed	42 (split + neg sampling)
Device	CUDA if available, else CPU
Num workers	0
Train/Val split	80/20 drug-level
Negative ratio	1:1 (equal positives and negatives)
Hard neg fraction	70%
Label smoothing	Enabled
Weighted loss	Disabled

1) Drug-Level Cold-Start Split: All unique drugs are split 80/20 into training and validation sets at the *drug level* (prior to pairing). Interaction pairs are then partitioned to train or validation exclusively when *both* constituent drugs belong to the same respective subset. This evaluation protocol measures generalization against completely unseen drugs. This strict setting is significantly more challenging and realistic than standard pair-level splitting, where individual drugs may trivially leak into both evaluation splits.

2) Hard Negative Sampling: For each split, negative (non-interacting) pairs are generated with a 70/30 hard/easy split. Hard negatives are non-interacting pairs scored by a "plausibility" metric: a weighted sum of shared enzymes ($\times 2$), targets ($\times 2$), transporters ($\times 1$), carriers ($\times 1$), and pathways ($\times 0.5$). The highest-scoring non-interacting pairs are retained as hard negatives. This prevents the model from learning trivial shortcuts while avoiding training collapse through the 30% easy negative fraction.

3) Symmetry Augmentation: During training, drug embeddings are randomly swapped with 50% probability per batch to enforce order-invariant predictions, i.e., $P(A, B) = P(B, A)$.

C. Evaluation Protocol

- **Primary metric**: ROC AUC (area under receiver operating characteristic curve) for model selection and early stopping
- **Secondary metrics**: Average Precision (AP), Accuracy, Confusion Matrix
- **Threshold**: 0.5 on sigmoid probability for binary classification
- **Best model**: Selected by highest validation AUC; checkpoint saved

V. Results and Discussion

A. Model Training Performance

TABLE VII
Training Performance (Per-Epoch Metrics)

Epoch	Train Loss	Val AUC	Val AP	Val Acc
1	0.6556	0.6945	0.6918	0.6394
2	0.6282	0.7055	0.6986	**0.6489**
3	0.6200	0.7002	0.6970	0.6442
4	0.6156	**0.7065**	0.7022	0.6485
5	0.6134	0.7001	0.6986	0.6480
6	0.6120	0.7047	**0.7024**	0.6429
7	0.6113	0.7042	0.7021	0.6455
8	0.6112	0.7041	0.6967	0.6422
9	**0.6047**	0.7023	0.6987	0.6446
10	0.6055	0.6996	0.6965	0.6411
Best	0.6156	**0.7065**	0.7022	0.6485

The model achieves a best validation AUC of 0.7065 under the cold-start evaluation protocol. This protocol ensures zero drug overlap between training and validation splits. Both drugs in every validation pair remain entirely unseen during training. This strict setting is substantially more challenging than the random pair-level splitting used by most methods. Random splitting allows algorithms to simply memorize individual drug structural signatures. The cold-start protocol eliminates this shortcut to measure true generalization.

B. Case Studies

To demonstrate DrugInsight's practical capability and the advantage of multi-source evidence fusion over single-database checkers, we present four clinically motivated drug pair evaluations. The output is confirmed using existing drugbank interaction checking software and cited references.

*1) **Case 1: Probenecid and Benzylpenicillin**:*

DrugInsight Output:
Predicted Interaction: Moderate
Risk Index: 42/100
Shared Biology: Pathway: Organic anion transporter 3

Pharmacological Mechanism: The excretion of Benzylpenicillin can be decreased when combined with Probenecid. This combination may cause clinically significant adverse effects requiring monitoring.

Clinical Recommendation: Use Probenecid and Benzylpenicillin together with caution. Watch for signs of enhanced pharmacodynamic effects. Consider dose reduction if adverse effects occur. *Fusion Layer Component Scores*: Drugbank 0.280, ML Model 0.736, TwoSides 0.0

Fusion weights: Rule 0.7 · ML 0.3 · TWOSIDES 0.0

Drugbank Drug Interaction Checker Output:
Severity: Moderate [14]
Description: The excretion of Benzylpenicillin can be decreased when combined with Probenecid.

Extended Description: Probenecid is a uricosuric and renal tubular blocking agent that inhibits the tubular secretion of penicillin and usually increases penicillin plasma levels by any route the antibiotic is given. A 2-fold to 4-fold elevation has been demonstrated for various penicillins. Probenecid is indicated to be given as an adjuvant to therapy to penicillins to enhance their plasma levels 1, concomitant use of probenecid and penicillins has been associated with increased penicillin-related adverse reactions including psychic disturbances [15].

*2) **Case 2: Amoxicillin and Clavulanic Acid**:* This is an intentional synergistic combination where clavulanic acid inhibits beta-lactamase enzymes, protecting amoxicillin from degradation. The system should identify this as a known, well-characterized interaction with a clear mechanism.

DrugInsight Output:
Predicted Interaction: Moderate
Risk Index: 63/100
Shared Biology: no shared biology

Inferred from Twosides: High PRR detected. Potential confounding, this signal may reflect individual drug toxicity rather than a true interaction.

Pharmacological Mechanism: Post-market surveillance data (TWOSIDES) shows a strong adverse event signal for this combination (PRR=125.0), suggesting real-world co-prescribing risks. This combination may cause clinically significant adverse effects requiring monitoring.

Clinical Recommendation: Use Amoxicillin and Clavulanic acid together with caution. Consider

dose reduction if adverse effects occur.

Fusion Layer Component Scores: Rule (DrugBank) 0.000, ML Model 0.882, TWOSIDES 0.609

Fusion weights: Rule 2.90190 · ML 3.93608 · TWOSIDES 0.04742

Drugbank Drug Interaction Checker Output:
Result: No Interactions Found [14]

*3) **Case 3: Doxorubicin and Sildenafil**:* This less-documented combination tests DrugInsight's ability to detect potential interactions for drug pairs with limited curated evidence, demonstrating the value of the ML component and TWOSIDES pharmacovigilance signals when DrugBank coverage is sparse.

DrugInsight Output:
Predicted Interaction: Moderate
Risk Index: 58/100
Shared Biology: ENZYMES CYP2D6, CYP3A4

Pharmacological Mechanism: The serum concentration of Doxorubicin can be increased when it is combined with Sildenafil. This combination may cause clinically significant adverse effects requiring monitoring.

Clinical Recommendation: Use Doxorubicin and Sildenafil together with caution. Monitor drug levels and clinical response. Consider dose reduction if adverse effects occur.

Fusion Layer Component Scores: DrugBank 0.500, ML Model 0.764, TWOSIDES 0.0

Fusion weights: Rule 0.7 · ML 0.3 · TWOSIDES 0.0

Drugbank Drug Interaction Checker Output:
Severity: Moderate [14]
Description: The serum concentration of Doxorubicin can be increased when it is combined with Sildenafil.

Extended Description: P-glycoprotein inhibitors increase the nuclear translocation of doxorubicin and cellular cytotoxicity resulting from doxorubicin treatment [16].

*4) **Case 4: Ritonavir and Darunavir**:* This is a well-known CYP3A4-mediated pharmacokinetic boosting interaction, where ritonavir inhibits CYP3A4 to increase darunavir levels. This pair tests the CYP knowledge base integration in the explainer module.

DrugInsight Output:
Predicted Interaction: Major
Risk Index: 94/100
Shared Biology: ENZYMES CYP2D6, CYP3A4; TARGETS Gag-Pol polyprotein, Pol polyprotein

Pharmacological Mechanism: Ritonavir is a known inhibitor of CYP3A4, the primary enzyme responsible for Darunavir metabolism. This inhibition reduces Darunavir clearance, increasing its plasma concentration and risk of toxicity. Both drugs act on shared pharmacological targets: Gag-Pol polyprotein, Pol polyprotein. Concurrent use may produce additive or synergistic effects at these targets, increasing the risk of exaggerated pharmacodynamic responses. This combination may cause serious or life-threatening adverse effects.

Clinical Recommendation: Avoid concurrent use of Ritonavir and Darunavir if possible. If co-administration is necessary, closely monitor for adverse effects and consider dose adjustment. Consult clinical guidelines or a pharmacist.

Fusion Layer Component Scores: Rule (DrugBank) 0.730, ML Model 0.932, TWOSIDES 0.000

Fusion weights: Rule 2.9019077471083583 · ML 3.9360892822619045 · TWOSIDES 0.0474253927807187

Drugbank Drug Interaction Checker Output:
Severity: [–] [14]
Description: [–]

Extended Description: [–]

*5) **Case 5: Acetaminophen and Lisinopril**:* This case highlights a common pair of medications with a minor interaction, demonstrating the system's ability to discriminate severity.

DrugInsight Output:
Predicted Interaction: Minor
Risk Index: 19/100
Shared Biology: None

Pharmacological Mechanism: Acetaminophen may decrease the excretion rate of Lisinopril which could result in a higher serum level. This combination may cause minor adverse effects unlikely to require intervention.

Clinical Recommendation: The interaction between Acetaminophen and Lisinopril is generally manageable. Standard monitoring is recommended. No immediate dose adjustment required.

Fusion Layer Component Scores: Rule (DrugBank) 0.825, ML Model 0.621, TWOSIDES 0.0

Fusion weights: Rule 0.7 · ML 0.3 · TWOSIDES 0.0

Drugbank Drug Interaction Checker Output:
Severity: Minor [14]
Description: Acetaminophen may decrease the excretion rate of Lisinopril which could result in a higher serum level.

Extended Description: The renal excretion of drugs is the overall result of a combination of kidney processes that include glomerular filtration, passive diffusion, tubular secretion, and tubular reabsorption [17].

6) ***Case 6: Hyaluronic Acid and Betaine****:* This case evaluates a pair with no structural overlap or human annotations where the system appropriately predicts no significant interaction relying heavily on the ML component.

DrugInsight Output:
Predicted Interaction: No Significant Interaction
Risk Index: 45/100
Shared Biology: None

Pharmacological Mechanism: No shared metabolic enzymes or pharmacological targets were identified that would suggest a clinically significant interaction.

Clinical Recommendation: Standard prescribing guidelines apply. No special precautions required based on available data.

Fusion Layer Component Scores: Rule (DrugBank) 0.000, ML Model 0.452, TWOSIDES 0.000

Fusion weights: Rule 0.0 · ML 1.0 · TWOSIDES 0.0

Drugbank Drug Interaction Checker Output:
Result: No Interactions Found [14]

C. Multi-Source Fusion Advantage

The adaptive evidence fusion layer provides a fundamental advantage over single-source DDI checkers:

- **Coverage**: DrugBank alone covers only manually curated pairs. When a drug pair lacks DrugBank records, the fusion layer shifts weight to the ML model (70%) and TWOSIDES (20%), enabling predictions beyond the curated database boundary.
- **Robustness**: By combining three independent evidence streams, the system is resilient to gaps in any single source. A drug pair may lack a DrugBank record but show strong TWOSIDES pharmacovigilance signals, or vice versa.
- **Transparency**: Component scores are reported separately (rule score, ML score, TWOSIDES score), enabling clinicians to understand which evidence sources contribute to each prediction.
- **Confidence estimation**: The system reports per-source confidence (high/moderate/low/weak) and an overall confidence assessment, enabling informed clinical decision-making.

VI. Limitations and Future Work

A. Current Limitations

- **Low epoch training**: The current model is trained with a limited number of epochs. Extended training with computational resources may improve the baseline ML component performance.
- **Severity head untrained**: The severity classification head is architecturally defined (3-class: Minor/Moderate/Major) but not trained due to the absence of ground-truth severity labels. Severity is currently derived from fusion-based risk index thresholds.
- **Single split evaluation**: Results are based on a single 80/20 drug-level split (seed=42) without cross-validation. K-fold cross-validation would provide more robust performance estimation.
- **Hand-tuned fusion weights**: The adaptive fusion weights (0.60/0.40/0.10 etc.) are manually designed based on clinical reasoning rather than learned from data.
- **No formal ablation**: A systematic ablation study comparing ML-only, fusion-only, and rule-only configurations is not yet conducted.

B. Future Directions

- **Training the severity head**: Incorporating interaction severity labels (e.g., from curated clinical databases) would enable end-to-end severity prediction.
- **Cross-attention variant**: A cross-attention model checkpoint exists in the repository, suggesting an alternative architecture for capturing inter-molecular dependencies which warrants systematic evaluation.
- **Learned fusion weights**: Replacing hand-tuned weights with learnable parameters or a meta-learning approach could improve evidence integration.
- **Extended training**: Utilizing more computational resources for additional training epochs and hyperparameter optimization.
- **Multi-class severity prediction**: Extending the binary classifier to multi-class DDI type prediction.

VII. Conclusion

This paper presents DrugInsight, a multi-source evidence fusion system for explainable drug-drug interaction prediction. DrugInsight combines AttentiveFP graph neural network molecular encodings with curated DrugBank evidence and real-world TWOSIDES pharmacovigilance signals. An adaptive weighted fusion layer dynamically integrates these independent scoring sources. The system demonstrates that such

hybrid approaches provide broader predictive coverage and far greater transparency than single-source checkers. Furthermore, an explicit rule-based explainer generates mechanism-grounded clinical explanations detailing metabolic competition, pharmacodynamic overlap, and safety signal strengths. This rule-based framework completely circumvents the hallucination risks frequently associated with large language models. The system successfully achieves a validation AUC of 0.7065 under a strict cold-start drug-level evaluation protocol. DrugInsight provides per-source confidence estimation alongside actionable clinical recommendations. It is packaged with an interactive Streamlit web interface, a FastAPI REST endpoint, and a command-line tool. This robust architecture addresses practical deployment needs for automated DDI screening in research and clinical environments.

REFERENCES

[1] J. Lazarou, B. H. Pomeranz, and P. N. Corey, "Incidence of adverse drug reactions in hospitalized patients: A meta-analysis of prospective studies," *Journal of the American Medical Association*, vol. 279, no. 15, pp. 1200–1205, 1998.

[2] D. S. Wishart *et al.*, "DrugBank 5.0: A major update to the DrugBank database for 2018," *Nucleic Acids Research*, vol. 46, no. D1, pp. D1074–D1082, 2018.

[3] J.-Y. Ryu, H. U. Kim, and S. Y. Lee, "Deep learning improves prediction of drug-drug and drug-food interactions," *Proceedings of the National Academy of Sciences*, vol. 115, no. 18, pp. E4304–E4311, 2018.

[4] X. Lin, Z. Quan, Z.-J. Wang, T. Ma, and X. Zeng, "KGNN: Knowledge graph neural network for drug-drug interaction prediction," in *Proc. IJCAI*, 2020, pp. 2739–2745.

[5] A. Nyamabo, Y. Yu, and J. Shi, "SSI-DDI: Substructure-substructure interactions for drug-drug interaction prediction," *Briefings in Bioinformatics*, vol. 22, no. 6, 2021.

[6] N. Xu, P. Wang, L. Chen, J. Tao, and J. Zhao, "MR-GNN: Multi-resolution and dual graph neural network for predicting structured entity interactions," in *Proc. IJCAI*, 2019, pp. 3968–3974.

[7] D. Deac, Y.-H. Huang, P. Velic˘kovic´, P. Lio`, and J. Tang, "Drug-drug adverse effect prediction with graph co-attention," in *Proc. ICML Workshop*, 2019.

[8] Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang, and M. Zheng, "Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism," *Journal of Medicinal Chemistry*, vol. 63, no. 16, pp. 8749–8760, 2020.

[9] N. P. Tatonetti, P. P. Ye, R. Daneshjou, and R. B. Altman, "Data-driven prediction of drug effects and interactions," *Science Translational Medicine*, vol. 4, no. 125, pp. 125ra31, 2012.

[10] R. C¸ elebi, H. Uyar, E. Yasar, O. Gu¨mus, O. Dikenelli, and M. D. Dumontier, "Evaluation of knowledge graph embedding approaches for drug-drug interaction prediction in realistic settings," *BMC Bioinformatics*, vol. 20, no. 1, pp. 1–14, 2019.

[11] G. Landrum, "RDKit: Open-source cheminformatics software," 2023. [Online]. Available: https://www.rdkit.org/

[12] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

[13] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[14] DrugBank, "Drug Interaction Checker." [Online]. Available: https://go.drugbank.com/drug-interaction-checker

[15] U.S. National Library of Medicine, "Probenecid Tablets, USP - DailyMed." [Online]. Available: https://dailymed.nlm.nih.gov/dailymed/fda/fdaDrugXsl.cfm?setid=14637e17-4265-4195-8182-7ca13796c684

[16] L. Bao, A. Haque, K. Jackson, S. Hazari, K. Moroz, R. Jetly, and S. Dash, "Increased expression of P-glycoprotein is associated with doxorubicin chemoresistance in the metastatic 4T1 breast cancer model," *American Journal of Pathology*, vol. 178, no. 2, pp. 838–852, 2011.

[17] C. A. van Ginneken and F. G. Russel, "Saturable pharmacokinetics in the renal excretion of drugs," *Clinical Pharmacokinetics*, vol. 16, no. 1, pp. 38–54, Jan. 1989. doi: 10.2165/00003088-198916010-00003. (PubMed ID 2650954).