

Informative Subgraph Extraction with Deep Reinforcement Learning for Drug-Drug Interaction Prediction

Jiancong Xie¹, Wentao Wei^{1,2}, Chi Zhang¹, Jiahua Rao^{1*}, Yuedong Yang^{1,3*}

¹School of Computer Science and Engineering, Sun Yat-sen University, China

²Pengcheng Laboratory, China

³Key Laboratory of Machine Intelligence and Advanced Computing, Sun Yat-sen University, China
{xiejc3, weiw8, zhangch588}@mail2.sysu.edu.cn, {raojh7, yangyd25}@mail.sysu.edu.cn

Abstract

Drug-drug interaction (DDI) prediction is pivotal for drug safety and clinical decision-making. Recently, subgraph-based methods utilizing knowledge graphs (KGs) and domain information have achieved promising results by extracting informative subgraphs for DDI prediction. However, existing subgraph extraction methods are typically coarse-grained and nonspecific, facing two key limitations: First, they are constrained by the vast and noisy nature of real-world KGs, making it challenging to identify the most informative substructures from the massive space of candidate subgraphs. Second, current methods often fail to exploit the molecular structural specificity of drugs to selectively extract relevant subgraphs, lacking effective integration of molecular structure information with knowledge graph context. To address these challenges, we propose **RISE-DDI**, a novel framework for **Reinforced-based Informative Subgraph Extraction** approach for drug-drug interaction prediction. Specifically, RISE-DDI formulates the subgraph extraction as a Markov Decision Process (MDP) and leverages a deep reinforcement learning (RL) agent to dynamically and adaptively extract the most informative and context-specific subgraphs for each drug pair. The agent is guided by a learnable structure-aware reward model that considers both the topological context from the knowledge graph and the molecular features of the drug pairs, thereby encouraging the selection of subgraphs that are both structurally relevant and biologically informative. Extensive experiments on DDI benchmark datasets demonstrate that our method outperforms state-of-the-art baselines in both transductive and inductive scenarios, achieving improvements of up to 20%. Furthermore, visualization analyses of the extracted subgraphs highlight the interpretability of our model, providing insights into the underlying mechanisms of drug interactions.

Code — <https://github.com/biomed-AI/RISE-DDI>

Introduction

Drug-drug interactions (DDIs) represent a significant challenge in clinical pharmacology and drug development (Percha and Altman 2013), as they may result in adverse drug reactions, diminished therapeutic efficacy, or even life-threatening events (Vilar et al. 2014; Xie et al. 2024a). The

growing prevalence of polypharmacy has made the accurate prediction of potential DDIs increasingly important for safeguarding patient safety and optimizing therapeutic strategies (Rao et al. 2024).

In recent years, computational approaches for DDI prediction have garnered significant attention. Notably, the incorporation of knowledge graphs (KGs), which encode rich semantic and structural information about drugs, targets, and their interactions, has shown great promise in modeling the intricate relationships between drugs (Lin et al. 2020; Su et al. 2022b,a). By leveraging the auxiliary information provided by KGs and additional domain-specific information, researchers have been able to capture more comprehensive and nuanced patterns underlying drug interactions.

However, knowledge graphs are often extremely large in scale, which inevitably introduces redundant and noisy information, making it challenging for DDI prediction models to extract meaningful signals from substantial noise (Du et al. 2024; Xie et al. 2024b). To address these challenges, subgraph-based methods have been proposed to efficiently extract informative local structures from the KG for each drug pair (Su et al. 2024). For example, SumGNN (Yu et al. 2021) and LaGAT (Hong et al. 2022) extracted the fixed k -hop neighborhoods for modeling drug pair features. CoSMIG (Rao et al. 2022) further selected informative subgraphs through the random walk algorithm. Despite their effectiveness, these coarse-grained methods lack adaptability, since the extracted subgraphs are fixed and non-learnable, which may introduce noise or overlook important information. To alleviate this issue, recent advances have introduced learnable subgraph extraction methods to adaptively select the most informative substructures from the knowledge graph (Wang, Yang, and Yao 2024; Du et al. 2024; Abdullahi et al. 2025). For instance, KnowDDI (Wang, Yang, and Yao 2024) utilizes node features to add or remove edges within a given subgraph for a specific drug pair. CSSE-DDI (Du et al. 2024) leverages neural architecture search (NAS) to automatically identify the optimal subgraph structures.

While these approaches have demonstrated notable performance, they still face several critical challenges. First, the subgraph extraction process is fundamentally constrained by the enormous search space inherent in real-world knowledge graphs, which often contain millions of triples (Him-

*Co-corresponding authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

melstein and Baranzini 2015; Zheng et al. 2021). The number of potential subgraph combinations grows exponentially with the size of the graph, making it computationally infeasible to exhaustively search for the optimal subgraph for each drug pair. To cope with this challenge, existing approaches typically resort to compromising strategies that restrict the search space through heuristic methods. For instance, CSSE-DDI explores only the ego-network of each drug by selecting k -hop neighbors with a flexible k value. While such simplification reduces computational complexity, it inevitably excludes informative nodes or includes noisy ones, leading to a suboptimal representation of drug pairs. Second, current subgraph extraction methods often overlook the integration of molecular structural information. In fact, the molecular structure features play an important role in revealing their biological activities and interaction mechanisms. Therefore, how to effectively combine molecular structural features with knowledge graph context during subgraph extraction and optimization remains a critical challenge.

To address these challenges, we propose **RISE-DDI**, a novel framework that leverages deep reinforcement learning (RL) to dynamically extract informative subgraphs for DDI prediction. Specifically, RISE-DDI formulates the subgraph extraction process as a Markov Decision Process (MDP), where an RL agent sequentially explores the vast space of possible subgraphs in a trial-and-error way. To guide the agent toward identifying the most informative subgraphs, we carefully designed a learnable reward model that evaluates the predictive utility of the extracted subgraph by considering both the structural context from the knowledge graph and the molecular features of the drug pairs. This enables the agent to adaptively select the most relevant nodes and edges for the prediction task, thereby reducing the risk of omitting important nodes or introducing noise. Extensive experiments on benchmark datasets demonstrate that RISE-DDI consistently outperforms state-of-the-art baselines in both transductive and inductive settings, achieving up to 20% improvement in the most challenging scenario where the test drugs are all unseen. Moreover, the RL-based subgraph selection enhances interpretability by explicitly highlighting the most relevant nodes and relations in the knowledge graph, providing valuable insights into drug interaction mechanisms.

Our main contributions are summarized as follows:

- We propose RISE-DDI, a novel framework that leverages deep reinforcement learning to extract informative subgraphs from knowledge graphs for DDI prediction.
- We introduce a structure-aware reward model that integrates both subgraph features and the molecular structure features to enhance subgraph extraction and DDI prediction.
- We conduct extensive experiments on benchmark datasets, demonstrating that RISE-DDI consistently outperforms state-of-the-art methods and provides interpretable insights into the mechanisms underlying drug interactions.

Related Work

Knowledge Graph-based DDI Prediction

The rapid growth of biomedical data has led to the emergence of KGs such as KEGG (Kanehisa et al. 2017), PharmKG (Zheng et al. 2021), and Hetionet (Himmelstein et al. 2017), offering structured information and semantic links for DDI prediction. Karim et al. (Karim et al. 2019) proposed Conv-LSTM, which leverages knowledge graph embedding techniques such as TransE and SimpleIE to learn effective representations of entities within KGs. To capture the high-order associations between entities, Lin et al. (Lin et al. 2020) introduced KGNN, which integrates graph convolutional networks (GCNs) on KGs to predict DDIs. Building upon these advances, KG2ECapsule (Su et al. 2022b) further enhances the expressive ability of KG-based models by employing a two-layer capsule network by learning the non-linear associations in KGs. Most recently, Su et al. (Su et al. 2022a) proposed DDKG, which incorporates graph attention networks (GATs) to identify the significance of different triplets within the KG. By assigning adaptive weights to different edges and entities, the model can focus on the most informative and relevant relationships, thereby improving the discriminative capability of DDI prediction models.

Subgraph-based DDI Prediction

Knowledge graphs often contain noise that can harm model performance. To address this issue, subgraph-based approaches have been proposed by extracting relevant subgraphs around drug pairs, reducing irrelevant information (Rao et al. 2025a). For example, SumGNN (Yu et al. 2021) extracts k -hop neighbors of drug node pairs and introduces a knowledge summarization module to identify effective interaction paths. Similarly, LaGAT (Hong et al. 2022) employs GAT within k -hop subgraphs to adaptively aggregate information from important neighbors. TIGER (Su et al. 2024) further investigates various subgraph extraction strategies by sampling subgraphs via random walks and PageRank-based probabilities. More recently, dynamic subgraph sampling methods have been proposed to better capture the most relevant local structures for specific drug pairs. KnowDDI (Wang, Yang, and Yao 2024) first learns drug representations through KG and constructs a knowledge subgraph for each drug-pair to interpret the predicted DDI. CSSE-DDI (Du et al. 2024) further leverages neural architecture search to automatically discover the optimal subgraph for accurate DDI prediction.

Methods

Figure 1 illustrates the overall architecture of RISE-DDI, which comprises two main components: the subgraph sampler f_{sampler} acting as an RL agent, and the interaction predictor $f_{\text{predictor}}$ acting as the reward model. The f_{sampler} identifies the most informative substructures \hat{G} from the knowledge graph G to learn contextual representations for drug entities, while the $f_{\text{predictor}}$ evaluates the predictive utility of the extracted subgraphs by integrating both their topological context and the molecular features of the drug pairs, providing feedback to guide the subgraph extraction

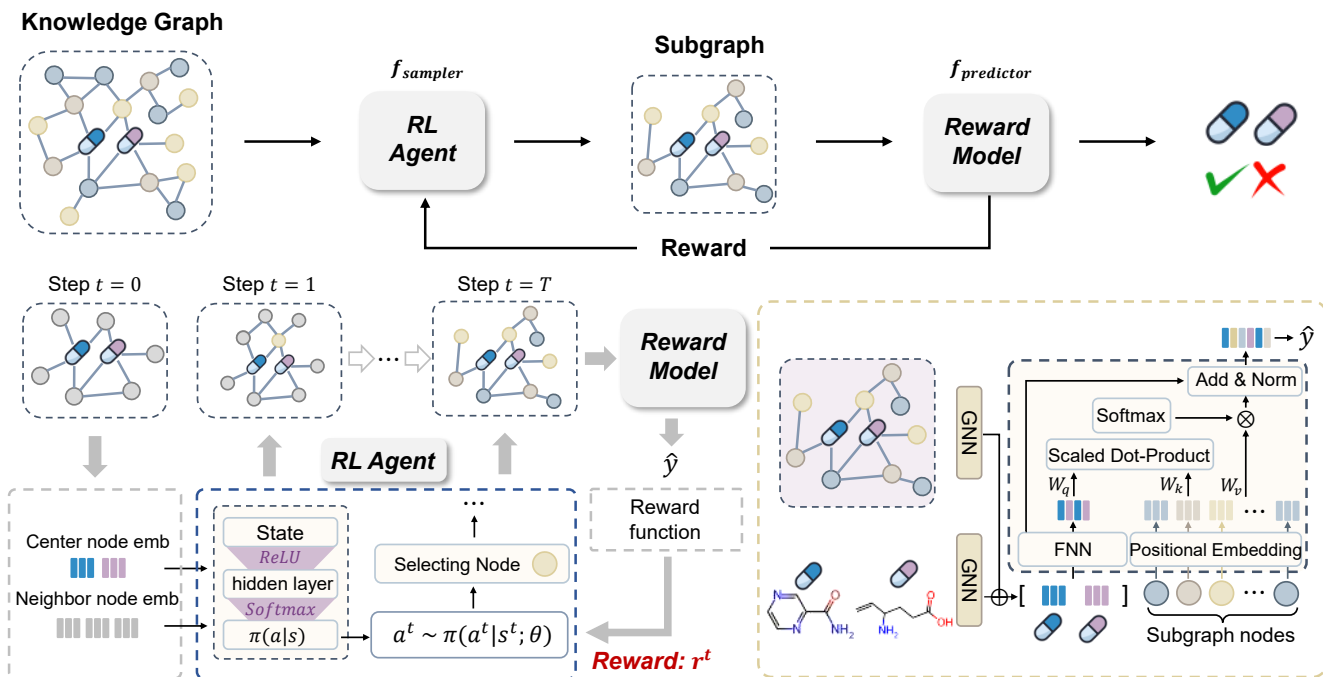


Figure 1: An overview of the RISE-DDI framework. Given a query drug pair, the RL agent first extracts informative subgraphs from the knowledge graph via reinforcement learning. The sampled subgraphs, together with molecular structures, are then fed into the structure-aware module to learn pairwise representations for drug-drug interaction prediction. The entire framework is optimized iteratively in an end-to-end manner.

process. Both modules are trained iteratively in an end-to-end manner.

Subgraph Sampling with Reinforcement Learning

Intuitively, the relevant background information for a drug pair (u, v) is embedded within a specific subgraph tailored to that pair. An ideal subgraph should (1) be information-rich, containing features that are beneficial for predicting the interaction between u and v , and (2) avoid redundant information. However, extracting such subgraphs from a large-scale knowledge graph is an NP-hard problem. To address this challenge, we propose to employ deep reinforcement learning to sample subgraphs in a trial-and-error manner.

Specifically, we formulate the subgraph sampling process for each drug pair (u, v) as a Markov Decision Process characterized by states, actions, transitions, and rewards, and train a subgraph sampling agent f_{sampler} as follows:

- **State:** At step t , the state s^t is represented as the current subgraph $\hat{G}_{u,v}^t$, which includes the drug pair nodes and the set of neighboring nodes that have already been sampled. Initially when $t=0$, $\hat{G}_{u,v}^0$ contains only u and v . The state vector \mathbf{e}_{s^t} is defined as the sum of embeddings of all nodes in the subgraph:

$$\mathbf{e}_{s^t} = \sum_{i \in \hat{G}_{u,v}^t} \mathbf{e}_i. \quad (1)$$

- **Action:** The action space of s^t consists of all neighboring entities of the current subgraph, as well as a virtual

node that serves as a termination node, formally defined as $\mathcal{N}(\hat{G}_{u,v}^t)$. The action is generated by the subgraph sampling agent:

$$a^t \sim f_{\text{sampler}}(\mathbf{e}_{s^t}, \{\mathbf{e}_i, i \in \mathcal{N}(\hat{G}_{u,v}^t)\}). \quad (2)$$

By taking an action, the agent adds a neighboring entity to the subgraph.

- **Transition:** The transition model P defines the probability of moving to a new state given the current state and action. In knowledge graphs, if state s^t transitions to s^{t+1} via action a , then $P(s^{t+1}|s^t) = 1$; otherwise, $P(s^{t+1}|s^t) = 0$.
- **Reward:** To ensure the quality of the sampled subgraph, we design a task-specific reward function. Specifically, the reward is defined based on whether the inclusion of the subgraph brings the prediction closer to the ground-truth label compared to the prediction without the subgraph, which is formulated as follows:

$$r^t = \alpha[y \cdot (p - p_0) + (1 - y) \cdot (p_0 - p)], \quad (3)$$

where α is a scaling coefficient that controls the magnitude of the reward, $y \in \{0, 1\}$ denotes the ground-truth label. p and p_0 are the prediction scores obtained with and without the subgraph, respectively, which are calculated as follows:

$$\begin{aligned} p &= f_{\text{predictor}}(\hat{G}_{u,v}^t, u, v) \\ p_0 &= f_{\text{predictor}}(\hat{G}_{u,v}^{(0)}, u, v) \end{aligned} \quad (4)$$

Training Policy Network. To optimize the sampler parameters Θ_s of f_{sampler} , we maximize the expected cumulative discounted reward as follows:

$$\max_{\Theta_s} \sum_{(u,v) \in \mathcal{T}} \mathbb{E} \left(\sum_{t=1}^T \lambda^{t-1} r^t \right), \quad (5)$$

where λ is the decay factor, \mathcal{T} is the training pairs. The optimal policy guides the sampler to extract subgraphs from the knowledge graph that are most informative for predicting drug-drug interactions, while maximizing the accumulated rewards.

During subgraph sampling, the RL agent iteratively selects neighboring nodes with the highest selection probabilities, expanding the subgraph until reaching a maximum of T steps or a termination node. The value of T determines the subgraph size. Additionally, candidate neighbors are limited to the k -hop neighborhood of the initial drug pair, with k set to match the number of layers in the GNN encoder. This ensures that sampled nodes are within the range that can influence the drug pair’s representation during message passing, as nodes beyond k hops would not contribute in the GNN.

Structure-aware Reward Model

In this section, we propose a structure-aware reward model $f_{\text{predictor}}$ that quantitatively evaluates the informativeness of extracted subgraphs by jointly considering the topological context from the knowledge graph and the molecular features of drug pairs. Given an extracted subgraph, $f_{\text{predictor}}$ computes the DDI interaction probability, which is used as the reward signal to guide the reinforcement learning agent.

Feature Encoder. Given the sampled subgraph $\hat{G}_{u,v} = f_{\text{sampler}}(u, v)$, we first employ a Graph Neural Network (GNN) to encode the subgraph and learn node representations:

$$\mathbf{h} = \text{GNN}_g(\mathbf{X}_g, \mathbf{R}_g, \hat{G}_{u,v}), \quad (6)$$

where \mathbf{X}_g and \mathbf{R}_g are the initial features of subgraph nodes and relationships, respectively. Similarly, a separate GNN is applied to encode the molecular structures M to extract molecular-level features:

$$\mathbf{m} = \text{GNN}_m(\mathbf{X}_m, \mathbf{R}_m, M), \quad (7)$$

where \mathbf{X}_m and \mathbf{R}_m are the initialization features of atoms and bonds, respectively. In our implementation, we adopt the relational graph transformer proposed by Su et al. (Su et al. 2024) to extract node and molecular features.

Structure-aware interaction Predictor. To enhance modality integration, we introduce an attention-based fusion mechanism that effectively integrates both pairwise molecular structural features and subgraph-level characteristics, leveraging both local and global contextual information.

Specifically, we fuse the features of subgraph nodes and the target drug pair in a weighted manner, with attention scores calculated between each subgraph node and the target drug pair. Firstly, to enrich graph structural information, we add the position embedding for each subgraph node:

$$\hat{\mathbf{h}}_i = \text{MLP}(\mathbf{h}_i + \mathbf{p}_{i,u,v}), \quad i \in \hat{G}_{u,v}, \quad (8)$$

where $\mathbf{p}_{i,u,v}$ is the relative positional encoding of a node i relative to a target link (u, v) . Inspired by previous work (Shomer et al. 2024), we adopt the Personalized PageRank (PPR) scores to measure the influence of one node on another:

$$\mathbf{p}_{i,u,v} = \text{MLP}(ppr(i, u) + ppr(i, v)). \quad (9)$$

Meanwhile, the Drug pair embedding of (u, v) is calculated by integrating the node features with their molecular structure features:

$$\mathbf{z}_{u,v} = \text{MLP}((\hat{\mathbf{h}}_u + \mathbf{m}_u) || (\hat{\mathbf{h}}_v + \mathbf{m}_v)), \quad (10)$$

where $||$ is the concatenate operation. Then, we fuse these two modalities of features through the Multi-Head Attention (MHA) mechanism:

$$w_p(i, u, v) = \frac{\exp(\phi_p(\hat{\mathbf{h}}_i, \mathbf{z}_{u,v}))}{\sum_{j \in \hat{G}_{u,v}} \exp(\phi_p(\hat{\mathbf{h}}_j, \mathbf{z}_{u,v}))} \quad (11)$$

where ϕ_p is a learnable network of head p . The attention weight $w_p(i, u, v)$ can be considered as the impact of the node i on the target link (u, v) at head p . We then aggregate the node embedding to capture relevant information from the subgraph nodes (Rao et al. 2025b).

$$\mathbf{s}_{u,v} = \sum_{p \in P} \sum_{i \in \hat{G}_{u,v}} w_p(i, u, v) \hat{\mathbf{h}}_i, \quad (12)$$

$$\hat{\mathbf{z}}_{u,v} = \text{Norm}(\mathbf{z}_{u,v}) + \mathbf{s}_{u,v}.$$

where P is the number of attention heads. The predicted interaction score is obtained as follows:

$$\hat{y}_{u,v} = \text{MLP}(\hat{\mathbf{z}}_{u,v}). \quad (13)$$

Model Optimization

Finally, we optimize the subgraph sampler f_{sampler} with parameters Θ_s and the reward model $f_{\text{predictor}}$ with parameters Θ_p in an iterative manner.

Predictor Optimization. We first freeze f_{sampler} and sample subgraph, which are then fed into $f_{\text{predictor}}$ and optimize Θ_p through the cross-entropy loss formulated as follows:

$$\mathcal{L}_p = - \sum_{(u,v) \in \mathcal{T}} y_{u,v} \log(\hat{y}_{u,v}) + (1 - y_{u,v}) \log(1 - \hat{y}_{u,v}). \quad (14)$$

Sampler Optimization. We then freeze $f_{\text{predictor}}$ and train f_{sampler} through the policy gradient based RL (REINFORCE) guided by the reward output from $f_{\text{predictor}}$. The gradients of Θ_s can be calculated as follows:

$$\begin{aligned} \nabla_{\Theta_s} \mathcal{L}_s &= \nabla_{\Theta_s} \sum_{(u,v) \in \mathcal{T}} \mathbb{E} \left(\sum_{t=1}^T \lambda^{t-1} r^t \right) \\ &\approx \sum_{(u,v) \in \mathcal{T}} \frac{1}{T} \sum_{t=1}^T [\lambda^{t-1} r^t \nabla_{\Theta_s} \log \mathbb{P}(a^t | s^t)]. \end{aligned} \quad (15)$$

	DrugBank	KEGG	OGB-biokg
#Nodes	391,116	129,910	93,773
#Relations	71	167	13
#Links	1,587,305	362,870	3,892,462
#Drugs	1,052	786	808
#DDIs	10,404	13,787	111,520
Avg. Deg.	128.5	55.4	97.7

Table 1: The statistics of three DDI benchmarks.

Experiments

In this section, we conduct extensive experiments on three DDI benchmarks under both transductive and inductive settings. We also perform an ablation study to evaluate the contribution of the two main modules, as well as an interpretability analysis.

Experimental Setup

Datasets. To evaluate the effectiveness of our method, we conduct experiments on three widely used DDI benchmarks with different scales, including DrugBank (Wishart et al. 2008), KEGG (Kanehisa et al. 2017), and OGB-biokg (Su et al. 2022a). We also collected SMILES for each drug and transformed them to molecular 2D graphs using rdkit (Landrum et al. 2013). Statistics of the benchmark datasets are presented in Table 1.

Baselines. We compare our method with a variety of DDI baselines that can be categorized as follows:

- **GNN-based Methods:** We consider three classic GNN methods: GCN (Kipf and Welling 2016), GAT (Veličković et al. 2017), and GIN (Xu et al. 2018). Note that these baselines are trained only using the DDI networks.
- **Knowledge graph-based Methods:** We select the representative knowledge graph-based models, including KG2ECapsule (Su et al. 2022b), and DDKG (Su et al. 2022a).
- **Subgraph-based Methods:** We adopt subgraph-based methods, including KnowDDI (Wang, Yang, and Yao 2024), CSSE-DDI (Du et al. 2024), and TIGER (Su et al. 2024). KnowDDI and CSSE-DDI employed an adaptive subgraph selection strategy, while TIGER used a fixed probability-based subgraph.

Experimental Settings. We evaluate our model under two scenarios: (1) **transductive scenario**, where all DDI pairs are randomly split with five-fold cross-validation, and (2) **inductive scenario**, where each DDI pair in the validation and test sets contains at least one drug that does not appear in the training set. This scenario, commonly referred to as the "cold start" setting, is particularly important for real-world applications where novel drugs frequently emerge. For both scenarios, we split the data into training, validation, and test sets with a ratio of 6:2:2, respectively. RISE-DDI is implemented using Pytorch v1.10.2 and trained on NVIDIA 4090

GPU. All trainable parameters are optimized by Adam algorithm with a learning rate of 0.001 for both the subgraph sampling and DDI prediction processes. We determined the subgraph layer L , RL step t , and dropout with grid search. Refer to the supplementary material for more details of the hyperparameters. For all baselines, they are retrained on the same machine with the same hyper-parameter settings reported in their original work.

Main Results

Performance in transductive scenario. Table 2 summarizes the overall results across all benchmarks in the transductive setting. As shown, RISE-DDI consistently outperforms all baseline methods on each dataset. Figure 2 further depicts the reward convergence of RISE-DDI on the test set, demonstrating its robustness and effectiveness in extracting informative subgraphs during training. Among the baselines, subgraph-based methods (KnowDDI, CSSE-DDI, and TIGER) achieve better performance than full-graph-based methods, underscoring the importance of local structural information for DDI prediction. Within the subgraph-based methods, CSSE-DDI outperforms TIGER in the DrugBank and OGB-biokg datasets, demonstrating the advantage of adaptive subgraph selection. Nevertheless, RISE-DDI surpasses these methods with an AUC improvement of 2.1%-4.7% on three datasets, which can be attributed to the proposed RL-based subgraph extraction strategy. Notably, the largest improvement is observed on the DrugBank dataset. One possible reason is that the DrugBank dataset contains a higher average number of neighbors per drug node, which provides a larger exploration space for the RL agent to discover optimal subgraphs.

Performance in inductive scenario. To further validate the robustness of RISE-DDI, we evaluate it and the subgraph-based baselines in an inductive scenario where test drug pairs include at least one drug that was not observed during training. As shown in Table 3, RISE-DDI consistently outperforms all baseline methods. In the most challenging setting, where both drugs in a query pair are unseen during training, RISE-DDI achieves an AUC improvement of 17.2%–20.3% over the best-performing baseline. We also observed that all methods exhibit performance degradation compared to the transductive setting. Among the baselines, two subgraph search-based methods (KnowDDI and CSSE-DDI) experience a notably larger decline than TIGER. This may be attributed to the limited generalization capability of existing subgraph extraction strategies when encountering unseen drug nodes, as well as the lack of integration of molecular structure information in these methods. In contrast, RISE-DDI demonstrates a substantially lower decline rate, highlighting its superior generalizability to novel drugs.

Ablation Study

Subgraph sampler. We first assess the effectiveness of the proposed RL-based subgraph extraction strategy by comparing it with three widely used fixed methods (Su et al. 2024), including: (1) **k -subtree-based Extractor** that constructs subgraphs with a node itself and its k -hop neighbors, (2)

Method	DrugBank			KEGG			OGB-biokg		
	AUC	AUPR	F1	AUC	AUPR	F1	AUC	AUPR	F1
GCN	0.613 \pm 0.57	0.561 \pm 0.65	0.581 \pm 1.21	0.735 \pm 0.62	0.693 \pm 0.81	0.701 \pm 1.09	0.792 \pm 0.91	0.821 \pm 0.13	0.701 \pm 2.01
GAT	0.627 \pm 0.61	0.553 \pm 0.69	0.572 \pm 1.18	0.763 \pm 0.58	0.701 \pm 0.85	0.713 \pm 1.07	0.803 \pm 0.94	0.833 \pm 0.14	0.719 \pm 2.13
GIN	0.641 \pm 0.59	0.573 \pm 0.67	0.593 \pm 1.15	0.749 \pm 0.61	0.713 \pm 0.82	0.724 \pm 1.11	0.815 \pm 0.89	0.841 \pm 0.12	0.718 \pm 2.09
KGNN	0.721 \pm 0.56	0.661 \pm 0.71	0.678 \pm 1.96	0.853 \pm 0.54	0.818 \pm 0.79	0.782 \pm 1.14	0.905 \pm 0.09	0.918 \pm 0.06	0.819 \pm 2.67
KG2ECapsule	0.713 \pm 0.92	0.656 \pm 0.58	0.669 \pm 0.74	0.896 \pm 0.44	0.869 \pm 0.69	0.828 \pm 0.49	0.908 \pm 0.10	0.920 \pm 0.07	0.823 \pm 0.20
DDKG	0.763 \pm 1.32	0.729 \pm 0.50	0.709 \pm 2.53	0.900 \pm 0.75	0.869 \pm 1.18	0.836 \pm 0.85	0.882 \pm 0.58	0.896 \pm 0.48	0.796 \pm 0.98
KnowDDI	0.782 \pm 0.81	0.799 \pm 0.73	0.755 \pm 0.89	0.914 \pm 0.24	0.881 \pm 0.37	0.853 \pm 0.28	0.902 \pm 0.16	0.929 \pm 0.18	0.838 \pm 0.22
CSSE-DDI	<u>0.879</u> \pm 0.79	<u>0.847</u> \pm 0.71	<u>0.821</u> \pm 0.91	0.933 \pm 0.24	0.927 \pm 0.37	0.871 \pm 0.25	<u>0.949</u> \pm 0.13	<u>0.960</u> \pm 0.17	<u>0.889</u> \pm 0.13
TIGER	0.867 \pm 0.77	0.837 \pm 0.68	0.803 \pm 0.94	<u>0.947</u> \pm 0.21	<u>0.935</u> \pm 0.35	<u>0.889</u> \pm 0.24	0.948 \pm 0.14	0.957 \pm 0.12	0.875 \pm 0.17
RISE-DDI (K-hop)	0.890 \pm 0.32	0.882 \pm 0.36	0.818 \pm 0.70	0.953 \pm 0.29	0.941 \pm 0.50	0.895 \pm 0.46	0.949 \pm 0.17	0.957 \pm 0.14	0.888 \pm 0.13
RISE-DDI (DW)	0.882 \pm 0.28	0.877 \pm 0.33	0.811 \pm 0.72	0.954 \pm 0.31	0.947 \pm 0.49	0.891 \pm 0.58	0.956 \pm 0.18	0.963 \pm 0.12	0.892 \pm 0.13
RISE-DDI (PR)	0.899 \pm 0.29	0.896 \pm 0.39	0.824 \pm 0.75	0.959 \pm 0.30	0.951 \pm 0.59	0.904 \pm 0.50	0.947 \pm 0.17	0.956 \pm 0.14	0.885 \pm 0.11
RISE-DDI	0.926 \pm 0.34	0.923 \pm 0.13	0.855 \pm 0.78	0.971 \pm 0.18	0.966 \pm 0.34	0.923 \pm 0.44	0.970 \pm 0.12	0.975 \pm 0.14	0.918 \pm 0.14
Improv. (%)	4.7	7.6	3.4	2.4	3.1	3.4	2.1	1.5	2.9

Table 2: Performance comparison on three datasets in the transductive setting, presented as the mean \pm standard deviation (%) across five-fold cross-validation. The best and second-best results are highlighted in bold and underlined, respectively. K-hop, DW, and PR are three variants of RISE-DDI that use k-hop subgraphs, DeepWalk-generated subgraphs, and PageRank-generated subgraphs, respectively.

		KnowDDI	CSSE-DDI	TIGER	RISE-DDI
DrugBank	One.	0.677 \pm 0.3	0.705 \pm 0.2	0.737 \pm 0.2	0.872 \pm 0.4
	Both.	0.593 \pm 0.5	0.622 \pm 0.4	0.651 \pm 0.2	0.786 \pm 0.3
	Avg.	0.668 \pm 0.2	0.693 \pm 0.2	0.725 \pm 0.1	0.859 \pm 0.2
	Δ	-14.6%	-21.1%	-16.4%	-7.24%
KEGG	One.	0.797 \pm 0.4	0.832 \pm 0.3	0.866 \pm 0.3	0.921 \pm 0.2
	Both.	0.665 \pm 0.4	0.697 \pm 0.3	0.729 \pm 0.2	0.893 \pm 0.3
	Avg.	0.781 \pm 0.2	0.814 \pm 0.2	0.848 \pm 0.1	0.917 \pm 0.2
	Δ	-14.5%	-12.8%	-10.5%	-5.46%
OGB-biokg	One.	0.687 \pm 0.5	0.718 \pm 0.3	0.749 \pm 0.2	0.819 \pm 0.2
	Both.	0.505 \pm 0.4	0.534 \pm 0.3	0.562 \pm 0.1	0.705 \pm 0.1
	Avg.	0.661 \pm 0.3	0.692 \pm 0.2	0.722 \pm 0.2	0.803 \pm 0.1
	Δ	-26.7%	-27.9%	-23.8%	-17.3%

Table 3: AUROC scores on three datasets in the inductive setting. "One" means one drug is unseen; "both" means both drugs are unseen.

DeepWalk-based subgraph Extractor that generates subgraphs via fixed-length random walks from each node, and **(3) Probability-based Extractor** that samples nodes based on PageRank probabilities. As shown in Table 2, RISE-DDI consistently outperforms all three fixed extraction strategies across the three benchmark datasets. While the performance of each fixed subgraph extraction method varies depending on the dataset, RISE-DDI achieves superior and more stable results in every case, demonstrating the effectiveness and robustness of the RL-based subgraph sampler.

Reward model. To evaluate the contributions of different components within the structure-aware reward model, we introduce three variants: (1) **RISE-DDI w/o MG**, which removes the molecular graph features; (2) **RISE-DDI w/o**




Figure 2: Reward convergence curves on the test sets of three datasets.

MHA, which removes the multi-head attention module and simply concatenates the subgraph and molecular features; and (3) **RISE-DDI w/o both**, which removes the entire structure-aware module. As shown in Figure 3, the full RISE-DDI consistently outperforms all its variants under both transductive and inductive settings. The removal of any individual module leads to a decline in performance, demonstrating the indispensability of each module. Furthermore, the performance degradation of the variant models is more pronounced in the inductive scenario, indicating that effectively integrating multi-view information is particularly crucial for robust generalization to unseen drugs.

Hyper-Parameter Studies

Effect of subgraph layers L . We first investigate the impact of subgraph layers L . Notably, the subgraph layers L is set to be the same as the GNN layers. As illustrated in Figure 4(a), increasing L initially leads to improved model performance, with optimal results observed when $L = 1$ or 2 . However, further increasing L results in a decline in performance. This decline can be attributed to the fact that a larger L exponentially expands the RL exploration space, thereby




Figure 3: Ablation study of the structure-aware pairwise learning module.




Figure 4: The result of RISE-DDI with varying value of subgraph layer L and RL step T .

increasing the complexity and difficulty of RL training and the computational cost.

Effect of RL steps T . The RL steps T determine the size of the sampled subgraphs. Figure 4(b) illustrates the effect of varying T . As observed, model performance continues to improve as T increases, since a small T limits the sampling space and hinders the model’s ability to identify optimal subgraphs. As the subgraph exploration space expands, performance gradually converges. Notably, the optimal subgraph size varies across different datasets. For instance, the optimal subgraphs identified on KEGG are generally smaller, while those on DrugBank and OGB-biokg are larger. These results demonstrate that our model can flexibly adapt to the structural properties of diverse datasets and automatically identify the most informative subgraphs in each case. Moreover, when the sampling step T is set to 40, the inference process remains efficient, incurring only an additional 1.5 seconds per batch compared to the baseline without sampling (with a batch size of 128).

Case Study

To illustrate the interpretability of RISE-DDI, we present two case studies from the KEGG dataset, both involving **Enalapril** but demonstrating distinct subgraph structures and mechanistic insights.

Figure 5(a) presents the subgraph centered on **Enalapril** and **Telmisartan**, both linked to RAAS-related pathways crucial in hypertension therapy (Te Riet et al. 2015). The subgraph also includes other RAAS inhibitors, such as

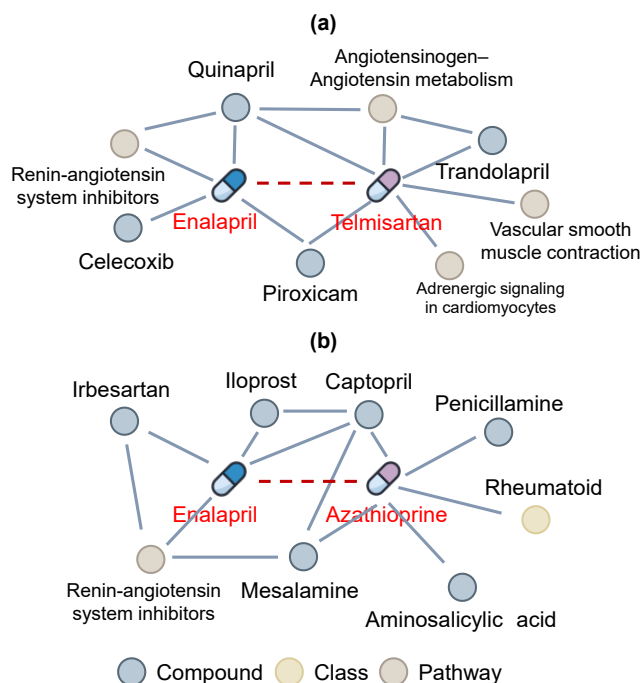


Figure 5: Case studies of subgraphs extracted for different drug pairs that share a common drug node.

Quinapril and Trandolapril (Balakumar and Jagadeesh 2015), highlighting the pharmacological similarity and overlapping mechanisms of Enalapril and Telmisartan. This overlap suggests that combined use of these drugs may heighten hypotension risk via cumulative RAAS inhibition (Investigators 2008).

Second, Figure 5(b) depicts the subgraph for the drug pair **Enalapril** and **Azathioprine**, with their interaction likely inferred through the path: Enalapril – Iloprost – Captopril – Azathioprine. Iloprost can enhance the anti-hypertensive effect of Enalapril or Captopril (Erre, Passiu et al. 2009), which indicates that Captopril and Enalapril have similar pharmacological actions. Given that Captopril–Azathioprine is associated with bone marrow suppression and anemia (Chan, Canafax, and Johnson 1987), Enalapril–Azathioprine may pose similar risks.

In summary, these results demonstrate the model’s ability to extract tailored subgraphs that capture the distinct pharmacological and mechanistic relationships underlying different drug combinations.

Conclusion

In this study, we propose RISE-DDI, a novel framework for DDI prediction that combines an RL-based subgraph extractor with a structure-aware feature integration mechanism. By leveraging RL, RISE-DDI adaptively explores the subgraph space to extract the most informative substructures around drug pairs. Extensive experiments on benchmark datasets show that RISE-DDI consistently outperforms other baselines in both transductive and inductive settings.

Acknowledgments

This study has been supported by the Guangdong S&T Program [2024B0101040005], Shenzhen Medical Research Fund [C2403001], the China Postdoctoral Science Foundation [2025M771540, GZB20250391], the Guangdong Basic and Applied Basic Research Foundation [2025A1515060011], and the Lingang Laboratory [LGL-8888].

References

- Abdullahi, T.; Gemou, I.; Nayak, N. V.; Murtaza, G.; Bach, S. H.; Eickhoff, C.; and Singh, R. 2025. K-paths: Reasoning over graph paths for drug repurposing and drug interaction prediction. *arXiv preprint arXiv:2502.13344*.
- Balakumar, P.; and Jagadeesh, G. 2015. Drugs targeting RAAS in the treatment of hypertension and other cardiovascular diseases. In *Pathophysiology and Pharmacotherapy of Cardiovascular Disease*, 751–806. Springer.
- Chan, G. L.; Canafax, D. M.; and Johnson, C. A. 1987. The therapeutic use of azathioprine in renal transplantation. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy*, 7(5): 165–177.
- Du, H.; Yao, Q.; Zhang, J.; Liu, Y.; and Wang, Z. 2024. Customized subgraph selection and encoding for drug–drug interaction prediction. *Advances in Neural Information Processing Systems*, 37: 109582–109608.
- Erre, G.; Passiu, G.; et al. 2009. Antioxidant effect of Iloprost: current knowledge and therapeutic implications for systemic sclerosis. *Reumatismo*, 61(2): 90–97.
- Himmelstein, D. S.; and Baranzini, S. E. 2015. Heterogeneous network edge prediction: a data integration approach to prioritize disease-associated genes. *PLoS computational biology*, 11(7): e1004259.
- Himmelstein, D. S.; Lizee, A.; Hessler, C.; Brueggeman, L.; Chen, S. L.; Hadley, D.; Green, A.; Khankhanian, P.; and Baranzini, S. E. 2017. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *elife*, 6: e26726.
- Hong, Y.; Luo, P.; Jin, S.; and Liu, X. 2022. LaGAT: link-aware graph attention network for drug–drug interaction prediction. *Bioinformatics*, 38(24): 5406–5412.
- Investigators, O. 2008. Telmisartan, ramipril, or both in patients at high risk for vascular events. *New England Journal of Medicine*, 358(15): 1547–1559.
- Kanehisa, M.; Furumichi, M.; Tanabe, M.; Sato, Y.; and Morishima, K. 2017. KEGG: new perspectives on genomes, pathways, diseases and drugs. *Nucleic acids research*, 45(D1): D353–D361.
- Karim, M. R.; Cochez, M.; Jares, J. B.; Uddin, M.; Beyan, O.; and Decker, S. 2019. Drug–drug interaction prediction based on knowledge graph embeddings and convolutional-LSTM network. In *Proceedings of the 10th ACM international conference on bioinformatics, computational biology and health informatics*, 113–123.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Landrum, G.; et al. 2013. RDKit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum*, 8(31.10): 5281.
- Lin, X.; Quan, Z.; Wang, Z.-J.; Ma, T.; and Zeng, X. 2020. KGNN: Knowledge graph neural network for drug–drug interaction prediction. In *IJCAI*, volume 380, 2739–2745.
- Percha, B.; and Altman, R. B. 2013. Informatics confronts drug–drug interactions. *Trends in pharmacological sciences*, 34(3): 178–184.
- Rao, J.; Lin, H.; Xie, J.; Wang, Z.; Zheng, S.; and Yang, Y. 2025a. Incorporating Retrieval-based Causal Learning with Information Bottlenecks for Interpretable Molecular Graph Learning. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 2398–2409.
- Rao, J.; Xie, J.; Yuan, Q.; Liu, D.; Wang, Z.; Lu, Y.; Zheng, S.; and Yang, Y. 2024. A variational expectation-maximization framework for balanced multi-scale learning of protein and drug interactions. *Nature Communications*, 15(1): 4476.
- Rao, J.; Xu, D.; Wei, W.; Chen, Y.; Yang, M.; and Yang, Y. 2025b. Quadruple Attention in Many-body Systems for Accurate Molecular Property Predictions. In *Forty-second International Conference on Machine Learning*.
- Rao, J.; Zheng, S.; Mai, S.; and Yang, Y. 2022. Communicative Subgraph Representation Learning for Multi-Relational Inductive Drug–Gene Interaction Prediction. In *Thirty-First International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence.
- Shomer, H.; Ma, Y.; Mao, H.; Li, J.; Wu, B.; and Tang, J. 2024. Lpformer: An adaptive graph transformer for link prediction. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, 2686–2698.
- Su, X.; Hu, L.; You, Z.; Hu, P.; and Zhao, B. 2022a. Attention-based knowledge graph representation learning for predicting drug–drug interactions. *Briefings in bioinformatics*, 23(3): bbac140.
- Su, X.; Hu, P.; You, Z.-H.; Yu, P. S.; and Hu, L. 2024. Dual-channel learning framework for drug–drug interaction prediction via relation-aware heterogeneous graph transformer. In *Proceedings of the AAAI conference on artificial intelligence*, 249–256.
- Su, X.; You, Z.; Huang, D.; Wang, L.; Wong, L.; Ji, B.; and Zhao, B. 2022b. Biomedical knowledge graph embedding with capsule network for multi-label drug–drug interaction prediction. *IEEE Transactions on Knowledge and Data Engineering*, 35(6): 5640–5651.
- Te Riet, L.; van Esch, J. H.; Roks, A. J.; van den Meiracker, A. H.; and Danser, A. J. 2015. Hypertension: renin–angiotensin–aldosterone system alterations. *Circulation research*, 116(6): 960–975.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Vilar, S.; Uriarte, E.; Santana, L.; Lorberbaum, T.; Hripsak, G.; Friedman, C.; and Tatonetti, N. P. 2014. Similarity-based

- modeling in large-scale prediction of drug-drug interactions. *Nature protocols*, 9(9): 2147–2163.
- Wang, Y.; Yang, Z.; and Yao, Q. 2024. Accurate and interpretable drug-drug interaction prediction enabled by knowledge subgraph learning. *Communications Medicine*, 4(1): 59.
- Wishart, D. S.; Knox, C.; Guo, A. C.; Cheng, D.; Shrivastava, S.; Tzur, D.; Gautam, B.; and Hassanali, M. 2008. DrugBank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic acids research*, 36(suppl_1): D901–D906.
- Xie, J.; Rao, J.; Xie, J.; Zhao, H.; and Yang, Y. 2024a. Predicting disease-gene associations through self-supervised mutual infomax graph convolution network. *Computers in Biology and Medicine*, 170: 108048.
- Xie, J.; Wang, Y.; Rao, J.; Zheng, S.; and Yang, Y. 2024b. Self-supervised contrastive molecular representation learning with a chemical synthesis knowledge graph. *Journal of Chemical Information and Modeling*, 64(6): 1945–1954.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yu, Y.; Huang, K.; Zhang, C.; Glass, L. M.; Sun, J.; and Xiao, C. 2021. SumGNN: multi-typed drug interaction prediction via efficient knowledge graph summarization. *Bioinformatics*, 37(18): 2988–2995.
- Zheng, S.; Rao, J.; Song, Y.; Zhang, J.; Xiao, X.; Fang, E. F.; Yang, Y.; and Niu, Z. 2021. PharmKG: a dedicated knowledge graph benchmark for biomedical data mining. *Briefings in bioinformatics*, 22(4): bbaa344.