IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

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 predicts drug-drug interactions (DDIs), returning an interaction probability, a severity tier, and a structured clinical explanation for any queried drug pair. The system combines AttentiveFP graph neural network encodings of molecular structure with curated pharmacological data from DrugBank and post-market adverse event signals from TWOSIDES. RDKit and PyTorch Geometric convert SMILES strings into molecular graphs; the GNN encodes each graph into a 256-dimensional structural embedding. These embeddings are concatenated with a 12-dimensional feature vector derived from shared biological targets, metabolizing enzymes, transporters, and carriers (DrugBank), and the Proportional Reporting Ratio from TWOSIDES. An MLP classifier processes the combined vector. A tiered fusion layer then weights the DrugBank rule score, MLP output, and TWOSIDES signal according to what evidence is available for the queried pair, producing a 0–100 risk index and a severity label.**

We evaluate the model under a drug-level cold-start split: neither drug in any validation pair appears during training. This tests whether the system can handle newly approved or understudied compounds rather than simply interpolating over seen chemical space. The best checkpoint reaches a validation AUC of 0.7065. Clinical explanations are generated by a rule-based module covering CYP enzyme-mediated metabolic competition, shared pharmacological targets, and pharmacovigilance signal strength; no language model is involved, so explanation outputs carry no hallucination risk. Per-source confidence scores accompany every prediction. The system is accessible via a Streamlit web interface, a FastAPI endpoint, and a command-line interface.

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

I. Introduction

Adverse drug events attributable to drug-drug interactions (DDIs) are a persistent problem in polypharmacy management. In the United States alone, DDIs account for an estimated 74,000 emergency department visits and 195,000 hospitalizations annually [1]. The patient populations most exposed—elderly individuals and those managing multiple chronic conditions—are also the least able to tolerate the consequences. Manually curated databases such as DrugBank [2], Medscape, and Lexicomp record interactions for known drug pairs with high precision, but their coverage thins out for recently approved compounds and off-label combinations that have not yet accumulated enough clinical evidence to be reviewed.

Graph neural networks offer one plausible route around this coverage problem. By encoding a drug's molecular graph directly, a GNN can generate structural representations for any compound with a valid SMILES string, including those that appear after a database's last update [3], [4]. The practical limitation is that most GNN-based DDI systems stop at a binary prediction. They do not explain *why* an interaction is expected, which makes adoption in clinical decision support difficult. Pharmacists and prescribers need to understand the mechanism—whether the risk comes from shared CYP3A4 metabolism, overlapping pharmacological targets, or a post-market safety signal—not just a probability score. Structure-only models also have no way to incorporate this pharmacological context because it lives in annotation databases, not in SMILES strings.

We built **DrugInsight** to address both issues. The system fuses three evidence sources: AttentiveFP GNN encodings for molecular structure, DrugBank pharmacological annotations for known interaction mechanisms, and TWOSIDES [9] post-market adverse event signals. An adaptive tiered fusion layer weights each source based on what is actually available for a given query pair. When DrugBank contains a curated record for the pair, that record anchors the prediction; when

it does not, the system shifts weight toward the ML output and any available pharmacovigilance signal. Explanations are generated by a rule-based module that reports metabolic competition via CYP enzyme overlap, pharmacodynamic risk via shared targets, and signal strength from TWOSIDES PRR values. Because the explainer is entirely rule-based, it does not inherit the hallucination risk of LLM-generated clinical text. Per-source confidence scores are reported alongside each prediction so that users can judge which evidence drove a particular result.

The model is tested under a strict drug-level split in which both members of every validation pair are withheld from training entirely. This evaluates generalization to unseen compounds rather than memorization of structural patterns for drugs the model has already processed. We report a best validation AUC of 0.7065 under this protocol.

The DrugInsight application is available at https://druginsight.streamlit.app/. Source code, documentation, and model architecture details are at https://github.com/AymanUzayr/DrugInsightv2.

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

Today, graph neural networks are widely used in molecular property prediction tasks. Although DeepDDI [3]used structural similarity profiles for DDI prediction, it was unable to perform learning directly on molecular graphs. In contrast, KGNN [4] combined knowledge graph embeddings with molecular information, while SSI-DDI [5]conducted modeli introduced multiple resolution graph representations, MHCADDI [7] took advantage of multi-head cross-attention to capture inter-molecular relationships. In our project, we use AttentiveFP [8], with graph attention and virtual super-nodes.

B. Knowledge-Based and Database Approaches

The DrugBank [2] database forms the main repository for curating drug interactions, with mechanisms being annotated at the pair level. Similar information can be obtained using commercial databases like Medscape and Lexicomp. These sources give high precision but suffer from the limitation that they only contain information on manually verified pairs and thus cannot cover the entire combination space; besides, information on new drugs is not abundant either.

C. Pharmacovigilance Signal Detection

FAERS and TWOSIDES [9] are data sources providing information about risks posed by drug interactions based on post-marketing monitoring. PRR represents the extent to which adverse drug interactions are more commonly reported than for single medications. Prior works have already employed such signals for retrospective analysis. On the other hand, DrugInsight regards them as prospective evidence sources to be incorporated into the prediction pipeline.

D. Explainability in DDI Prediction

Most DDI prediction models produce predictions without generating explanations. Attention mechanisms [8] allow us to compute feature importance scores implicitly, but they cannot give interpretable explanations directly understandable by clinicians. Knowledge-based solutions [10] have explored ontology reasoning; however, they necessitate extensive knowledge engineering efforts. The proposed method, DrugInsight, incorporates the interactions between CYP enzymes, DrugBank's drug action, and pharmacovigilance reports into an explanation system without relying on LLMs.

III. Methodology

A. System Architecture Overview

Fig. 1 is the architecture of DrugInsight . It all starts when the user enters two drugs, the system then searches through DrugBank to find the SMILES of the entered drugs. The retrieved SMILES are converted to molecular graphs, encodes each graph using the AttentiveFP architecture. The next step involves extracting pharmacological pair features, classifies interaction probability via an MLP, fuses the prediction with curated evidence and pharmacovigilance signals, and produces a structured clinical explanation.

B. Data Sources and Preprocessing

1) Extraction from the DrugBank database: The DrugBank [2] databases can be extracted to form CSV files consisting of various aspects including drug details, interactions, enzyme information, targets, transporters, carriers, pathways, and SMILES. Drugs that lack proper SMILES (and therefore cannot be parsed using RDKit [11] MolFromSmiles) are filtered out, leaving a total of about 3,803 drugs with chemical structure information.

2) TWOSIDES Pharmacovigilance Data: The TWOSIDES database [9] contain adverse event reports after market release along with information such as PRR values for drug pairs. Matching across databases uses DrugBank's RxCUI IDs, but uses drug normalization where the RxCUI ID does not match.

3) Interaction Enrichment: Each interaction pair is enriched with 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 limit outlier influence; duplicates are deduplicated by keeping the highest-PRR record. The final dataset contains approximately 936,178 interaction pairs.

C. Molecular Graph Representation

Each drug's SMILES string is converted into a PyTorch Geometric [12] `Data` object using RDKit [11]. Hydrogen atoms are explicitly added before feature extraction. The

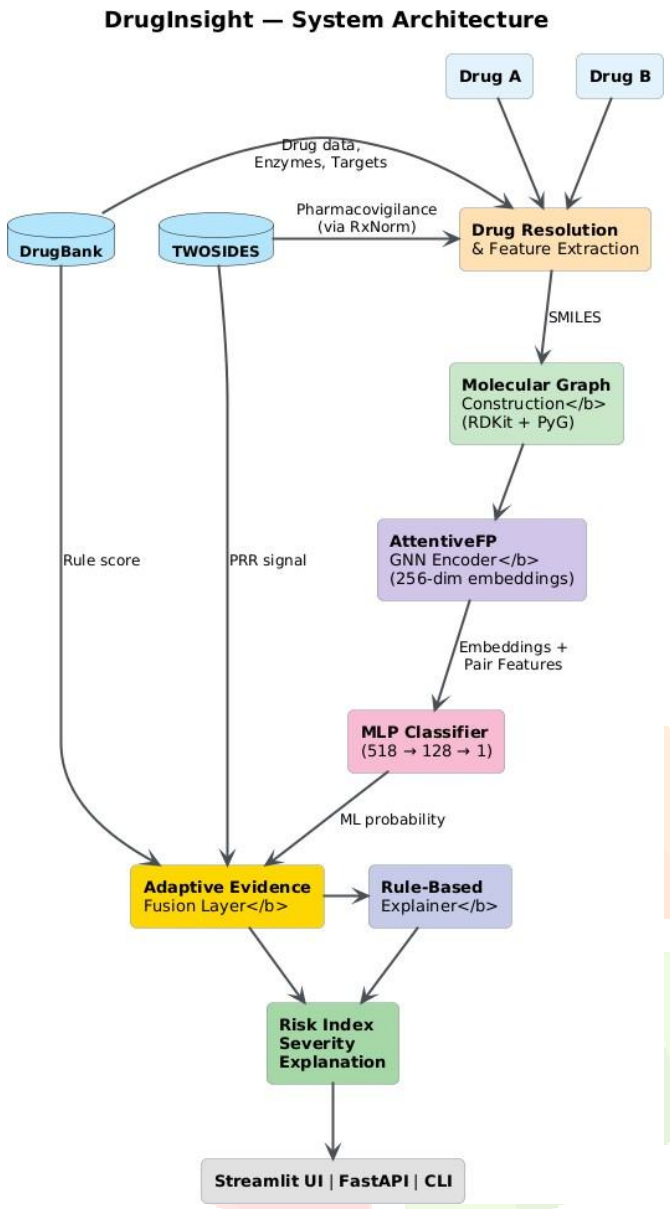

Fig. 1. End-to-end DrugInsight pipeline. Two drug identifiers are resolved, encoded as molecular graphs by AttentiveFP, classified by the MLP, and the prediction is fused with DrugBank rules and TWOSIDES signals before explanation generation.

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.

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.

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	/245.05	Peak adverse safety signal
TWOSIDES signals count	TWOSIDES	/1882	Breadth of safety signals
TWOSIDES found flag	TWOSIDES	binary	Signal availability

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.

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:

- **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):** $P_{fused} = 1.0 \times ML$.

When curated evidence is available, it anchors the prediction; when it isn't, the system falls back to ML alone. A risk index is derived as $fused\text{–}prob \times 100$ (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.

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

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 (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}
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
Train/Val split	80/20 drug-level
Negative ratio	1:1
Hard neg fraction	70%

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

2) Hard Negative Sampling: Negative 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 ($\times 0.5$).

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

C. Evaluation Protocol

The primary metric is ROC AUC, used for model selection and early stopping. Average Precision, Accuracy, and Confusion Matrix are tracked as secondary metrics.

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, which allows algorithms to memorize individual drug structural signatures.

B. Case Studies

To demonstrate DrugInsight's practical capability and the advantage of multi-source evidence fusion, we present six clinically motivated drug pair evaluations.

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.
Fusion Scores: DrugBank 0.280, ML 0.736, TWOSIDES 0.0
Fusion Weights: Rule 0.7 · ML 0.3 · TWOSIDES 0.0
DrugBank Checker Output [14]: Moderate — The excretion of Benzylpenicillin can be decreased when combined with Probenecid. Probenecid inhibits tubular secretion of penicillin, typically elevating plasma levels 2- to 4-fold [15].

2) Case 2: Amoxicillin and Clavulanic Acid: This intentional synergistic combination tests whether the system appropriately flags co-administration risks via TWOSIDES when DrugBank has no curated record.

DrugInsight Output:
Predicted Interaction: Moderate *Risk Index*: 63/100
TWOSIDES: High PRR detected (PRR=125.0)
Fusion Scores: DrugBank 0.000, ML 0.882, TWOSIDES 0.609
DrugBank Checker Output [14]: No Interactions Found.

3) Case 3: Doxorubicin and Sildenafil:

DrugInsight Output:
Predicted Interaction: Moderate *Risk Index*: 58/100
Shared Biology: ENZYMES CYP2D6, CYP3A4
Fusion Scores: DrugBank 0.500, ML 0.764, TWOSIDES 0.0
DrugBank Checker Output [14]: Moderate — The serum concentration of Doxorubicin can be increased when combined with Sildenafil. P-glycoprotein inhibitors increase nuclear translocation of doxorubicin [16].

4) Case 4: Ritonavir and Darunavir:

DrugInsight Output:
Predicted Interaction: Major *Risk Index*: 94/100
Shared Biology: ENZYMES CYP2D6, CYP3A4; TARGETS Gag-Pol polyprotein
Mechanism: Ritonavir is a known inhibitor of CYP3A4, the primary enzyme responsible for Darunavir metabolism. This inhibition reduces Darunavir clearance, increasing plasma concentration and toxicity risk.
Fusion Scores: DrugBank 0.730, ML 0.932, TWOSIDES 0.000
DrugBank Checker Output [14]: No separate record (this interaction is the intended pharmacokinetic boosting effect used therapeutically in HIV regimens).

5) Case 5: Acetaminophen and Lisinopril:

DrugInsight Output:
Predicted Interaction: Minor *Risk Index*: 19/100
Mechanism: Acetaminophen may decrease the excretion rate of Lisinopril.
Fusion Scores: DrugBank 0.825, ML 0.621, TWOSIDES 0.0
DrugBank Checker Output [14]: Minor — Acetaminophen may decrease the excretion rate of Lisinopril which could result in a higher serum level [17].

6) Case 6: Hyaluronic Acid and Betaine:

DrugInsight Output:
Predicted Interaction: No Significant Interaction

Risk Index: 45/100
Shared Biology: None
Fusion Scores: DrugBank 0.000, ML 0.452, TWOSIDES 0.000
Fusion Weights: Rule 0.0 · ML 1.0 · TWOSIDES 0.0
DrugBank Checker Output [14]: No Interactions Found.

C. Multi-Source Fusion Advantage

Adaptive evidence fusion has an inherent advantage over the single-source DDI checkers for reasons that include, but are not limited to:

- **Scope**:DrugBank includes only manually curated pairs. If a particular drug pair does not have entries in DrugBank, the evidence fusion module tilts its balance towards the ML-based predictor (70%) and TWOSIDES (20%), thus allowing predictions outside the manually curated database.
- **Reliability**: The use of three independent evidence streams makes the algorithm resistant to gaps in any individual source. There could be cases in which the drug pair has no data in the DrugBank, but has good signals in TWOSIDES, or vice versa.
- **Clarity**: Individual scores from the component algorithms are provided separately—the rule score, ML score, and TWOSIDES score—allowing clinicians to see the evidence that goes behind each particular prediction.
- **Confidence**: The algorithm gives confidence scores for each source (high/moderate/low/weak), as well as overall confidence.

VI. Limitations and Future Work

A. Current Limitations

- **Moderately discriminative performance with respect to cold start evaluation**: An AUC of 0.7065 is obtained when evaluating on drug level cold start pairs, which is significantly more difficult than random pair-level splits used by other competing approaches. Though this highlights good generalization to unseen compounds, the exact number shows that the current architecture of AttentiveFP along with a 12-dimensional pharmacological feature vector does not fully represent available interaction signals. The maximum obtainable result via this method of evaluation has not been characterized by any comparable cold-start benchmarks.
- **Pair symmetry assumption**:The current method assumes that the interaction between drugs A and B is equivalent to interaction between drugs B and A. This assumption leads to training using drug swap augmentation. In reality, some interactions are asymmetric. E.g., Drug A can inhibit CYP3A4 activity regarding metabolism of Drug B but Drug B might not affect Drug A. The molecular graph encoder takes each drug separately, thus the architecture cannot learn asymmetrical interactions.
- **Untrained severity classifier**:The severity prediction head of the architecture is specified as a linear classifier of shape Linear(128 -¿ 3). It is untrained since there are no standardized machine-usable labels in terms of drug interaction severity. Therefore, severity classes are determined by a threshold over calculated risk indices only.
- **Heuristic fusion layer design**: Tier-switching boundaries and confidence ratios are static heuristics not optimized from data.
- **Sparse TWOSIDES coverage**: The RxNorm bridge fails for a substantial proportion of query pairs, limiting pharmacovigilance channel contribution.
- **Single split evaluation**: No k-fold cross-validation or formal ablation study has been performed.
- **SMILES dependency**: The GNN pipeline requires a parseable SMILES string, limiting coverage to approximately 3,803 drugs.

B. Future Directions

- Supervised severity categorization based on structured clinical pharmacological data.
- Directionality learning of interaction can be done through cross-attention on drug graph embeddings.
- Data integration for evidence can be achieved through a lightweight gating mechanism.
- Polypharmacy drug interaction detection uses multi-drug hypergraphs as drug encoding.
- Drug interaction category classification detects CYP-dependent metabolism, pharmacodynamics targets and transport inhibitors.
- Risk stratification per population considers context vectors of patients.
- Clinical validation involves prospective validation on patient records.

VII. Conclusion

The foundation of DrugInsight lies in two premises: no single data source includes all clinically important drug pairs, and binary predictions without explanations based on mechanisms of actions are of little use to clinicians. These problems are addressed in turn by using an attentive neural layer to combine molecular structural information provided by AttentiveFP GNNs with pharmacological features from DrugBank and real-world evidence from TWOSIDES, dynamically adjusted in weight depending on the presence of evidence for each specific drug pair, and by generating explanations in the form of rules, bypassing a language generation component altogether.

The cold-start AUC score on novel drug pairs of 0.7065 achieved during validation demonstrates not only that generalization to unseen chemical space takes place, but that there is much room for improvement. Namely, the cold start strategy implemented here is significantly harder than random pairwise splits common to comparable approaches, and, furthermore, the existing 12-dimensional vector of pharmacological features used as part of input does not cover all important variables.

Among the key problems left open are the untrained severity head (currently, risk index thresholds are set up manually), symmetry implied by training augmentation, and sparsity of coverage by RxNorm, leaving the use of TWOSIDES limited mostly to a subset of predictions. Future development should take care of those issues alongside expanding the analysis from pairs to polypharmacy as a whole and conducting a clinical validation experiment.

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, 1989.