HeTAN: Heterogeneous Graph Triplet Attention Network for Drug Repurposing

Farhan Tanvir
Department of Computer Science
Georgia State University
Atlanta, Georgia, USA
ftanvir@gsu.edu

Khaled Mohammed Saifuddin
Department of Computer Science
Georgia State University
Atlanta, Georgia, USA
ksaifuddin1@student.gsu.edu

Tanvir Hossain
Department of Computer Science
Georgia State University
Atlanta, Georgia, USA
thossain5@student.gsu.edu

Arunkumar Bagavathi
Department of Computer Science
Oklahoma State University
Stillwater, Oklahoma, USA
abagava@okstate.edu

Esra Akbas
Department of Computer Science
Georgia State University
Atlanta, Georgia, USA
eakbas1@gsu.edu

***Abstract*—Modeling the interactions between drugs, targets, and diseases has significant implications for drug discovery, precision medicine and personalized treatments. Current computational approaches consider pairwise interaction, including drug-target or drug-disease interaction individually. On the other hand, within human metabolic systems, the interaction of drugs with protein targets in cells influences target activities. Moving beyond binary relationships and exploring tighter relationships together as triple is essential to understanding drugs' mechanism of action (MoAs). Moreover, considering the heterogeneity of drugs, targets, and diseases, along with their distinct characteristics, it is critical to model these complex interactions appropriately. To address these challenges, we develop a novel Heterogeneous Graph Triplet Attention Network (`HeTAN`) by modeling the interconnectedness of all entities in a heterogeneous graph. `HeTAN` introduces a novel triplet message passing and triplet-wise attention mechanism within this heterogeneous graph structure. In contrast to focusing only on pairwise attention as the importance of an entity for the other, we define triplet attention to model the importance of pairs for the other in the drug-target-disease triplet prediction problem. We perform extensive experiments on real-world datasets and our results show that `HeTAN` outperforms several baselines, demonstrating its superior performance in uncovering novel drug-target-disease relationships.**

***Index Terms*—Drug Discovery, Heterogeneous Graph Neural Network, Graph Neural Network, Representation Learning, Graph Attention, Triplet Prediction**

I. Introduction

Understanding drugs' mechanism of action (MoA) is crucial for drug repurposing, a promising approach to accelerating drug discovery and offering avenues for personalized medicine and targeted therapies. However, traditional drug discovery is time-consuming and expensive [1]. To address this challenge, computational methods have emerged as invaluable tools for leveraging large-scale chemical and genomic data [2].

Recent machine learning advancements have enhanced the study of drugs' MoAs through various learning tasks like drug behavior analysis, target activity evaluation, and disease modeling. [3]. Among these tasks, predicting the relations of drugs with other entities, such as drug-disease and drug-target prediction, have gained significant attention [4], [5]. While existing methods have made progress in predicting the relations of drugs with other entities, they often treat these tasks as isolated tasks, leading to limitations in capturing the interconnected nature of drugs with other entities. Crucially, a drug's therapeutic effect hinges on its interplay with biological targets within complex pathways and the overall metabolic system [1]. Drugs interact with protein targets in cells to modulate target activities, altering biological pathways to treat diseases. This activity integrates higher-order relationships among multiple entities. Therefore, a more comprehensive triple relationship involving drugs, targets, and diseases must be considered to capture the interplay between these entities.

Tensor factorization has emerged as a popular approach for drug-target-disease triplet prediction problems. They infer missing entries in drug-target-disease tensors via extracting latent structures from high-dimensional data [6]. NeurTN [7] combines tensor algebra and deep neural networks to learn the intrinsic relationships among drugs, targets, and diseases. However, traditional tensor models like Canonical Polyadic (CP) decomposition and Tucker decomposition suffer from issues, including linearity and data sparsity. Nonlinear tensor factorization methods have shown promise in capturing the complexities of the data, but they often rely on prior Gaussian processes that are challenging to estimate [8]. Moreover, incorporating auxiliary information into tensor models requires tedious feature engineering, making it challenging to handle large-scale healthcare data [9]. Furthermore, while many graph-based machine-learning models are common for drug-related problems, the tensor model does not utilize graph machine-learning models to predict new triplets.

Heterogeneous graphs, also called Heterogeneous Information Networks (HIN) [10], provide a robust framework for representing diverse entities and interactions in drug discovery.

In these graphs, nodes represent entities like drugs, proteins, pathways, chemical substructures, ATC codes, and diseases, while edges capture interactions between them. While many models are developed to represent the relationship between drug, target, and disease, they focus on predicting pairwise relations between drug and other entities such as drug-drug, drug-disease and drug-target [11]–[15]. These methods base their predictions on established drug-drug similarity, target-target similarity as well as known drug-target associations. However, there is no HIN-based triplet prediction model.

To address these limitations and model the complex interactions between drugs, targets, and diseases more effectively, we propose a novel **He**terogeneous Graph **T**riplet **A**ttention **N**etwork (`HeTAN`). `HeTAN` leverages the power of heterogeneous graphs, representing diverse entities and their interactions, and employs a novel triplet attention mechanism to capture higher-order interactions within the drug-target-disease triplets. We capture higher-order interactions between drug, target, and disease through a triplet-wise attention mechanism. This gives us a more comprehensive understanding of drug MoAs and can accelerate drug repurposing for personalized medicine. While it is defined for drugs, targets, and diseases triplets, it is a generic model that can be applied to other triplets. Our main contributions are as follows:

- **Utilizing heterogeneous graph neural network for drug-target-disease triplet prediction:** We propose a novel approach that models the complex interactions between drugs, targets, and diseases using a heterogeneous graph neural network (HGNN). By incorporating different types of nodes and edges, our approach effectively captures the rich information embedded in the interactions between these entities, leading to improved prediction performance.
- **Introducing the `HeTAN` model:** We develop a novel model, `HeTAN`, by proposing a novel triplet message passing and triplet-wise attention mechanisms on different types of entities in a heterogeneous graph. Our model goes beyond the pair-wise interaction and captures higher-order triplet-wise interactions to make triplet predictions on the heterogeneous graph. While triplet message passing enables passing the information among three different entities (drug-target-disease), the triplet attention mechanism enables the model to focus on the most relevant pairs for an entity instead of the most relevant neighbor. These enhance its predictive accuracy and its ability to capture intrinsic and complex interactions among three entities. No prior work in GNN and HGNN has explored triplet-wise message-passing and attention mechanisms.
- **Extensive Experiments:** We conduct extensive experiments to show the effectiveness of our model on two different datasets. We also compare the proposed `HeTAN` model with several baseline models. The results with different accuracy measures show that our method significantly surpasses the baseline models. In addition, different case studies denote that different datasets and external literature evidence can validate our model's predictions.

The remainder of this paper is organized as follows: Section II reviews related works. Section III describes the creation of a heterogeneous graph. Details of the `HeTAN` model are presented in Section III. Experiments and results are discussed in Section IV. Finally, Section V concludes the paper.

II. Related Works

This section provides an overview of existing research in computational predictions of drugs, targets, and diseases, specifically on triplet prediction for high-dimensional structured data.

A. Modeling drug-target-disease

Treating human diseases involves interactions among drugs, biological targets, and disease pathways. Computational pharmacology seeks to uncover associations among these entities and understand drugs' mechanisms of action (MoAs) [1]. A common technique involves network-based inference models, such as bipartite networks with distinct layers for drugs and diseases (targets). Various machine learning methods, including random walks, matrix factorization, and support vector machines [16], have been used to predict new drug-disease and drug-target interactions.

DTINet [15] integrates diverse drug-related information to build a heterogeneous network and employs a compact feature learning algorithm to derive low-dimensional vector representations of nodes. This model uses a known set of drug-target associations as a reference to determine the optimal projection from the drug space onto the protein space, ensuring that the projected feature vectors of drugs closely align with the feature vectors of their known targets. Chen et al. [16] integrated a protein-protein similarity network, a drug-drug similarity network, and a drug-target interaction network into a heterogeneous network. Using a random walk algorithm, they inferred new drug-target connections without directly modeling drug-disease relationships, focusing on predicting drug-target interactions by learning a transformation matrix from known interactions. Similarly, Fu et al. [17] utilized known drug-target connections from various data sources but did not explicitly use drug-disease-target triples, thus only predicting drug-target interactions. Zheng et al. [18] developed a matrix factorization method based on the similarity of chemical structures and protein sequences to establish drug-target relationships. These methods rely on chemical structure similarity through structural fingerprints and protein sequence-based similarity. However, these approaches treat drug-disease and drug-target predictions as separate tasks, limiting a comprehensive understanding of the interconnected drug-target-disease relationships.

B. Triplet Prediction

Triplet prediction has broad applications, ranging from drug repurposing to natural language processing and computer vision. Zhang et al. [19] introduced an attention mechanism based on transformers to capture relationships between three

entities (query, key, and value) for improved language understanding and generation. In natural language processing, triplet prediction has been used for tasks like relationship extraction, where models use sentence-level attention and entity descriptions to predict relationships between entity triplets in text [20]. In computer vision, triplet prediction techniques have been used for face recognition and person re-identification [21].

In drug-target-disease prediction, various models have been developed, such as collective matrix factorization [22] and neural tensor networks [7], to capture nonlinear dependencies within triplets. Recent research explores the interdependence of drugs, targets, and diseases through event-graph modeling and neural tensor network models [7], [23]. Meanwhile, recent advancements have explored the potential of hypergraphs and hypergraph neural networks in biomedical problems [24], [25], including drug-microbe-disease associations, [26]. A hypergraph is a unique graph with hyperedges. Unlike a regular graph where the degree of each edge is 2, hyperedge is degree-free; it can connect an arbitrary number of nodes. While these approaches are significant, the application of heterogeneous graph neural networks (HGNNs) and triplet attention mechanisms remains largely unexplored.

Despite many proposed models for triplet prediction, most focus on homogeneous entities, and none have applied graph-based models to drug-related problems. `HeTAN` distinguishes itself in drug repurposing by pioneering the combined use of HGNNs and triplet attention mechanisms. While defined for drug-target-disease triplets, `HeTAN` can be applied to other triplets as well.

III. Methodology

Given that a triplet includes a drug, a target, and a disease, our goal is to predict whether the triplet has an interaction. In this paper, we propose a novel approach that leverages the power of heterogeneous information networks (HIN) and introduces the concept of triplet attention. To achieve this, we develop the Heterogeneous Graph Triplet Attention Network (`HeTAN`), which employs an end-to-end encoder-decoder architecture. The encoder integrates a triplet attention mechanism to determine the significance of pairs (e.g., target-disease) for the other entity (e.g., drug) while learning embeddings of all entities and triplets. Moreover, `HeTAN` incorporates a decoder that learns and predicts the interaction between entities of triplets. The system architecture of `HeTAN` is outlined in Figure 1. We optimize the model parameters with a cross-entropy loss function.

Our proposed model consists of the following steps:

1) Heterogeneous Graph Construction & Node's Feature Extraction
2) Heterogeneous Graph Triplet Attention Network Architecture
   - Encoder: Triplet Attention-based node representation learning
   - Decoder: Drug-Target-Disease triplet prediction

A. Heterogeneous Graph Construction & Node's Feature Extraction

The first step in our approach is to construct a heterogeneous graph that captures the complex relationships among drugs, proteins (targets), and diseases. The graph consists of three types of nodes: *drugs*, *proteins*, and *diseases*. We establish edges between these nodes based on known drug-target interactions and drug-disease associations. This construction allows us to represent the rich interactions and dependencies between different entities in the graph.

Graph neural network (GNN) models can optimize and refine node representations with an iterative learning process. These models transform initial node attributes or features with message passing and aggregation mechanisms from the node's neighbors to generate enriched and effective node vector representations. In our next step, we utilize the structural properties of drugs and targets to extract node features from our heterogeneous graph. Specifically, we focus on the chemical substructures of drugs and targets, represented as SMILES strings [27] and Amino Acid sequences, respectively. We employ the Explainable Substructure Partition Fingerprint (ESPF) [28] algorithm to create drug and target features. ESPF decomposes the SMILES string and Amino Acid sequence into frequent substructures, selecting the most significant ones based on a frequency threshold. These substructures provide informative features for the drugs and targets utilized in the subsequent steps of the `HeTAN` model. We represent disease nodes with one-hot encoded representations. After constructing the heterogeneous graph and extracting node features, we propose the `HeTAN` architecture for learning representations that capture complex relationships among drugs, targets, and diseases.

B. Heterogeneous Graph Triplet Attention Network

The core goal of our research is to address the challenge of predicting drug-target-disease interactions. To achieve this, our model `HeTAN` leverages the rich information in heterogeneous networks and captures the complex relationships among drugs, targets, and diseases. The model is trained using an end-to-end approach to predict drug-target-disease interactions. The model is responsible for aggregating information from neighboring nodes and learning higher-order relationships in the graph using a triplet attention mechanism, the critical component of our model. The triplet attention mechanism calculates attention coefficients based on the features of all three nodes in a triplet (i, j, and k), where i is the central node, and j and k are neighboring nodes. When aggregating information, these coefficients are used to weigh the importance of neighboring node pairs for the central nodes in each triplet. The attention mechanism is applied to all graph-generated triplets, enabling the model to capture complex, interconnected relationships among the different node types. Afterward, the encoded representations of nodes are obtained by aggregating information from neighboring nodes, weighted by the attention coefficients. This process is repeated for each layer of

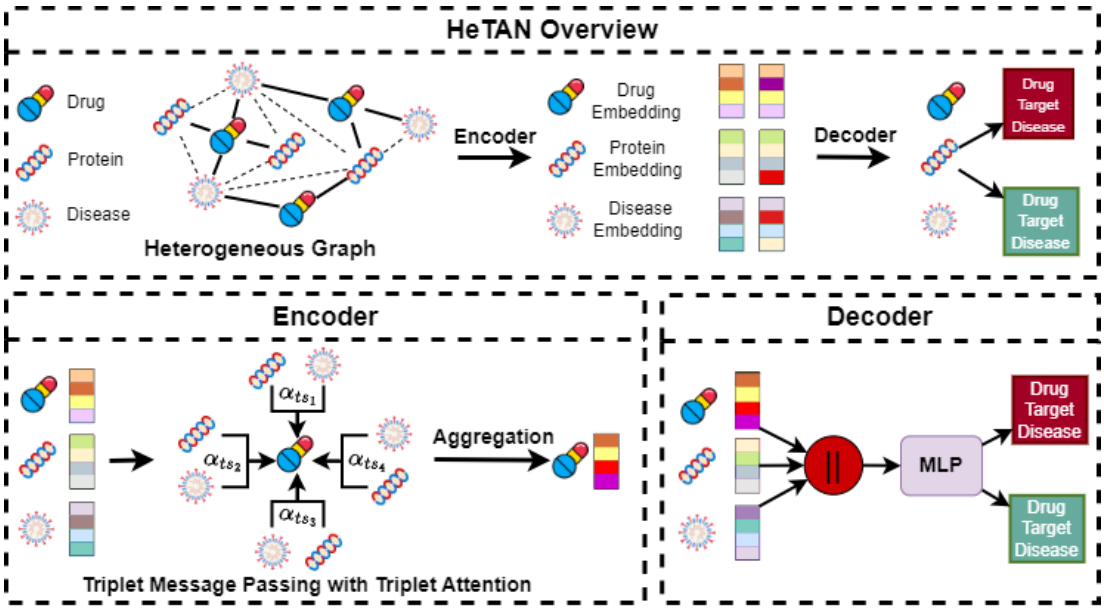

Fig. 1: The HeTAN workflow comprises three steps: Heterogeneous Graph Construction, Encoder, and Decoder. Initially, a heterogeneous network is built with drug, target, and disease nodes connected by drug-target or drug-disease edges. Target-disease connections (dashed lines) are inferred from shared drug associations. We introduce Triplet Message Passing (TMP) and Triplet-wise attention to generate node representations. Finally, using concatenated node representations, a Multi-Layer Perceptron (MLP) predicts drug-target-disease interactions.

the HeTAN model, allowing the model to learn increasingly complex patterns and dependencies across multiple layers.

Our proposed model consists of the following steps:

- Encoder: Triplet Attention-based Node Representation learning
- Decoder: Predicting Drug-Target-Disease Interactions

*1) **Node Representation Learning (Encoder)**:* The encoder component of our model focuses on learning informative node representations. In a heterogeneous graph, nodes and edges belong to different types, and each type of node has its own distinct feature space. To effectively learn informative node representations, we must align these diverse feature spaces into a common one. This enables meaningful comparisons and interactions among nodes of different types. To overcome this challenge, we introduce a type-specific transformation matrix M, which projects the features of different nodes into a common feature space as $h_i^{'} = M \circ h_i$

Traditional Graph Convolutional Networks (GCNs) rely on pairwise message passing, where neighboring nodes pass messages to each other. However, this approach falls short in capturing intricate dependencies beyond pairwise relationships, especially in our case, where understanding drug-target-disease interactions requires considering the complex relations inherent in drug-target-disease triplets.

To solve this limitation, we define the **Triplet Message Passing** function (TMP), a novel mechanism that leverages node triplets for representation learning. Instead of pairwise interactions, **TMP** considers neighboring node pairs and their influence on the central node. For a central node of type i, we define its neighbors as node pairs (N_i) comprising node pairs of type j and k as $N_i = \{(j_1, k_1), ..., (j_n, k_n)\}$. We pass messages from these neighbor pairs to the central node i. This allows the model to capture richer contextual information and complex relationships in drug-target-disease interactions. The triplet message passing function is defined as:

$$z_i^l = TMP(z_i^{l-1}, N_i) \tag{1}$$

For one central node, there are several node pairs as the neighbors. However, it is essential to note that not all neighbor pairs are equally crucial for the central node. Message passing should consider these varying levels of importance. We design a novel **Triplet-wise Attention** mechanism to incorporate the importance of neighbor pairs for a central node into message passing. This attention mechanism utilizes the features of all three nodes in a triplet and assigns attention coefficients, signifying the relative importance of the neighbor pairs for the central node. Based on the features of all three nodes in a triplet, the attention coefficient e_{ijk} is defined as follows:

$$\begin{aligned} e_{ijk} &= a(h_i^{'}, h_j^{'}, h_k^{'}) \\ &= LeakyRELU(NN(h_i^{'} || h_j^{'} || h_k^{'})) \end{aligned} \tag{2}$$

In Eq. 2, a denotes the triplet-wise attention mechanism, and $||$ denotes the concatenation operation. We employ a neural network in the attention mechanism, denoted as NN. This neural network is designed to capture essential relationships and dependencies among the nodes in a triplet. Additionally,

to capture the nonlinear dependencies among drug-target-disease data, we apply the LeakyReLU activation function. LeakyReLU is chosen for its ability to introduce nonlinearity in the model, allowing it to capture complex relationships critical for accurately predicting drug-target-disease interactions.

It is vital to make attention coefficients easily comparable across different nodes. Therefore, the attention coefficients are then normalized using a Softmax function. This step ensures that the model appropriately weighs the attention of each neighbor when aggregating information. The normalized attention coefficient α_{ijk} is defined as follows;

$$\begin{aligned}\alpha_{ijk} &= softmax_j(e_{ijk}) \\ &= \frac{\exp(e_{ijk})}{\sum_{l,m\in N(i)} \exp(e_{ilm})}\end{aligned} \tag{3}$$

During the triplet message passing process, it is imperative to consider the message from neighbor pairs. To generate messages from pairs, we concatenate the representations of the nodes within the pair. We then pass this concatenated feature vector of size $2d$ through a single-layer feedforward neural network to transform it into a feature vector of size d. After multiplying each pair message with calculated pair-wise attention, these messages are aggregated and combined with the central node's representation using a self-attention mechanism. This mechanism considers the node's features and aggregated information, allowing the model to capture its unique influence within the heterogeneous graph. So, the triplet message passing function, **TMP** in Eq 1 is defined as follows:

$$z_i = \delta(h_i^{'} + W \circ \sum_{j,k\in N(i)} (\alpha_{ijk} \circ NN(h_j^{'} || h_k^{'}))) \tag{4}$$

where $\circ$ represents multiplication operator and W is a trainable parameter. To incorporate self-attention, we use W, which determines the weight or importance of node i's embedding in the aggregation process. The purpose of W is to control the influence of node i's features and the aggregated features of its neighboring nodes on the overall representation.

We employ multi-head attention to capture more complex patterns and relationships, each focusing on different aspects of the data. This further enhances the model's ability to learn intricate patterns and relationships within the heterogeneous graph. In multi-head attention, multiple attention mechanisms (K) are used individually to transform the features, and the outputs are concatenated $(||)$ to obtain the final representation. So, the final triplet message passing with multi-head attention is defined as follows;

$$z_i = ||_{k=1}^{K} \delta(h_i^{'} + W \circ \sum_{j,k\in N(i)} (\alpha_{ijk} \circ NN(h_j^{'} || h_k^{'}))) \tag{5}$$

By integrating these equations and steps, our model learns informative node representations within a heterogeneous graph. This ensures it captures complex relationships, intricate patterns, and crucial interactions for predicting drug-target-disease interactions.

*2) **Drug-Target-Disease Triplet Prediction (Decoder)**:* Our model has a decoder component that predicts the likelihood of interactions between drugs, targets, and diseases as new triplets based on the representations of entities obtained from the encoder. Decoder, in particular, assigns a score to drug, target, and disease triplet (v_i,v_j,v_k), expressing how likely it is that drug v_i target v_j, and disease v_k are interacting. The corresponding entities' features are concatenated and passed through a multilayer perceptron (MLP).

$$pred_{x,y,z} = MLP(z_x || z_y || z_z) \tag{6}$$

The MLP outputs a prediction score, Y', between 0 and 1. A score close to 1 indicates a high likelihood of interaction among the triplets, whereas a score close to 0 indicates less likely interaction.

*3) **Model Optimization**:* We train our entire encoder-decoder architecture as a binary classification problem by minimizing a binary cross-entropy loss function specified as follows:

$$L = -\sum_{i=1}^{N} Y_i \log Y_i^{'} + (1 - Y_i) \log(1 - Y_i^{'}) \tag{7}$$

where N is the total number of triplets, Y_i is the actual label indicating the presence or absence of an interaction for the triplet, and $Y_i^{'}$ is the predicted score for the triplet.

C. Analysis of HeTAN model

Here we give the analysis of HeTAN as follows:

- **Handling Diverse Nodes and Relationships:** HeTAN effectively handles different types of nodes and relationships, integrating rich semantics within a heterogeneous graph. We facilitate message passing among neighbor node pairs and a given node. During this message passing, we incorporate the importance of neighbor pairs for a central node through a novel Triplet-wise Attention mechanism. Leveraging Triplet Message Passing and Triplet-wise Attention allows for enhanced integration, promotion, and improvement of diverse node embeddings.
- **Efficiency and Complexity:** The proposed HeTAN is highly efficient and can be easily parallelized. The complexity of HeTAN can be analyzed based on its main components: heterogeneous graph construction, triplet attention mechanism, and message passing. The initial step of constructing a heterogeneous graph involves nodes (drugs, targets, diseases) and edges (interactions), with a complexity of $O(|V| + |E|)$. For each triplet of nodes (i, j, k), the triplet attention mechanism computes attention scores and normalizes them, resulting in a complexity of $O(|E|ff' + |V|df')$, where f is the initial feature dimension, f' is the output feature dimension, and d is the average degree of nodes.

  HeTAN leverages multi-head attention to capture complex patterns, scaling the computation by the number

TABLE I: Statistics of Dataset

# of Instances	DrugBank (DB)	DrugBank and CTD
Drugs	531	450
Targets	836	708
Diseases	279	1, 267
Triplets	27,238	175, 288

TABLE II: Hyper-parameter Settings

Parameter	Values
Learning rate	1e-2, 5e-2, 1e-3, 5e-3, 1e-5
Number of heads per layer	8, 16, 32
Hidden units	8, 16, 32, 64, 128
Dropout	0.1, 0.3, 0.5, 0.6
Weight decay	0.01, 0.001

of heads K. The message passing and aggregation process, combined with multi-head attention, contributes to an overall complexity of $O(K(|E|ff' + |V|df'))$. This efficient handling of heterogeneous graphs and higher-order interactions enables HeTAN to effectively capture intricate relationships among drugs, targets, and diseases, demonstrating its capability in drug-target-disease triplet prediction.

IV. Experiment

To evaluate our HeTAN model, we conduct experiments involving negative sampling and random dataset splitting into train and test sets. Our performance assessment include Recall, Precision, F1-score, AUROC, and the commonly used top-n metric hit@n. This section summarizes our experimental parameters, evaluation protocols, and analysis of results.

A. Datasets, Parameter Settings & Baselines

Our study utilizes data from DrugBank and CTD, providing insights into drug-related information. Two dataset configurations are employed. One uses data exclusively from DrugBank, encompassing details about drug-target interactions and drug-disease associations. The other configuration integrates information from DrugBank (concerning drug-target interactions) with data from CTD (providing drug-disease associations). This integration provides a comprehensive view of <drug, target, disease> triplets. Subsequently, the DrugBank and combined datasets will be referred to as DB and DB&C, respectively. Table I summarizes the vital characteristics of nodes and edges in the heterogeneous graph.

Datasets are split into random training (80%) and testing (20%) subsets for five iterations. This splitting process is repeated five times, and the average accuracy metrics are calculated and reported in the results section. The optimal hyper-parameters are obtained by grid search based on the validation set. The ranges of grid search are shown in Table II. We train the HeTAN model using the cross-entropy loss and optimize the model parameters using the Adam optimizer. The optimal learning rate is determined to be 1e-5, and the optimal dropout rate is found to be 0.6 to prevent overfitting. Training runs for 2000 epochs with early stopping after 200 consecutive epochs without validation loss improvement.

To construct negative triplets, we employ negative sampling by randomly replacing one or all nodes in positive triplets, ensuring they are absent from the actual data. We assess HeTAN's performance through a diverse set of metrics encompassing accuracy, precision, F1-score, and AUROC. We also utilize the commonly used top-n metric hit@n and NDCG@n, as proposed by [29], [30]. Hit@n measures whether a test triplet appears within the top-n ranked predictions, while NDCG@n prioritizes higher-ranked matches. We rank triplets in descending order based on model prediction scores, prioritizing those most likely to represent valid interactions.

To evaluate HeTAN's effectiveness, we compare it to a range of state-of-the-art models categorized by their approaches:

- *Tensor Decomposition Methods:* CP and Tucker are famous tensor models with diverse variants that are being successfully applied in health data analysis [31]. They both adopt multilinear assumptions.
- *Attention-based Methods:* We use transformer, a robust deep learning architecture that captures complex relationships and patterns in the data. It utilizes self-attention mechanisms to effectively learn and represent the interactions between drugs, targets, and diseases. For these models, the embeddings of the triplet nodes (drug, target, disease) are concatenated, and the combined embeddings are used to predict interactions.
- *Graph Neural Network (GNN):* We use GNN architectures on our heterogeneous graph to learn the representation of nodes. We select three standard GNN-based methods: GIN [32], GAT [33], and GraphSAGE [34]. Among these GNN models, GAT [33] uses pairwise attention to generating node representation. Similar to attention-based methods, GNN models concatenate the embeddings of the triplet nodes and use these concatenated embeddings for interaction prediction.
- *Heterogeneous Graph Neural Network (HGNN):* For this baseline, we use the commonly used HGNN model heterogeneous graph transformer (HGT) [35]. HGT incorporates pairwise attention on a heterogeneous graph to learn the representation of nodes. The triplet node embeddings are concatenated and used to predict interactions.
- *NeurTN:* Neural Tensor Network (NeurTN) [7] combines tensor algebra and deep neural networks, offering a more powerful way to capture the nonlinear relationships among drugs, targets, and diseases. Both NeurTN and HeTAN combine drug-target and drug-disease interactions from DrugBank and CTD.
- *DDTE:* Moon et al. [36] construct a heterogeneous knowledge graph including various drug-related information and utilize TransE [30] model to infer drug–disease–target relationships.
- *MHGNN:* MHGNN-DTI [37] builds the model with a dual-channel architecture to learn drug and target embeddings, respectively, using a graph attention mechanism and metapath techniques. It proposes building correlation graphs to exploit high-order relations. Finally, it performs

TABLE III: Comparing performance of HeTAN with other baseline models on DB

Model	Method	F-1 Score	Precision	Recall	ROC-AUC	AUPR
Tensor-based	TD	47.00	48.51	45.59	49.19	48.98
	CPD	52.91	52.19	56.19	49.84	50.06
Attention-based	Transformer	52.31	62.96	51.18	60.62	59.46
GNN-based	GraphSAGE	72.24	61.3	83.94	66.4	59.92
	GIN	74.06	71.18	77.2	73.08	66.31
	GAT	72.63	62.65	82.34	67.64	60.92
HGNN-based	HGT	80.44	82.71	79.41	83.13	83.32
	MHGNN	81.9	87.16	83.86	92.63	91.57
	HeTAN	**86.31**	**88.43**	**84.34**	**93.46**	**93.07**

TABLE IV: Comparing performance of HeTAN with other baseline models on DB&C

Model	Method	F-1 Score	Precision	Recall	ROC-AUC	AUPR
Tensor-based	TD	53.17	62.45	47.86	60.65	62.04
	CPD	57.23	63.72	52.19	59.82	60.76
Attention-based	Transformer	83.05	85.24	81.09	82.04	75.36
GNN-based	GraphSAGE	83.31	78.97	90.04	75.98	70.72
	GIN	83.98	78.16	83.49	76.45	71.52
	GAT	85.17	82.44	82.11	83.76	77.15
HGNN-based	HGT	85.22	87.08	84.55	87.98	84.07
	MHGNN	87.7	85.88	88.79	95.64	94.45
	HeTAN	**90.91**	**93.12**	**89.88**	**98.01**	**97.75**

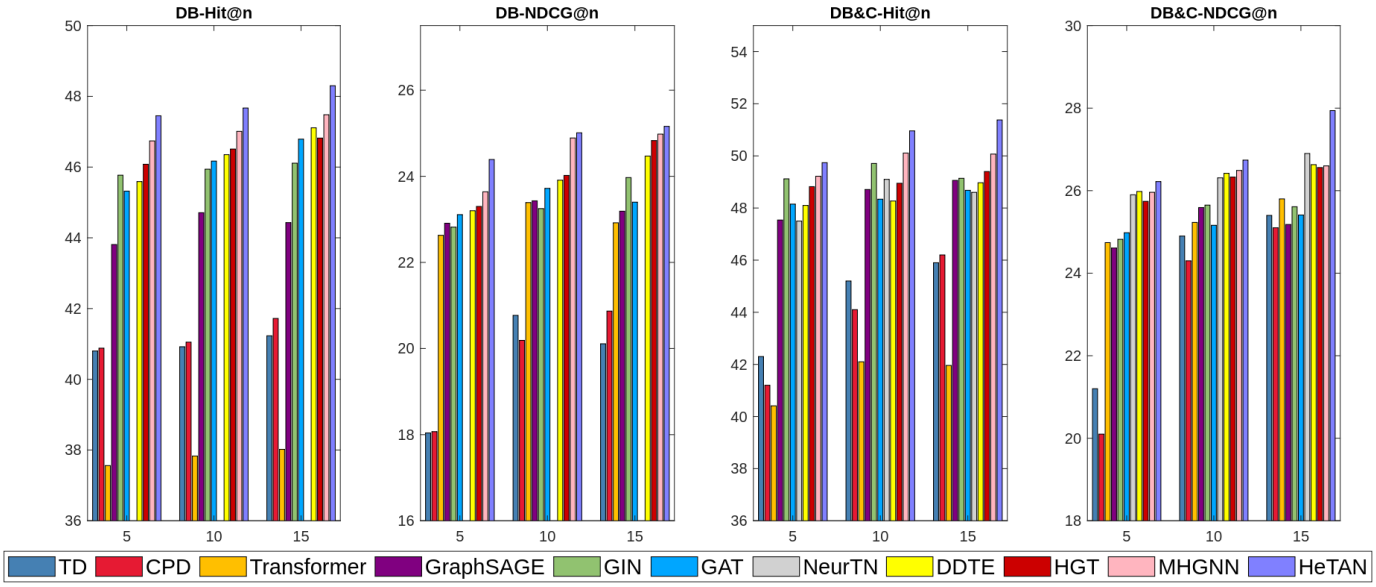

Fig. 2: Evaluation of top-n performance for HeTAN and other baseline models in terms of a) Hit@n and b) NDCG@n on DB and DB&C

pairwise drug-target interaction prediction.

B. Comparison with baselines

In this study, we conduct a comprehensive performance analysis of HeTAN compared to a selection of state-of-the-art baseline models. We employ diverse performance metrics to assess these models' efficacy. Specifically, we report the F-1 Score, Precision, Recall, ROC-AUC, and AUPR results in Table III for DB and Table IV for DB&C. Both tables refer to the tensor-based baselines, tucker decomposition and CP decomposition as TD and CPD, respectively. Our model, HeTAN, outperforms all baseline models for both datasets, showcasing its exceptional predictive capabilities.

For instance, on DB, HeTAN achieves impressive F-1 score, ROC-AUC, and AUPR of 86.31%, 93.46%, and 93.07%, representing significant improvements over the best-performing baseline, MHGNN, which achieves F-1 score, ROC-AUC, and AUPR of 81.9%, 92.63%, and 91.57%, respectively. The superior performance of HeTAN is further evident in the DB&C dataset, where it attains F-1 score, ROC-AUC, and AUPR of 90.91%, 98.01%, and 97.75%, surpassing the performance of other models by a considerable margin.

TABLE V: Novel Triplet Predictions by HeTAN FROM DB&C

Drug	Target	Disease	DB&C Label	Prediction	DB Label
Carbamazepine	NR1I2-HUMAN	Osteoporosis	0	0.99	1
Testosterone	ERR3-RAT	Myocardial infarction	0	0.98	1
Nefazodone	DRD2-HUMAN	Schizophrenia	0	0.97	1
Raloxifene	ERR3-RAT	Obesity	0	0.93	1
Fenofibrate	MMP19-HUMAN	Psoriatic arthritis	0	7e-09	0

TABLE VI: Novel Triplet Predictions by HeTAN FROM DB

Drug	Target	Disease	DB Label	Prediction	DB&C Label
Cyclobenzaprine	5HT2C-HUMAN	Muscle Spasm	0	0.99	1
Cyclobenzaprine	AA2AR-HUMAN	Gout	0	0.98	1
Imipramine	ADA1D-HUMAN	Interstitial Lung Disease	0	0.97	1
Quetiapine	HRH1-HUMAN	Schizophrenia	0	9.9e-10	0
Verapamil	CAC1S-HUMAN	Cluster headache	0	5e-07	0

TABLE VII: Top five drug-target pairs predicted by our proposed HeTAN for depression

Drug (DrugBank)	Target (UniProt)	Evidence
Amitriptyline	Sodium-dependent serotonin transporter	Kim Lawson [38]
Nortriptyline	5-hydroxytryptamine receptor 2A	Pierre Blier [39]
Imipramine	Sodium-dependent serotonin transporter	Dempsey et al. [40]
Nortriptyline	Muscarinic acetylcholine receptor M5	Philip et al. [41]
Nortriptyline	MD(2) dopamine receptor	Pierre Blier [39]

In addition to these performance metrics, we adopt the top-n metrics, Hit@n and NDCG@n, as illustrated in Figure 2. These metrics are particularly critical in triplet prediction, as they assess the ranking quality of the model's predictions. HeTAN's top-n metrics (Hit@n and NDCG@n) performance showcases its superior ranking ability, which is crucial for accurate triplet prediction. On DB, HeTAN achieves a Hit@15 score of 50.11% and NDCG@15 of 27.36%, significantly exceeding the top baseline (MHGNN) by over 6% and 3%, respectively.

Tensor-based models exhibit solid performance but often lag in recall and ranking. Attention-based methods, like the Transformer, improve Precision and Recall. NeurTN, combining tensor and attention models, excels in top-n metrics. Since different accuracy results like F-1 score, Precision and Recall are unavailable on NeurTN paper, we could not present and analyze these results with other baseline models. Similarly, we could not obtain these results from DDTE.

GNN and HGNN-based models consistently achieve F-1 scores surpassing 70%, emphasizing the pivotal role of graph structural information. GNN and HGNN-based models represent interactions between graph nodes and capture graph dependence through message passing. Comparing graph attention-based models, GAT and HGT rely on pairwise attention, and HeTAN utilizes triplet-wise attention. HeTAN consistently performs better than GAT and HGT. For example, GAT achieved F-1 scores of 72.63% on DB and 85.17% on DB&C, while HGT scored 80.44% on DB and 85.21% on DB&C. One notable heterogeneous graph neural network, MHGNN achieves F-1 scores of 81.9% on DB and 87.7% on DB&C, demonstrating strong performance. MHGNN's strength lies in its dual-channel architecture and meta-path techniques to exploit high-order relations. Still, MHGNN falls short of HeTAN's results. HeTAN achieve 86.31% and 90.91% on DB and DB&C, respectively.

HeTAN effectively manages diverse nodes and relationships, integrating rich semantics within a heterogeneous graph. By using triplet message passing and triplet-wise attention, the model captures intricate patterns and dependencies, offering a comprehensive understanding of drug-target-disease associations. Multi-head attention enhances its ability to learn from complex data, ensuring robust predictions. Overall, HeTAN significantly improves prediction accuracy, positioning itself as a powerful tool for drug discovery and personalized medicine.

C. Prediction and Validation of Triplets

This study evaluates HeTAN's ability to predict drug-target-disease interactions using real-world datasets. To determine its effectiveness in predicting missing interactions, we compare HeTAN's predictions with data from two distinct datasets.

We start by selecting triplets from Dataset DB&C, which lack interaction data in DB&C but possess relevant association information in DB. We train HeTAN on the DB&C dataset, ensuring that the selected triplets are used exclusively in the test set to minimize potential bias. The predicted scores for these triplets, presented in Table V, consistently exceed 90%, suggesting that these triplets are likely to exhibit interactions despite the absence of explicit interaction data in DB&C. To validate these predictions further, we cross-reference them with the information in DB. Remarkably, this comparison confirms the interactions between these triplets, reinforcing the predictive power and accuracy of HeTAN.

To expand our validation process and test HeTAN's generalizability, we select another set of five drug triplets from DB, which lack interaction information within DB but contain such data in DB&C, as highlighted in Table VI. We train HeTAN using the DB dataset for this validation, tailoring the model specifically to this unique dataset configuration. Subsequently, we validate the predicted scores by cross-referencing them with DB&C, which serves as an independent

validation set. Validating predicted scores against DB&C emphasizes HeTAN's reliability and generalizability, showcasing its adaptability across datasets and reinforcing its real-world predictive capabilities.

D. Case Study on Depression

Personalized treatment is a core objective in our medical research, particularly in identifying effective drugs for specific diseases and understanding their biological targets. HeTAN has been employed to uncover new drug-target combinations relevant to depression—a complex condition with various molecular factors. By focusing on triplets where the disease is depression, HeTAN was trained on a heterogeneous graph from DrugBank and CTD datasets. For this experiment, we filter our predicted triplets to focus on those where the disease is depression. Table VII enlists the top five pairs of (drug, target) corresponding to depression and literature evidence supporting these predictions. For depression, these pairs are the highest-ranked predictions based on the model's scoring and have corresponding evidence in the literature, demonstrating their potential relevance and validity. These results underline HeTAN's potential in identifying clinically relevant drug-target pairs, marking a significant step toward personalized medicine. The model's reliable predictions offer a promising approach to revolutionizing treatments for complex diseases like depression.

E. Ablation Study

To assess the contribution of each component in HeTAN, we perform an ablation study with five variants:

- `HeTAN-Sum (HeTAN-S)`: This variant employs summation instead of concatenation and neural network transformations for neighbor embedding in Eq 4.
- `HeTAN-Concat (HeTAN-C)`: Three neighbor node embeddings are concatenated in Eq 2 and then reduced in dimension.
- `HeTAN-Elem-Prod (HeTAN-EP)`: This variant uses the element-wise product on neighbor node embeddings in Eq 4.
- `HeTAN-Triplet-Attention-Sum (HeTAN-TAS)`: In this variant, three neighbor node embeddings are summed in Eq 2 to get triplet-wise attention.
- `HeTAN-Triplet-Attention-Elem-Prod (HeTAN-TAEP)`: This variant applies the element-wise product on three neighbor node embeddings in Eq 2 to get triplet-wise attention.

In comparing the model variants with the original HeTAN, HeTAN-Sum and HeTAN-Concat demonstrate weaker performance, likely due to their use of summation or concatenation, which may not capture complex relationships as effectively as the original approach. Similarly, HeTAN-TAEP and HeTAN-TAS underperform compared to HeTAN, highlighting the efficacy of applying a neural network for concatenated embeddings. As shown in Figure 3, the original HeTAN model consistently surpasses its variants across key metrics like F1-score, Recall, and ROC-AUC, underscoring the effectiveness of its triplet-wise attention and message passing mechanisms. This analysis confirms HeTAN's robustness in managing complex, heterogeneous data in biomedical research.

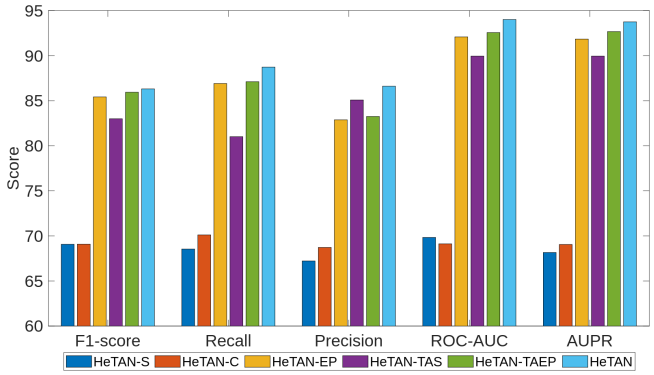

Fig. 3: Performance Comparison of HeTAN with its variants

V. Conclusion

HeTAN stands out as a powerful model for modeling drug-target-disease interactions thanks to its dedicated HGNN architecture and innovative triplet-attention mechanism. This approach effectively addresses limitations encountered in previous models, leading to significant improvements in performance. The novel triplet-attention mechanism holds broad potential for application beyond drug discovery, extending to diverse domains involving heterogeneous graphs and higher-order interactions.

While HeTAN is currently defined for drug-target-disease triplets, future research could further enhance its capabilities by applying it to different triplet combinations and incorporating additional elements, such as drug-target-pathway-disease interactions. This expansion could lead to a deeper understanding of drug mechanisms and improved predictive accuracy. Moreover, integrating multi-omics data and exploring more complex graph structures are promising avenues for boosting HeTAN's predictive power and providing a more comprehensive view of biological processes. These advancements can significantly contribute to progress in personalized medicine and drug development, ultimately benefiting patient outcomes and healthcare systems.

References

[1] A. S. Hauser, M. M. Attwood, M. Rask-Andersen, H. B. Schiöth, and D. E. Gloriam, "Trends in gpcr drug discovery: new agents, targets and indications," *Nature Reviews Drug Discovery*, vol. 16, pp. 829–842, 2017.

[2] H. Chen and J. Li, "Drugcom: Synergistic discovery of drug combinations using tensor decomposition," *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 899–904, 2018.

[3] M. Zitnik, M. Agrawal, and J. Leskovec, "Modeling polypharmacy side effects with graph convolutional networks," *Bioinform.*, vol. 34, no. 13, pp. i457–i466, 2018. [Online]. Available: https://doi.org/10.1093/bioinformatics/bty294

[4] A. Ezzat, P. Zhao, M. Wu, X. Li, and C. K. Kwoh, "Drug-target interaction prediction with graph regularized matrix factorization," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 14, pp. 646–656, 2017.

[5] X.-Y. Yan, S. Zhang, and C. He, "Prediction of drug-target interaction by integrating diverse heterogeneous information source with multiple kernel learning and clustering methods," *Computational biology and chemistry*, vol. 78, pp. 460–467, 2019.

[6] H. Chen and J. Li, "Modeling relational drug-target-disease interactions via tensor factorization with multiple web sources," *The World Wide Web Conference*, 2019.

[7] ——, "Learning data-driven drug-target-disease interaction via neural tensor network," in *International Joint Conference on Artificial Intelligence*, 2020.

[8] S. Zhe, K. Zhang, P. Wang, K.-c. Lee, Z. Xu, Y. Qi, and Z. Ghahramani, "Distributed flexible nonlinear tensor factorization," *Advances in neural information processing systems*, vol. 29, 2016.

[9] Y. Shan, T. R. Hoens, J. Jiao, H. Wang, D. Yu, and J. Mao, "Deep crossing: Web-scale modeling without manually crafted combinatorial features," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 255–262.

[10] C. Shi, Y. Li, J. Zhang, Y. Sun, and P. S. Yu, "A survey of heterogeneous information network analysis," *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, pp. 17–37, 2017.

[11] F. Tanvir, M. I. K. Islam, and E. Akbas, "Predicting drug-drug interactions using meta-path based similarities," *2021 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*, pp. 1–8, 2021.

[12] F. Tanvir, K. M. Saifuddin, M. I. K. Islam, and E. Akbas, "Predicting drug-drug interactions using heterogeneous graph attention networks," *Proceedings of the 14th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:263621744

[13] F. Tanvir, K. M. Saifuddin, M I. K. Islam, and E. Akbas, "Ddi prediction with heterogeneous information network - meta-path based approach," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, pp. 1–12, 2024.

[14] F. Wan, L. Hong, A. Xiao, T. Jiang, and J. Zeng, "Neodti: Neural integration of neighbor information from a heterogeneous network for discovering new drug-target interactions," *bioRxiv*, 2018. [Online]. Available: https://api.semanticscholar.org/CorpusID:49658098

[15] Y. Luo, X. Zhao, J. Zhou, J. Yang, Y. Zhang, W. Kuang, J. Peng, L. Chen, and J. Zeng, "A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information," *Nature Communications*, vol. 8, 2017.

[16] X. Chen, M.-X. Liu, and G.-Y. Yan, "Drug–target interaction prediction by random walk on the heterogeneous network," *Molecular BioSystems*, vol. 8, no. 7, pp. 1970–1978, 2012.

[17] G. Fu, Y. Ding, A. Seal, B. Chen, Y. Sun, and E. Bolton, "Predicting drug target interactions using meta-path-based semantic network analysis," *BMC Bioinformatics*, vol. 17, 2016.

[18] X. Zheng, H. Ding, H. Mamitsuka, and S. Zhu, "Collaborative matrix factorization with multiple similarities for predicting drug-target interactions," *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013.

[19] H. Zhou, J. Li, J. Peng, S. Zhang, and S. Zhang, "Triplet attention: Rethinking the similarity in transformers," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2378–2388.

[20] G. Ji, K. Liu, S. He, and J. Zhao, "Distant supervision for relation extraction with sentence-level attention and entity descriptions," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 31, no. 1, 2017.

[21] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *arXiv preprint arXiv:1703.07737*, 2017.

[22] M. Žitnik and B. Zupan, "Data fusion by matrix factorization," *IEEE transactions on pattern analysis and machine intelligence*, vol. 37, no. 1, pp. 41–53, 2014.

[23] J. Qu, B. Wang, Z. Li, X. Lyu, and Z. Tang, "Understanding multivariate drug-target-diseaseinterdependence via event-graph," in *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2021, pp. 1685–1687.

[24] K. M. Saifuddin, C. May, F. Tanvir, M. I. K. Islam, and E. Akbas, "Seq-hygan: Sequence classification via hypergraph attention network," *ArXiv*, vol. abs/2303.02393, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:257365316

[25] K. M. Saifuddin, B. Bumgardnerr, F. Tanvir, and E. Akbas, "Hygnn: Drug-drug interaction prediction via hypergraph neural network," *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pp. 1503–1516, 2022. [Online]. Available: https://api.semanticscholar.org/CorpusID:250072419

[26] L. Liu, F. Huang, X. Liu, Z. Xiong, M. Li, C. Song, and W. Zhang, "Multi-view contrastive learning hypergraph neural network for drug-microbe-disease association prediction," in *International Joint Conference on Artificial Intelligence*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:260845944

[27] System, "System, D. C. I. 2015. Smiles tutorial."

[28] K. Huang, C. Xiao, L. Glass, and J. Sun, "Explainable substructure partition fingerprint for protein, drug, and more," *NeurIPS Learning Meaningful Representation of Life Workshop*, 2019.

[29] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 173–182.

[30] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," *Advances in neural information processing systems*, vol. 26, 2013.

[31] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, pp. 455–500, 2009.

[32] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" *arXiv preprint arXiv:1810.00826*, 2018.

[33] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio', and Y. Bengio, "Graph attention networks," *ArXiv*, vol. abs/1710.10903, 2018.

[34] W. L. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *NIPS*, 2017.

[35] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous graph transformer," in *Proceedings of the web conference 2020*, 2020, pp. 2704–2710.

[36] C. Moon, C. Jin, X. Dong, S. M. Abrar, W. Zheng, R. Y. Chirkova, and A. Tropsha, "Learning drug-disease-target embedding (ddte) from knowledge graphs to inform drug repurposing hypotheses," *Journal of biomedical informatics*, p. 103838, 2021. [Online]. Available: https://api.semanticscholar.org/CorpusID:235425831

[37] M. Li, X. Cai, S. Xu, and H. Ji, "Metapath-aggregated heterogeneous graph neural network for drug-target interaction prediction," *Briefings in bioinformatics*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:255462340

[38] K. Lawson, "A brief review of the pharmacology of amitriptyline and clinical outcomes in treating fibromyalgia," *Biomedicines*, vol. 5, no. 2, p. 24, 2017.

[39] P. Blier *et al.*, "Neurotransmitter targeting in the treatment of depression," *The Journal of clinical psychiatry*, vol. 74, no. suppl 2, p. 12763, 2013.

[40] C. M. Dempsey, S. M. Mackenzie, A. Gargus, G. Blanco, and J. Y. Sze, "Serotonin (5ht), fluoxetine, imipramine and dopamine target distinct 5ht receptor signaling to modulate caenorhabditis elegans egg-laying behavior," *Genetics*, vol. 169, no. 3, pp. 1425–1436, 2005.

[41] N. S. Philip, L. L. Carpenter, A. R. Tyrka, and L. H. Price, "Nicotinic acetylcholine receptors and depression: a review of the preclinical and clinical literature," *Psychopharmacology*, vol. 212, pp. 1–12, 2010.