

# Bi-Level Graph Neural Networks for Drug-Drug Interaction Prediction

Yunsheng Bai<sup>\*1</sup> Ken Gu<sup>\*1</sup> Yizhou Sun<sup>1</sup> Wei Wang<sup>1</sup>

## 1. Bi-Level Graph Neural Networks

We introduce BI-GNN (*Bi-Level Graph Neural Networks*) for modeling biological link prediction tasks such as drug-drug interaction (DDI) (Vilar et al., 2012) and protein-protein interaction (PPI) (Keskin et al., 2008). Taking drug-drug interaction as an example, existing methods using machine learning either only utilize the link structure between drugs without using the graph representation of each drug molecule (Zitnik et al., 2018; Ma et al., 2018), or only leverage the individual drug compound structures without using graph structure for the higher-level DDI graph (Deac et al., 2019). Readers are referred to a recent survey (Sun et al., 2019) for more details.

We demonstrate our model using the drug-drug interaction prediction task as an example. Drug-drug interactions occur when the presence of one drug changes the effect of another drug producing an observed side effect (Wishart et al., 2018). Specifically, we consider the transductive setting of drug repurposing (Pushpakom et al., 2019), in which predicting interactions between existing drugs along with leveraging interaction data can help infer similar physiological effects for other existing drugs (Zhou et al., 2015). This is especially important with the recent COVID-19 pandemic, as drug repurposing techniques aim to find existing drugs that may be helpful in treading COVID-19 (Hamilton, 2020). Our framework is also extendable to other biological link prediction tasks with different interacting biological entities, e.g. proteins (Fout et al., 2017; Huan et al., 2005).

The key idea is to fundamentally view the data as a bi-level graph, in which the highest level is a representing the interaction between biological entities (interaction graph), and each biological entity itself is further expanded to its intrinsic graph representation (representation graphs), in which the graph is either flat like a drug compound (Kearnes et al., 2016) or hierarchical like a protein with amino acid level graph, secondary structure, tertiary structure, etc (Wikipedia contributors, 2020). Our BI-GNN not only allows the usage of information from both the high-level interaction graph

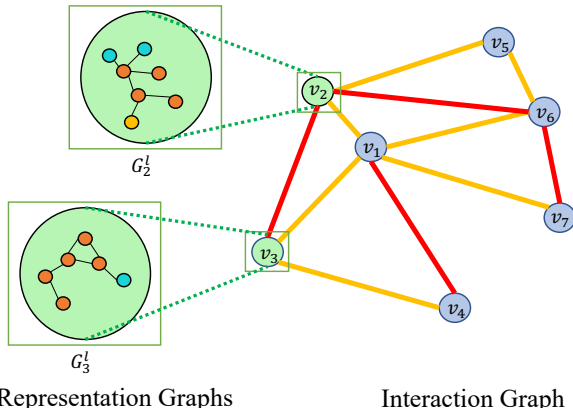


Figure 1. A bi-level graph view of the drug-drug interaction data, in which the number of levels is 2. Existing GNN methods operate only on either the representation graphs or the single interaction graph without utilizing both under the GNN framework. Node colors in the representation graphs denote molecular level element types. Edge colors in the interaction graph denote drug interactions types.

and the low-level representation graphs, but also offers a baseline for future research opportunities to address the bi-level nature biological interaction networks.

**Definitions** We denote an undirected, unweighted graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $N = |\mathcal{V}|$  nodes. In addition, we let  $\mathcal{G}^h = (\mathcal{V}^h, \mathcal{E}^h)$  denote the higher level interaction network with node set  $\mathcal{V}^h$  and edge set  $\mathcal{E}^h$ . Moreover each  $e_i^h \in \mathcal{E}^h$  can be of different types. Let  $\mathcal{R}$  denote the set of all edge types. Each  $v_i \in \mathcal{V}^h$  is also a graph which we denote as  $\mathcal{G}_i^l = (\mathcal{V}_i^l, \mathcal{E}_i^l)$ . The set of all graphs in the lower level is  $\{\mathcal{G}_1^l, \mathcal{G}_2^l, \dots, \mathcal{G}_{|\mathcal{V}^h|}^l\}$ . Fig 1 shows an example with  $|\mathcal{V}^h| = 7$ .

Node features for the  $i^{th}$  representation graph are summarized in a  $N_i \times D_l$  matrix  $\mathbf{H}_i^l$  where  $N_i = |\mathcal{V}_i^l|$ . Likewise node features for the interaction graph are denoted by a  $N \times D_h$  matrix  $\mathbf{H}^h$ . Taking DDI as an example in which the representation graphs are molecular structures of drugs,  $\mathbf{H}_i^l$  would represent atom features such as the atomic number, whether an atom is aromatic, its hybridization, the number of hydrogen atoms attached to the atom etc. of  $\mathcal{G}_i^l$ . Currently, we do not consider edge features in the lower level graph but this can be easily extended in future work.

**BI-GNN** BI-GNN consists of the following sequential

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Computer Science, University of California, Los Angeles, USA. Correspondence to: Yunsheng Bai <yba@ucla.edu>, Ken Gu <kengu13@ucla.edu>.

Table 1. Comparison of baseline methods.  $GAT^l$  runs on the representation graphs while  $GAT^h$  runs on the interaction graph

Method	Low-Level Model	High-Level Model	Feature of Drugs
FP-PRED	N/A	N/A	Molecular Structure
LL-GNN	$GAT^l \times 5$	N/A	Molecular Structure
MHCADDI (Deac et al., 2019)	$\{GAT-GMN\}^l \times 3$	N/A	Molecular Structure
DECAGON-R (Zitnik et al., 2018)	N/A	$GAT^h \times 3$	Random
DECAGON-FP (Zitnik et al., 2018)	N/A	$GAT^h \times 3$	Molecular Structure
DECAGON-OH (Zitnik et al., 2018)	N/A	$GAT^h \times 3$	Learned Transductively
BI-GNN (this paper)	$GAT^l \times 5$	$GAT^h \times 3$	Molecular Structure

stages: 1) *Lower Level Representation Graph Neural Network* generates vector representations for each representation graph; 2) *Higher Level Interaction Graph Neural Network* further propagates information from the lower level graph embeddings to neighboring nodes in the interaction graph, which provides the final graph representations to a *fully connected network* to obtain a final link prediction score. In this scenario, the representation graph may contain more than one level, one at the amino acid level and another at the secondary structure level.

### 1.1. Stage I: Lower Level Representation Graph Embedding

In Stage I, we generate the representation graphs' graph embeddings. In general this can be obtained through any graph embedding method. For example, one can use node embedding models such as Graph Convolutional Networks (Kipf & Welling, 2017), Graph Attention Networks (Veličković et al., 2018), or Graph Isomorphism Networks (Xu et al., 2019) followed by a READOUT function, a function that takes the embeddings of nodes as input and outputs a single embedding for the graph. Additionally, one can use hierarchical graph representation models such as DIFFPOOL (Ying et al., 2018) or GRAPH-U-NET (Gao & Ji, 2019).

In this work, we use Graph Attention Networks (GAT) with multi-scale READOUT of the updated node embeddings. Formally, the  $k$ -th layer of GAT is defined as:

$$\mathbf{x}_i^{(k+1)} = \text{ReLU}(\alpha_{i,i} \hat{\mathbf{x}}_i^{(k)} + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j} \hat{\mathbf{x}}_j^{(k)}) \quad (1)$$

$$\alpha_{i,j} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}^{(k)T} [\hat{\mathbf{x}}_i^{(k)} \parallel \hat{\mathbf{x}}_j^{(k)}]))}{\sum_{m \in \mathcal{N}(i)} \exp(\text{LeakyReLU}(\mathbf{a}^{(k)T} [\hat{\mathbf{x}}_i^{(k)} \parallel \hat{\mathbf{x}}_m^{(k)}]))} \quad (2)$$

Here,  $\hat{\mathbf{x}}_i$  is the transformed node embedding from initial feature:

$$\hat{\mathbf{x}}_i = \mathbf{W} \mathbf{x}_i. \quad (3)$$

$\mathcal{N}(i)$  is the set of all first-order neighbors of node  $i$  plus node  $i$  itself;  $\mathbf{W}^{(l)} \in \mathbb{R}^{D^l \times D^{l+1}}$  is the weight matrix associated with the  $n$ -th GAT layer;  $\alpha_{i,j}$  is a scalar attention

weight that node  $i$  gives to node  $j$ ;  $\mathbf{a}^{(l)} \in \mathbb{R}^{2D^l}$  is the attention weight vector; and  $\parallel$  is vector concatenation.

**Multi-Scale READOUT** Inspired by GIN (Xu et al., 2019), we concatenate the node representations across all layers of the GAT. This allows the model to consider all structural information, at various levels of granularity. For  $K$  GAT layers the representation graph embeddings  $\mathbf{x}_G$  can be expressed as:

$$\mathbf{x}_G = \text{CONCAT}(\text{READOUT}(\mathbf{x}_v^{(k)} | v \in G) | k = 1, 2, \dots, K). \quad (4)$$

### 1.2. Stage II: Higher Level Interaction Node Embedding

Graph embeddings from Stage I as initial node features  $\mathbf{H}^h \in \mathbb{R}^{N \times D_h}$  for the interaction graph. This is motivated by the intuition that the lower level network presents a useful initial representation from which the higher level network further enhances for the task. Similar to Stage I, many different node embedding methods may be used to refine the entity representation. We continue to use a different set of GAT layers for the higher level node embedding. Because only the final node representation in the interaction graph is needed, there is no multi-scale READOUT. Additionally, in the case of multiple edge types, the GAT propagation becomes

$$\mathbf{x}_i^{(t+1)} = \text{ReLU}(\sum_{r \in \mathcal{R}} (\alpha_{i,i}^r \hat{\mathbf{x}}_i^{r(t)} + \sum_{j \in \mathcal{N}^r(i)} \alpha_{i,j}^r \hat{\mathbf{x}}_j^{r(t)})) \quad (5)$$

where  $\mathcal{N}^r(i)$  is the set of neighbors of node  $i$  that have edge type  $r$ ;  $\hat{\mathbf{x}}_i^r = \mathbf{W}_r \mathbf{x}_i^r$ ; and  $\alpha_{i,j}^r$  is calculated as in Equation 2 but parameterized by  $\mathbf{a}_r$ .

For link prediction, we concatenate the representations of a pair of entities after  $T$  upper level GAT layers and feed it into fully connected layers to predict a link prediction score:

$$\text{pred}(v_i, v_j) = \text{MLP}(\mathbf{x}_i \parallel \mathbf{x}_j). \quad (6)$$

We choose this decoder over simpler alternatives such as dot product between drug embeddings, because empirically we find this yields better performance.

For multiclass link prediction (multiple edge types), the

fully connected layers output  $C$  scores for  $C$  classes. We use cross entropy loss with logits for the loss function.

## 2. Experiments

We train BI-GNN on two DDI datasets, DRUGBANK (Marinka Zitnik & Leskovec, 2018; Wishart et al., 2018) and DRUGCOMBO (Liu et al., 2020) whose details can be found in the supplementary material.

### 2.1. Baseline Methods

To fairly compare the uses of different levels of information we aim to use similar architectures as BI-GNN, summarized in Table 1 and detailed in the supplementary material.

**Representation Graph Models** For models that only use the lower level representation graph we focus on three variants. The first which we call FP-PRED, feeds extended connectivity fingerprint (ECFP) (Rogers & Hahn, 2010) representations directly into the prediction layers. ECFPs are representations of chemical structures which capture relevant molecular features and molecular structure. The second is LL-GNN which is the lower level model outlined in section 1.1. The last model is MHCADDI (Deac et al., 2019) which uses intra-graph message passing and inter-graph co-attention. We implement a similar model that uses Graph Matching Networks (GMN) (Li et al., 2019) for inter-graph attention.

### Interaction Graph Models

Under models that only use the interaction graph, we consider DECAGON (Zitnik et al., 2018), which considers protein-protein networks and protein-drug networks on top of the DDI as the input graph. We adopt DECAGON for the DDI network only. More specifically, we focus on three different initializations of  $H^h \in \mathbb{R}^{N \times D^h}$ , the input drug features: (1) one hot initialization which we call DECAGON-OH, (2) random initialization denoted as DECAGON-R, and (3) ECFP (Rogers & Hahn, 2010) initialization which we refer to as DECAGON-FP.

Table 2. Overall link prediction accuracy of all the methods on two datasets. Top two results are highlighted in bold.

Method	DRUGBANK		DRUGCOMBO	
	ROC	PR	ROC	F1
FP-PRED	0.630	0.647	0.648	0.495
LL-GNN	0.822	0.800	0.813	0.661
MHCADDI	0.864	0.844	0.824	0.681
DECAGON-R	0.898	0.896	0.853	<b>0.753</b>
DECAGON-FP	0.905	0.903	0.844	0.731
DECAGON-OH	<b>0.940</b>	<b>0.939</b>	<b>0.864</b>	0.748
BI-GNN	<b>0.933</b>	<b>0.929</b>	<b>0.859</b>	<b>0.767</b>

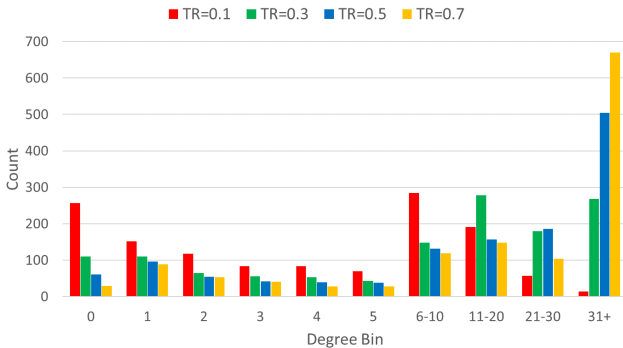


Figure 2. The distribution of node degrees across different training data ratios (percentage of DDI pairs in the training set denoted as “TR”).

### 2.2. Results

From Table 2 we observe the effectiveness of different levels of network information. For models that use only the lower level representation graphs, inter-graph communication provides better performance (LL-GNN vs MHCADDI) and utilizing GNNs allow better learning of the graph representations suit the link prediction task (LL-GNN vs FP-PRED). However, DECAGON baselines, which only use the higher level interaction network, significantly outperform the lower level models. Finally, utilizing both levels of information with BI-GNN is better than DECAGON-FP and DECAGON-R and is competitive with DECAGON-OH. Nevertheless, it must be noted that despite the strong performance by one hot initialization, it is not scalable to larger interaction networks. Specifically, because  $D^h$  grows linearly with the number of the nodes, the weight matrix becomes larger. In addition, as described in section 2.3 one hot initialization is unable to generalize representations for unseen nodes or nodes which have few neighbors.

### 2.3. When BI-GNN Performs the Best

Since the overall performance shows that BI-GNN is either the best or the second best, we conduct additional evaluation to investigate under what circumstances BI-GNN performs better than all the baseline methods. Specifically, we conjecture that the high performance achieved by DECAGON-OH is due to its ability to learn drug representations in a transductive way. From Equation 3, if the initial feature is one-hot, after multiplying with the learnable weight matrix  $W$ , each drug is essentially represented as a randomly initialized but learnable embedding. On one hand, this inherently prohibits the using of DECAGON-OH to inductive setting where a drug in the testing set can be unseen during training whose embedding is then unknown. On the other hand, however, this also allows DECAGON-OH to take advantage of the transductive nature of the drug re-purposing task and often achieve the best performance.

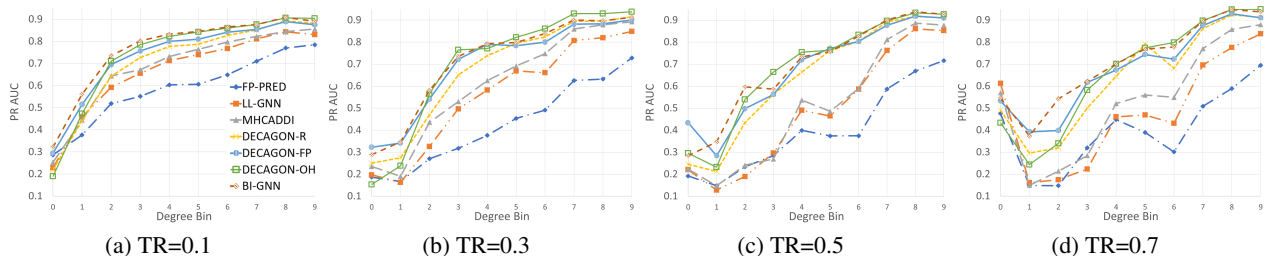


Figure 3. Performance of all methods on DRUGBANK under different training data ratios. Under each case, further breakdown of performance under different node degree splits (Figure 2) are shown.

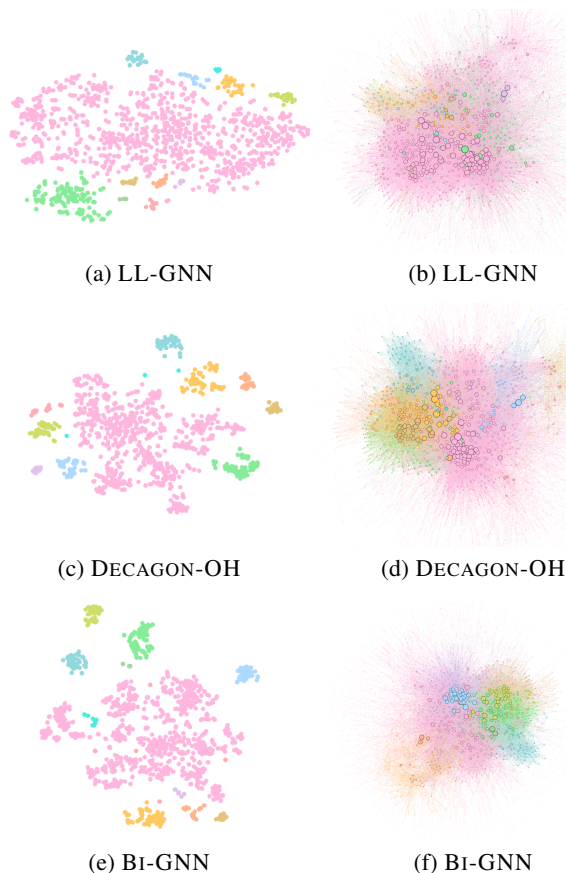


Figure 4. Left: Drug embeddings generated by different models with color indicating clustering. Right: The DRUGBANK interaction graph with node colors corresponding to the clustering and circle size indicating node degree.

What if the link structure is not available or very sparse? To DECAGON-OH, sparse neighborhood would limit its learning ability to adjust each drug representation, and lower performance would be expected. In contrast, our BI-GNN may still leverage the inherent molecular structure of each drug compound. Based on this hypothesis, we split the drugs into “bins” based on node degrees. Figure 2 shows the distribution of number of nodes in different bins, where bins with lower indices correspond to sparse regions of the DDI network. We then evaluate the performance of each

model under each bin. As shown in Figure 3, BI-GNN indeed outperforms DECAGON-OH when node degree is low. In practice, drugs may come with different amount of known interacting drugs, and our framework is especially good for drugs with little to no information available.

#### 2.4. Correlation between Drug Embeddings and Drug-Drug Interactions

To gain a better understanding of what kind of representations are useful for the link prediction task, we perform the following visualization procedure. For LL-GNN, DECAGON-OH, and BI-GNN, we take the final drug embeddings before the final prediction layer, and project them into a 2D plane using t-SNE (Maaten & Hinton, 2008). We then use DBSCAN (Ester et al., 1996) to cluster the embeddings and color the drug nodes in the interaction network according to the clustering assignment, resulting in Figure 4. It is clear that DECAGON-OH and BI-GNN learn useful drug embeddings which highly correlate to the link structure in the DDI graph, while the embeddings generated by LL-GNN which only relies on the molecular structure do **not** lead to good performance for link prediction in the interaction graph. Based on Section 2.3 and 2.4, we conclude that the message passing across the **higher** level graph contributes the most to the performance of GNN models on the DDI prediction task.

### 3. Conclusion and Future Work

We propose a multi-level GNN framework for biological entity link prediction by constructing a bi-level graph with higher level representing the interaction graph between biological entities while the lower levels representing the individual biological entity such as drug, protein, etc. In the future, we plan to extend our method to the protein-protein interaction task where each protein is represented as a hierarchical graph with amino acids, secondary structures, etc. Additionally, to improve the accuracy of our BI-GNN, we will explore the integration of graph matching and the introduction of additional input features into our framework.

## ACKNOWLEDGMENTS

This work is partially supported by NSF DGE-1829071, NSF DBI-1565137, NIH R35-HL135772, NSF III-1705169, NSF CAREER Award 1741634, NSF #1937599, Okawa Foundation Grant, and Amazon Research Award.

## References

- Deac, A., Huang, Y.-H., Veličković, P., Liò, P., and Tang, J. Drug-drug adverse effect prediction with graph co-attention, 2019.
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pp. 226–231, 1996.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Fout, A., Byrd, J., Shariat, B., and Ben-Hur, A. Protein interface prediction using graph convolutional networks. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6530–6539. Curran Associates, Inc., 2017.
- Gao, H. and Ji, S. Graph u-nets. *CoRR*, abs/1905.05178, 2019. URL <http://arxiv.org/abs/1905.05178>.
- Hamilton, W. Graph methods for covid-19 response. <https://grlplus.github.io/files/graphs-against-covid.pdf>, 2020.
- Huan, J., Bandyopadhyay, D., Wang, W., Snoeyink, J., Prins, J., and Tropsha, A. Comparing graph representations of protein structure for mining family-specific residue-based packing motifs. *Journal of Computational Biology*, 12: 657–671, 2005.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. Molecular graph convolutions: moving beyond fingerprints. *Journal of Computer-Aided Molecular Design*, pp. 1–14, 2016. URL <http://link.springer.com/article/10.1007%2Fs10822-016-9938-8>. doi:10.1007/s10822-016-9938-8 <http://dx.doi.org/10.1007/s10822-016-9938-8>.
- Keskin, O., Nussinov, R., and Gursoy, A. Prism: protein-protein interaction prediction by structural matching. In *Functional Proteomics*, pp. 505–521. Springer, 2008.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Li, Y., Gu, C., Dullien, T., Vinyals, O., and Kohli, P. Graph matching networks for learning the similarity of graph structured objects. *ICML*, 2019.
- Liu, H., Zhang, W., Zou, B., Wang, J., Deng, Y., and Deng, L. DrugCombDB: a comprehensive database of drug combinations toward the discovery of combinatorial therapy. *Nucleic Acids Res.*, 48(D1):D871–D881, 01 2020.
- Ma, T., Xiao, C., Zhou, J., and Wang, F. Drug similarity integration through attentive multi-view graph auto-encoders. In *IJCAI, IJCAI'18*, pp. 3477–3483. AAAI Press, 2018. ISBN 9780999241127.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov): 2579–2605, 2008.
- Marinka Zitnik, Rok Sosič, S. M. and Leskovec, J. BioSNAP Datasets: Stanford biomedical network dataset collection. <http://snap.stanford.edu/biodata>, August 2018.
- Pushpakom, S., Iorio, F., Eyers, P. A., Escott, K. J., Hopper, S., Wells, A., Doig, A., Guillems, T., Latimer, J., McNamee, C., et al. Drug repurposing: progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.
- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. Model.*, 50(5):742–754, 2010. doi: 10.1021/ci100050t. URL <https://doi.org/10.1021/ci100050t>.
- Sun, M., Zhao, S., Gilvary, C., Elemento, O., Zhou, J., and Wang, F. Graph convolutional networks for computational drug development and discovery. *Briefings in Bioinformatics*, 21(3):919–935, 06 2019. ISSN 1477-4054. doi: 10.1093/bib/bbz042. URL <https://doi.org/10.1093/bib/bbz042>.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- Vilar, S., Harpaz, R., Uriarte, E., Lourdes, S., Rabadan, R., and Friedman, C. Drug–drug interaction through molecular structure similarity analysis. *Journal of the American Medical Informatics Association : JAMIA*, 19: 1066–1074, 2012.

Wikipedia contributors. Protein — Wikipedia, the free encyclopedia. <https://en.wikipedia.org/w/index.php?title=Protein&oldid=960063647>, 2020. [Online; accessed 5-June-2020].

Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., Assempour, N., Iynkkaran, I., Liu, Y., Maciejewski, A., Gale, N., Wilson, A., Chin, L., Cummings, R., Le, D., Pon, A., Knox, C., and Wilson, M. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic Acids Res.*, 46(Database-Issue):D1074–D1082, 2018. doi: 10.1093/nar/gkx1037. URL <https://doi.org/10.1093/nar/gkx1037>.

Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.

Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*, 2018.

Zhou, B., Wang, R., Wu, P., and Kong, D.-X. Drug repurposing based on drug–drug interaction. *Chemical biology & drug design*, 85(2):137–144, 2015.

Zitnik, M., Agrawal, M., and Leskovec, J. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13):457–466, 2018.

## A. Dataset Description

We run experiments on 2 real word datasets on drug-drug interaction. Moreover, for each drug, we obtain its SMILE string from which we can derive the molecular structure. We one hot encode the lower level representation graph’s node features. These include the atomic number, the number of attached hydrogen atoms, the hybridization, whether it is an acceptor or donor and whether it is aromatic.

### A.1. DRUGBANK

DRUGBANK is a dataset of drug-drug interactions from the DRUGBANK database containing detailed drug data and its interaction information<sup>1</sup>. We take the pairs mined by (Marinka Zitnik & Leskovec, 2018) and keep the pairs given as DRUGBANK Ids from which we can find the drugs’ molecular structure information. After preprocessing, there are 1309 drug representation graphs and 41072 drug interactions. For training and evaluation we perform negative sampling on the positive drug pairs.

<sup>1</sup><https://www.drugbank.ca/>

### A.2. DRUGCOMBO

DRUGCOMBO is a database containing drug-drug combinations obtained from various sources including external databases, manual curations, and experimental results<sup>2</sup>. We take the pairs of drug combinations that have been classified as exhibiting synergistic or antagonistic effects. We cross-reference PubChem<sup>3</sup>, a database of chemical molecules for the molecular structures. Similar to before, we keep the pairs in which molecular structure information is found. In total, there are 3242 drug representation graphs, 34335 drug pairs classified as having synergistic effects, and 15057 pairs classified as having antogonistic effects. We also treat the different classes as edge features in the interaction graph and use it in our higher level propagation. Similar to DRUGBANK we perform negative sampling for each known effect.

## B. Parameter Settings

For each lower level and higher level GAT we have hidden dimensions of 64. The input dimensions for DECAGON-FP and DECAGON-R are 64 dimensions while the input dimension for DECAGON-OH is the number of drugs in the DDI. We use 3 higher level GAT layers for DECAGON models. For LL-GNN we use 5 lower level GAT layers while for MHCADDI we use 3 blocks of GAT followed by GMN co-attention. Experiments are ran on an Intel i7-6800K CPU and Nvidia Titan GPU. We split the dataset into training validation and testing sets based on the known pairs. For training, for all models except MHCADDI, we have a batch size of 64 graphs and select all adjacent pairs in the training set. For MHCADDI, we have a batch size of 128 pairs as each graph is dependent on its paired graph only for its own representation. Additionally, we use the Adam optimizer (Kingma & Ba, 2015) with the learning rate set to 0.001 and the the number of iterations set to 10000. We use the mean aggregator as our READOUT function. We evaluate on the validation set after every 100 iterations and employ early stopping if performance does not improve over a window size of 15. Finally, all experiments were implemented with the PyTorch and PyTorch Geometric libraries (Fey & Lenssen, 2019).

## C. Baseline Details and Related Work

### C.1. MHCADDI vs GMN

In this work, we mention the use of inter-graph attention in MHCADDI with GMNs (Li et al., 2019). For two graphs  $\mathcal{G}_1^l, \mathcal{G}_2^l$  and node  $x_i \in \mathcal{G}_1^l$  the model as described in (Deac

<sup>2</sup><http://drugcombdb.denglab.org/main>

<sup>3</sup><https://pubchem.ncbi.nlm.nih.gov/>

et al., 2019) computes the inter-graph representation as

$$\mathbf{x}_{i,inter}^{(k+1)} = \sum_{j \in \mathcal{G}_2^l} \alpha_{i,j}^{(k+1)} \mathbf{W}_1^{(k+1)} \mathbf{x}_j^{(k)} \quad (7)$$

$$\alpha_{i,j}^{(k+1)} = \frac{\exp(\langle \mathbf{W}_2^{(k+1)} \mathbf{x}_i^{(k)}, \mathbf{W}_2^{(k+1)} \mathbf{x}_j^{(k)} \rangle)}{\sum_{j' \in \mathcal{G}_2^l} \exp(\langle \mathbf{W}_2^{(k+1)} \mathbf{x}_i^{(k)}, \mathbf{W}_2^{(k+1)} \mathbf{x}_{j'}^{(k)} \rangle)}. \quad (8)$$

Here,  $\langle \cdot, \cdot \rangle$  is the dot product. In our model, we utilize a similar attention mechanism adopted from Graph Matching Networks (Li et al., 2019). Specifically we compute the inter-graph message as follows

$$\mathbf{x}_{i,inter}^{(k+1)} = \sum_{j \in \mathcal{G}_2^l} \alpha_{i,j} (\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(k)}) \quad (9)$$

$$a_{i,j} = \frac{\exp(\cos(\mathbf{x}_i^{(k)}, \mathbf{x}_j^{(k)}))}{\sum_{j \in \mathcal{G}_2^l} \exp(\cos(\mathbf{x}_i^{(k)}, \mathbf{x}_j^{(k)}))}. \quad (10)$$

In both cases, MHCADDI and GMN update the next step node representation,  $\mathbf{x}_i^{(k+1)}$ , using a combination of the representation from intra-graph message passing and inter-graph message passing.

## C.2. DECAGON vs BI-GNN

In the main text, we compare with three versions of DECAGON (Zitnik et al., 2018), but here we would like to point out that the original work of DECAGON uses additional input features for the drugs as well as a heterogeneous interaction network including both drugs and proteins, whereas our three versions use the same amount of input features as BI-GNN for fair comparison. In addition, drug-drug similarity has also been explored for DDI prediction in literature (Ma et al., 2018). In future, we plan to incorporate these additional information to further improve the prediction accuracy.