

Communicative Subgraph Representation Learning for Multi-Relational Inductive Drug-Gene Interaction Prediction

Jiahua Rao^{1,†}, Shuangjia Zheng^{1,4,*}, Sijie Mai² and Yuedong Yang^{1,3,*}

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

²School of Electronic and Information Technology, Sun Yat-sen University

³Key Laboratory of Machine Intelligence and Advanced Computing, Sun Yat-sen University

⁴Galixir Technologies Ltd, Beijing

{raojh6,zhengshj9,maisj}@mail2.sysu.edu.cn, yangyd25@mail.sysu.edu.cn

Abstract

Illuminating the interconnections between drugs and genes is an important topic in drug development and precision medicine. Currently, computational predictions of drug-gene interactions mainly focus on the binding interactions without considering other relation types like agonist, antagonist, etc. In addition, existing methods either heavily rely on high-quality domain features or are intrinsically transductive, which limits the capacity of models to generalize to drugs/genes that lack external information or are unseen during the training process. To address these problems, we propose a novel **Communicative Subgraph** representation learning for **Multi-relational Inductive drug-Gene** interactions prediction (CoSMIG), where the predictions of drug-gene relations are made through subgraph patterns, and thus are naturally inductive for unseen drugs/genes without retraining or utilizing external domain features. Moreover, the model strengthened the relations on the drug-gene graph through a communicative message passing mechanism. To evaluate our method, we compiled two new benchmark datasets from DrugBank and DGIdb. The comprehensive experiments on the two datasets showed that our method outperformed state-of-the-art baselines in the transductive scenarios and achieved superior performance in the inductive ones. Further experimental analysis including LINCS experimental validation and literature verification also demonstrated the value of our model.

1 Introduction

Illuminating the interactions between drugs and genes is crucial in the drug discovery or repositioning process [Pritchard *et al.*, 2017], aiming to discover effective therapeutic drugs or identify new druggable targets [Strittmatter, 2014] and “off-targets” that may cause adverse drug reactions [Malki and Pearson, 2020]. Although experimental studies remain

the most reliable approach for determining drug-gene interactions (DGIs), the characterization of all possible DGIs is a daunting task due to the vast costs involved in experiments [Stachel *et al.*, 2014]. Hence, computational prediction has emerged as an attractive alternative to discover new drug-gene interactions.

Currently, computational predictions of DGIs mainly focus on the binding interaction between drugs and genes to discover drugs with a high binding affinity toward their target [Tsubaki *et al.*, 2019; Zheng *et al.*, 2020; Bagherian *et al.*, 2021]. However, these methods are unable to illuminate the multiple relational information (i.e., agonist, antagonist, etc.), which is also important for revealing the mechanism behind the interplay between drugs and genes [Cotto *et al.*, 2018]. For example, an agonist is a drug that activates the receptor to produce a biological response, while an antagonist is a drug that blocks agonist-mediated responses upon binding to a target gene. Therefore, it is critical to determine the multi-relational interactions between drugs and genes.

In addition, the existing methods fell roughly into two categories: feature-based and network-based methods. The feature-based methods [Öztürk *et al.*, 2018; Torng and Altman, 2019] mainly build the embedding function (i.e. graph representation) with the manually crafted features but often lead to inferior performance when high-quality domain features are not available. Furthermore, although much effort has been devoted to developing the network-based methods [Luo *et al.*, 2017; Zeng *et al.*, 2019] without domain features, these methods are intrinsically transductive, where the learned latent features cannot generalize to drugs/genes unseen during the training process. To make inductive predictions for unseen entities, one common scenario in the recommendation system studies, scientists utilize subgraphs around the entity pairs without leveraging any global information specific to the whole graph [Zhang and Chen, 2020; Teru *et al.*, 2020]. Such local subgraphs contain rich graph pattern information, enabling accurate predictions of target interactions on both the transductive and inductive scenarios.

In this study, we propose a novel **Communicative Subgraph** representation learning for **Multi-relational Inductive drug-Gene** interactions prediction (CoSMIG) to address the above problems. We firstly compiled two new challenging benchmarks for multi-relational DGI prediction. To the best

*Corresponding authors.

†Work done during an internship at Galixir.

of our knowledge, this is the first benchmark of this kind. Furthermore, our model is naturally inductive because we extracted a subgraph by the random walk with restart algorithm for each training interaction without leveraging any domain and global information specific to the drug-gene graph. Therefore, our model could be applied to unseen drugs/genes without retraining. And we also proposed a communicative message passing mechanism to strengthen the role of multi-relational information on the drug-gene graph for accurate multi-relational DGI predictions. Finally, the proposed model is shown to outperform state-of-the-art approaches over two public DGI datasets on both transductive and inductive settings. An independent test on the LINCS L1000 dataset indicated that our predicted increases/decreases of gene expression caused by drugs are highly consistent with the experimental values. More importantly, our method was proven to be able to identify new drug-gene interactions that are not in the current datasets but have been reported in the literature, demonstrating that our model may provide new insights into the understanding of the mechanisms of DGIs.

In brief, the main contributions are listed below:

- We proposed a novel inductive subgraph representation learning framework without leveraging any domain and global information of the given drug-gene graph, which could be applied to unseen drugs/genes without retraining.
- We made accurate multi-relational DGI predictions by strengthening multi-relational information through our communicative message passing mechanism.
- We compiled two new challenging benchmarks for multi-relational DGI prediction and conducted extensive experiments to demonstrate the effectiveness of CoSMIG and its interpretability in understanding the multiple relational information.

2 Related Work

In this section, we introduced the related works in two areas. First, we focus on reviewing the computational methods for drug-gene binary interaction prediction. Second, we summarized several state-of-the-art recommendation methods that have been adopted as baselines for multiple relational DGIs prediction in our work.

2.1 Drug-Gene Interactions

The key issue with drug-gene interactions is to learn a low-dimensional representation of drugs and genes. Once learned, representations can then be utilized to predict the probability scores of drug-gene binary interactions. Generally, these existing methods can be split into the following two categories: feature-based and network-based methods.

The key idea of feature-based models is to adopt a popular deep learning architecture to represent drugs and proteins/targets. For example, DeepDTA [Öztürk *et al.*, 2018] employed a convolution neural network (CNN) to learn representations from the raw protein sequences and SMILES strings and combine them to predict the binding affinities of drug-target interactions, while Graph-CNN [Torng and Altman, 2019] extracts features from pocket graphs and 2D ligand graphs through the graph convolutional framework. As

their input was based on the manually crafted features, they often lead to inferior performance when high-quality domain features are not available.

Another research line is to formulate DGI prediction as a link prediction task within the drug-gene graph [Luo *et al.*, 2017; Zeng *et al.*, 2019]. They often perform the network diffusion algorithms (i.e. random walk) to obtain informative and low-dimensional vector representations of drugs and genes in the network. After that, they infer new interactions of drugs and genes based on the known drug-gene pairs. However, they are transductive because the trained model could not be generalized to the unseen drugs/genes during the training process.

However, these methods still focused on discovering the binary interaction of drugs and proteins/targets, such as the high binding affinity of drugs and their targets, and fail to determine the interaction types between drugs and genes. In fact, an increasing number of people discovered that the interaction type between drugs and human genes has become an important subject in the research on drug efficacy and human genomics [Cotto *et al.*, 2018]. Therefore, illuminating the interconnections between drugs and genes and determining their relationship types could provide new insights into drug repositioning.

2.2 Recommendation Methods

Taking inspiration from recent progress in recommendation systems, we could re-formulate DGI prediction as a multi-relational recommendation task with drugs and genes respectively representing users and items.

Matrix completion is one of the widely-used recommender methods, where a matrix with rows and columns respectively representing users and items is used to predict users' interest in items for filling the missing entries of the rating matrix. However, it is transductive so it often requires retraining when adding a new drug/gene unseen before. To alleviate the inductive problem, Inductive Matrix Completion (IMC) [Jain and Dhillon, 2013] has been proposed, which leverages the domain features of users and items. But IMC method still has strong constraints on the quality of the domain features.

Due to the superior graph representation learning abilities, the recommendation tasks have been studied using Graph neural networks (GNNs), showing their superiority and effectiveness [Berg *et al.*, 2017]. For example, [Berg *et al.*, 2017] proposed the graph convolutional matrix completion (GC-MC) which directly applies a GNN to the user-item bipartite graph for extracting user and item latent features. But these kinds of model are still transductive.

A recent inductive graph-based recommender system, PinSage [Ying *et al.*, 2018a], uses content as initial node features and has successfully recommended related pins on Pinterest. Furthermore, IGMC [Zhang and Chen, 2020] proposed a graph-level GNN to learn representations for subgraphs and use the subgraph embedding to predict the ratings between users and items. However, these GNN models are limited by the node-to-node message passing mechanism, which cannot explicitly consider the relation embeddings in the embedding propagation layer. This set of methods, together with traditional MF-based methods constitute our baselines.

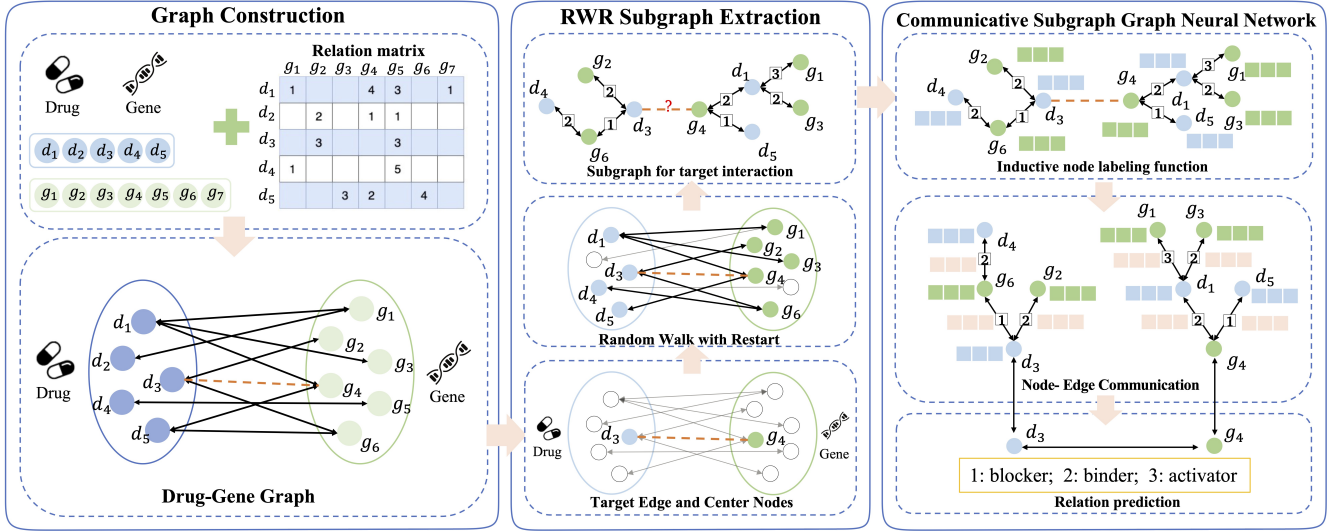


Figure 1: The framework of CoSMIG. We first extract a subgraph around each interaction and train a communicative subgraph graph neural network to map subgraphs to interactions. Each subgraph is induced by the drug and gene associated with the target interaction as well as their h -hop neighbors (here $h = 3$). Finally, the learned subgraph embedding of each interaction is used to predict the various interactions between drug and gene.

3 Methods

As shown in Figure 1, we predict DGIs through CoSMIG consisting of three main components: 1) subgraph extraction layer, which extracts the h -hop subgraph by random walk with restart (RWR); 2) embedding layer, which encodes the h -hop subgraph through a graph-level GNN; and 3) prediction layer, which predicts the interactions between two subgraph.

3.1 RWR Subgraph Extraction

For one given pair of gene g and drug d , we first constructed a subgraph by extracting h -hop neighbors around the pair (d, g) . Since our model represents DGI through a bipartite graph, h has to be an odd number for drug-gene interaction and h is usually greater than 2 for accurate prediction. Therefore, full extraction of all h -hop neighbors may cause intractably large subgraphs, and we extract subgraphs through the random walk with restart (RWR) instead of the breadth-first search by previous works [Zhang and Chen, 2020]. More specifically, given an interaction between drug and gene, we conduct random walk with restart from the two nodes around the interaction and have to find a solution to the equation:

$$p = cAD^{-1}p + (1 - c)e \quad (1)$$

where c is a number in the range $(0, 1)$ called restart probability and p is a column vector with p_i denoting the probability at node i . D is the degree matrix of adjacency matrix A with each diagonal value $D_{ii} = \sum_j A_{ij}$. The restart probability controls whether the next walk is jumping to a randomly selected neighbor (with probability c) or going to the starting node (with probability $1 - c$). For the starting vector e , we set $e_i = 1$ if node i is the starting node else 0, and thus the starting vector e allows us to preserve the node's local topological structure and AD^{-1} allows us to further visit their neighborhoods.

3.2 Communicative Subgraph Neural Network

After extracting h -hop subgraph around the pair (d, g) , we designed a communicative subgraph neural network, inspired by previous works [Song et al., 2020; Mai et al., 2021], to model the inductive subgraph by iteratively communicating the relation and node embeddings and strengthening the role of multi-relational information.

Given an extracted subgraph \mathcal{G} with n nodes and e edges, we represent N for the node embeddings, E for the edge embeddings, and R for the relation embeddings. We also defined the adjacency matrix in the subgraph as A^{ne} and A^{re} which represents the node-to-edge and relation type-to-edge adjacency matrix, respectively.

Before feeding to CoSMIG, we first apply an inductive node labeling function to it, which uses different labels to mark nodes' different roles in the subgraph without leveraging any external domain features and global information. Our node labeling function is defined as $(2i + j)$ where i is the hop number of the node and j is the node type with 0 representing the drug nodes and 1 representing the gene nodes. The one-hot encoding of these node labels will be treated as the initial node features of the subgraph, denoted as N_0 . The initial relational feature R_0 is the one-hot encoding of interaction types and the initial edge features E_0 is the initial relation feature of each edge.

Therefore, we firstly mapped the node and edge representations to the same dimensionality f :

$$N^{(0)} = \sigma(N_0 W_n^{(0)}), E^{(0)} = \sigma(E_0 W_e^{(0)}) \quad (2)$$

where σ denotes the nonlinear activation function, $W_n^{(0)}$ and $W_e^{(0)}$ are the learnable parametric matrix, $N^{(0)} \in R^{n \times f}$ and $E^{(0)} \in R^{e \times f}$ are the transformed node and relation embeddings, respectively.

Node embedding aggregation. In our node-edge interaction mechanism, relation embedding is required for updating the node embedding at each iteration. Firstly, we use the node embedding (N_i, N_j) and the relation embedding ($E_{i,j}$) to calculate the weight for the associations i, j through edge attention scores by:

$$\alpha_{i,j}^{(k)} = \sigma(\sigma([N_i^{(k)} || N_j^{(k)} || E_{i,j}^{(k)}] W_{a_0}^{(k-1)})) W_{a_1}^{(k)} \quad (3)$$

where $W_{a_0}^{(k)}$ and $W_{a_1}^{(k)}$ denotes the learnable parametric matrix at the iteration k and the operator $||$ denotes feature concatenation.

Then we use the attentive edge embedding $E_{i,j}^{(k-1)}$ to update the node representation:

$$E_{i,j}^{(k-1)} = \alpha_{i,j}^{(k-1)} E_{i,j}^{(k-1)} \quad (4)$$

$$N_{agg}^{(k)} = A^{ne} E^{(k-1)} \quad (5)$$

$$N^{(k)} = \sigma((N_{agg}^{(k)} + N^{(k-1)}) W_n^{(k)}) \quad (6)$$

where $N_{agg}^{(k)}$ denotes the node aggregation information from its neighbors at the iteration k , $W_n^{(k)}$ represents the parametric matrix for node embedding at the iteration k , $E_{i,j}^{(k-1)}$ is the attentive edge embedding for edge (i, j) . Therefore, By using the relation embedding to update node embedding, the node embedding can aggregate all the relations along with their relative information in the subgraph, providing powerful relational inference ability.

Relation embedding updating. The relation embedding is updated for totally $l - 1$ iterations and the node embedding is required to update the relation embedding. We first aggregate information from node to edge and relation to edge by:

$$E_{agg}^{(k)} = (A^{ne})^T N^{(k)} + (A^{re})^T R^{(k)} \quad (7)$$

where T denotes matrix transpose, $(A^{ne})^T N^{(k)}$ aggregates the head information to edge, and $(A^{re})^T R^{(k)}$ aggregates the relation information to edge.

Then we use the aggregation information to update edge representation:

$$E^{(k)'} = \sigma(E^{(k-1)} + \sigma(E_{agg}^{(k)})) \quad (8)$$

$$E^{(k)} = \sigma(E^{(k)'} W_e^{(k)} + E^{(0)}) \quad (9)$$

where σ denotes nonlinear activation function, $W_e^{(k)}$ denotes the learnable parametric matrix, and $E^{(0)}$ is the original transformed relation embedding to perform residual learning.

3.3 Relation Prediction

In this section, we pool the node representations into a graph-level feature vector to represent the interaction (d, g) . There are many choices such as summing, averaging, SortPooling [Zhang *et al.*, 2018], DiffPooling [Ying *et al.*, 2018b]. In this work, we only concatenate the final representations of the target drug d and gene g as the graph representation:

Dataset	DrugBank	DGIdb
Number of Drug	425	1185
Number of Gene	11284	1664
Interactions	80924	11366
Interaction types	2	14

Table 1: Statistics of two Drug-Gene Interaction datasets

$$h = \text{concat}(N_d, N_g) \quad (10)$$

where N_d and N_g denote the final representations of the target drug and gene, respectively. We empirically verified its superior performance because highlighting the target drug and gene is important for DGIs predictions.

After getting the final graph representation, we use an MLP to output the predicted interactions:

$$\hat{r} = w_y^T \sigma(W_h h) \quad (11)$$

where σ is an activation function, and w_y^T and W_h are parameters of the MLP for mapping the graph representation h to a scalar prediction \hat{r} .

3.4 Model Training

Following [Zhang and Chen, 2020], we minimize the mean squared error (MSE) between the predictions and the ground truth interactions:

$$\mathcal{L} = \left(\frac{1}{|(d, g)| \Omega_{d, g} = 1} \right) \sum_{(d, g): \Omega_{d, g} = 1} (r_{d, g} - \hat{r}_{d, g})^2 \quad (12)$$

where $r_{d, g}$ and $\hat{r}_{d, g}$ denote the true interaction and predicted interaction of (u, v) , respectively, and Ω is a 0/1 mask matrix indicating the observed entries of the interaction matrix R .

4 Experiments

4.1 Experimental Setup

Datasets. To evaluate the effectiveness of CoSMIG, we compiled the multi-relational datasets from DGIdb [Cotto *et al.*, 2018] and DrugBank [Wishart *et al.*, 2018], respectively (Table 1). (1) DrugBank: This is the pharmacotranscriptomics dataset obtained from DrugBank. Wherein, the up/down-regulation of genes due to the metabolism of pharmaceutical compounds represents the interactions between drugs and genes. Therefore, this dataset has two types of interactions, increased and decreased. (2) DGIdb: This dataset is adopted from DGIdb, containing over 1664 genes and 1185 drugs involved in over 11,366 drug-gene interactions and 14 types of relationship. To ensure the quality of the dataset, we dropped duplicates and retained the drugs and genes with at least five interactions on the two datasets.

Experimental settings. We implemented CoSMIG using pytorch geometric [Fey and Lenssen, 2019], which is available at <https://github.com/biomed-AI/CoSMIG>. We tuned model hyperparameters based on cross validation results on DrugBank and used them across all datasets. The hop number h was set to 3. The depth of model was set to 4. For each

	Methods	Features	DrugBank		DGIdb	
			Validation ACC	Ind. Test ACC	Validation ACC	Ind. Test ACC
MF-based	MC	no	-	0.518 ± 0.013	-	0.559 ± 0.009
	GRALS	yes	-	0.532 ± 0.021	-	0.578 ± 0.016
	F-EAE	no	-	0.566 ± 0.004	-	0.623 ± 0.003
GNN-based	GC-MC	yes	-	0.586 ± 0.008	-	0.601 ± 0.005
	sRGCNN	yes	-	0.602 ± 0.010	-	0.689 ± 0.007
	PinSage	yes	-	0.629 ± 0.004	-	0.713 ± 0.005
	IGMC	no	-	0.634 ± 0.003	-	0.803 ± 0.006
Proposed	CoSMIG-w/GCN	no	0.562 ± 0.004	0.581 ± 0.004	0.778 ± 0.023	0.803 ± 0.009
	CoSMIG-w/GraphSAGE	no	0.584 ± 0.003	0.602 ± 0.008	0.807 ± 0.014	0.814 ± 0.010
	CoSMIG-w/RGCN	no	0.614 ± 0.004	0.637 ± 0.005	0.821 ± 0.013	0.832 ± 0.002
	CoSMIG-w/AvgPooling	no	0.619 ± 0.003	0.643 ± 0.006	0.822 ± 0.006	0.835 ± 0.003
	CoSMIG-w/SumPooling	no	0.625 ± 0.004	0.655 ± 0.003	0.824 ± 0.007	0.839 ± 0.004
	CoSMIG-w/SortPooling	no	0.639 ± 0.002	0.667 ± 0.004	0.833 ± 0.003	0.841 ± 0.005
	CoSMIG	no	0.658 ± 0.008	0.678 ± 0.003	0.840 ± 0.011	0.852 ± 0.012

Table 2: The comparison of different methods by the overall accuracy on the DrugBank and DGIdb datasets in transductive scenario.

subgraph, we randomly dropped out its adjacency matrix entries with a probability of 0.1 during the training. The training process lasted 80 epochs on a Nvidia GeForce RTX 3090 GPU. The impacts of several hyperparameters of our method such as embedding size and learning rate have been shown in Appendix Figure S1. As a multi-relational link prediction problem, we calculate the accuracy score (ACC) between the actual drug-gene interactions and the predictions to evaluate our method. We trained the model for five times and averaged the results to obtain the final performance.

4.2 Performance Comparison

Performance in transductive scenario. The compared methods included three MF-based (MC, GRALS, F-EAE) and five GNN-based (GC-MC, sRGCNN, PinSage, IGMC) methods. As shown in Table 2, CoSMIG achieved the highest accuracies on both the DrugBank and DGIdb. The poor performance of MF-based methods (MC, GRALS, F-EAE) may be ascribed to the fact that the simple low rank (or sparse) matrix approximation is insufficient to capture the complex interaction types, whereas the performance of GC-MC and sRGCNN verify that the node-level of graph neural network can improve the representation learning and benefit the prediction of drug-gene interactions. IGMC, instead of learning transductive node-level features, learns local graph patterns related to the interactions inductively based on relational graph neural network (R-GCN), showing highly competitive performance compared to the other baselines. Nevertheless, strengthening the role of interaction embeddings in the subgraph modeling, CoSMIG improves over IGMC by 6.49% and 5.75% on the DrugBank and DGIdb, respectively. Furthermore, CoSMIG continues to outperform other baseline methods when evaluating on each interaction type (Fig. 3A).

Performance in inductive scenario. Note that GRALS and sRGCNN could not be applied to the inductive scenario due to their model settings. As shown in Table 3, CoSMIG consis-

Methods		DrugBank		DGIdb	
		ACC	Δ	ACC	Δ
MF-based	IMC	0.441	14.8%	0.424	24.2%
	F-EAE	0.474	16.3%	0.532	14.6%
GNN-based	GC-MC	0.513	12.4%	0.553	7.99%
	PinSage	0.567	9.86%	0.654	8.27%
	IGMC	0.612	3.47%	0.778	3.11%
Proposed	CoSMIG	0.672	0.88%	0.842	1.17%

 Table 3: Performance on the inductive scenario. Δ represents the decline rate between the transductive scenario and the inductive scenario of each model.

tently outperformed other baselines on the inductive datasets. The MF-based methods still achieved the poorest performances in inductive scenario and there is a huge performance gap between the transductive and inductive scenario. Moreover, IGC-MC and PinSage also show a significant decrease (the rate of decline ranging from 7.99% to 12.4%) in the inductive scenario, likely due to their only node-level message passing mechanism. Compared to the node-level GNN methods, IGMC decreases slightly because their extracted enclosing subgraph contains rich graph pattern information without any global information in the drug-gene graph. Nevertheless, CoSMIG achieves consistent outperforms IGMC by 8.93% and 7.60% on the DrugBank and DGIdb, respectively. And the decline rate of CoSMIG is significantly lower than other methods. The remarkable improvement can be attributed to its RWR subgraph extraction methods and communicative message passing mechanism.

4.3 Ablation Study

We conducted ablation studies on two benchmark datasets to investigate factors that influence our performance. As shown in Table 2, when using different types of GNN aggregators,




Figure 2: Evaluating the effect of Subgraph Extraction on DrugBank and DGIdb.

all variants were inferior to all CoSMIG. The best baseline, CoSMIG-RGCN, obtained accuracy score of 0.821 and 0.832 for the validation and test set on DGIdb, which are 2.26% and 2.35% lower than our proposed model. For three pooling layer variants, CoSMIG-w/SortPooling performs the best, while its accuracy is 0.83% and 1.29% lower than CoSMIG on the validation set and test-set of DGIdb, respectively. The consistent results between the validation and independent test sets also indicated the robustness of the CoSMIG model.

To indicate the influence of subgraph extraction, we also compared our RWR subgraph extraction with the enclosing subgraph extraction method proposed in IGMC. And we varied the hop number of neighbors to investigate whether CoSMIG can benefit from multiple hop neighbors. The results are shown in Figure 2. We observe that the performance of RWR subgraph extraction method consistently outperforms the enclosing subgraph method both on the DrugBank and DGIdb datasets. And as the hop number increases, the Accuracy of CoSMIG increases rapidly due to the more informative neighbors and interactions in the subgraph, and the model reaches its best performance at the hop number of 3. Subsequently, the performance slowly descends as the hop number continues to increase, indicating that an overlarge hop number will bring redundant and harmful information.

4.4 Experimental Analysis

Visualization analysis. We extracted the learned representations of drug-gene interaction pairs from the trained CoSMIG model, and project them into 2D space using tSNE as shown in Figure 3B-C. Obviously, CoSMIG could clearly distinguish each interaction type of DGIs, indicating that CoSMIG allows accurate representations for DGIs predictions.

Validation on the LINCS L1000 database. We further performed large-scale computation of up/down regulating genes and evaluated the predictions through the LINCS L1000 database. As detailed in Appendix Section 2, these results highlight the superiority of CoSMIG and our predicted increase/decrease of gene expression caused by drugs are highly consistent with the experimental values.

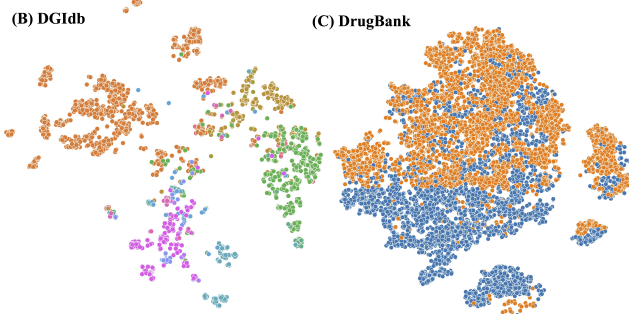
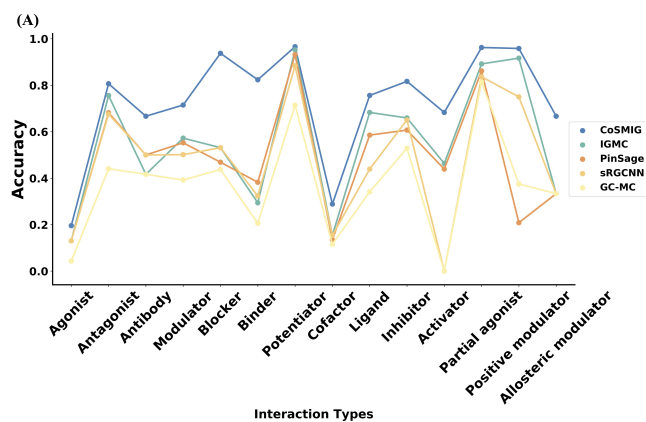


Figure 3: (A) The accuracy of each interaction type for DGIdb, and the pairs projected by t-SNE for (B) DGIdb and (C) DrugBank.

Discovering novel drug-gene interactions. We also demonstrated the new potential interactions that are not in the current datasets but have been reported in the literature as detailed in Appendix Section 3.

5 Conclusion

In this work, we firstly compiled two new challenging benchmarks for multi-relational drug-gene interactions prediction instead of focusing solely on the binding interaction. Furthermore, our model is naturally inductive based on the communicative subgraph neural network without leveraging any domain knowledge and global information specific to the drug-gene graph. Therefore, our model could be applied to unseen drugs/genes without retraining. Finally, we conducted extensive experiments to demonstrate the effectiveness of our model and its interpretability in understanding the multiple relational information. To demonstrate the value of our model, we evaluated the interactions types of increased and decreased between drugs and genes with the gene expression profile from LINCS L1000 database. Moreover, we demonstrated the new potential interactions that are not in the current datasets but have been reported in the literature.

Acknowledgments

This study has been supported by the National Key R&D Program of China [2020YFB0204803], National Natural Science Foundation of China [61772566, 62041209], Guangdong Key Field R&D Plan [2019B020228001,

2018B010109006], Introducing Innovative and Entrepreneurial Teams [2016ZT06D211], and Guangzhou S&T Research Plan [202007030010].

References

- [Bagherian *et al.*, 2021] Maryam Bagherian, Elyas Sabeti, Kai Wang, Maureen A Sartor, Zaneta Nikolovska-Coleska, and Kayvan Najarian. Machine learning approaches and databases for prediction of drug–target interaction: a survey paper. *Briefings in bioinformatics*, 22(1):247–269, 2021.
- [Berg *et al.*, 2017] Rianne van den Berg, Thomas N Kipf, and Max Welling. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- [Cotto *et al.*, 2018] Kelsy C Cotto, Alex H Wagner, Yang-Yang Feng, Susanna Kiwala, Adam C Coffman, Gregory Spies, Alex Wollam, Nicholas C Spies, Obi L Griffith, and Malachi Griffith. Dgidb 3.0: a redesign and expansion of the drug–gene interaction database. *Nucleic acids research*, 46(D1):D1068–D1073, 2018.
- [Fey and Lenssen, 2019] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [Jain and Dhillon, 2013] Prateek Jain and Inderjit S Dhillon. Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*, 2013.
- [Luo *et al.*, 2017] Yunan Luo, Xinbin Zhao, Jingtian Zhou, Jinglin Yang, Yanqing Zhang, Wenhua Kuang, Jian Peng, Ligong Chen, and Jianyang Zeng. A network integration approach for drug–target interaction prediction and computational drug repositioning from heterogeneous information. *Nature communications*, 8(1):1–13, 2017.
- [Mai *et al.*, 2021] Sijie Mai, Shuangjia Zheng, Yuedong Yang, and Haifeng Hu. Communicative message passing for inductive relation reasoning. *Association for the Advancement of Artificial Intelligence (AAAI)*, 2021.
- [Malki and Pearson, 2020] Mustafa Adnan Malki and Ewan Robert Pearson. Drug–drug–gene interactions and adverse drug reactions. *The pharmacogenomics journal*, 20(3):355–366, 2020.
- [Öztürk *et al.*, 2018] Hakime Öztürk, Arzucan Özgür, and Elif Ozkirimli. Deepdta: deep drug–target binding affinity prediction. *Bioinformatics*, 34(17):i821–i829, 2018.
- [Pritchard *et al.*, 2017] Jayne-Louise E Pritchard, Tracy A O’Mara, and Dylan M Glubb. Enhancing the promise of drug repositioning through genetics. *Frontiers in pharmacology*, 8:896, 2017.
- [Song *et al.*, 2020] Ying Song, Shuangjia Zheng, Zhangming Niu, Zhang-Hua Fu, Yutong Lu, and Yuedong Yang. Communicative representation learning on attributed molecular graphs. In *IJCAI*, volume 2020, pages 2831–2838, 2020.
- [Stachel *et al.*, 2014] Shawn J Stachel, John M Sanders, Darrell A Henze, Mike T Rudd, Hua-Poo Su, Yiwei Li, Kausik K Nanda, Melissa S Egbertson, Peter J Manley, Kristen LG Jones, et al. Maximizing diversity from a kinase screen: identification of novel and selective pan-trk inhibitors for chronic pain. *Journal of Medicinal Chemistry*, 57(13):5800–5816, 2014.
- [Strittmatter, 2014] Stephen M Strittmatter. Overcoming drug development bottlenecks with repurposing: old drugs learn new tricks. *Nature medicine*, 20(6):590–591, 2014.
- [Teru *et al.*, 2020] Komal Teru, Etienne Denis, and Will Hamilton. Inductive relation prediction by subgraph reasoning. In *International Conference on Machine Learning*, pages 9448–9457. PMLR, 2020.
- [Torng and Altman, 2019] Wen Torng and Russ B Altman. Graph convolutional neural networks for predicting drug–target interactions. *Journal of chemical information and modeling*, 59(10):4131–4149, 2019.
- [Tsubaki *et al.*, 2019] Masashi Tsubaki, Kentaro Tomii, and Jun Sese. Compound–protein interaction prediction with end-to-end learning of neural networks for graphs and sequences. *Bioinformatics*, 35(2):309–318, 2019.
- [Wishart *et al.*, 2018] David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids research*, 46(D1):D1074–D1082, 2018.
- [Ying *et al.*, 2018a] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 974–983, 2018.
- [Ying *et al.*, 2018b] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. *Advances in neural information processing systems*, 31, 2018.
- [Zeng *et al.*, 2019] Xiangxiang Zeng, Siyi Zhu, Xiangrong Liu, Yadi Zhou, Ruth Nussinov, and Feixiong Cheng. deepdr: a network-based deep learning approach to in silico drug repositioning. *Bioinformatics*, 35(24):5191–5198, 2019.
- [Zhang and Chen, 2020] Muhan Zhang and Yixin Chen. Inductive matrix completion based on graph neural networks. In *International Conference on Learning Representations*, 2020.
- [Zhang *et al.*, 2018] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [Zheng *et al.*, 2020] Shuangjia Zheng, Yongjian Li, Sheng Chen, Jun Xu, and Yuedong Yang. Predicting drug–protein interaction using quasi-visual question answering system. *Nature Machine Intelligence*, 2(2):134–140, 2020.