

Heterogeneous Causal Metapath Graph Neural Network for Gene-Microbe-Disease Association Prediction

Kexin Zhang¹, Feng Huang¹, Luotao Liu¹, Zhankun Xiong¹, Hongyu Zhang^{1,2},
Yuan Quan^{1,2*} and Wen Zhang^{1,2*}

¹ College of Informatics, Huazhong Agricultural University, Wuhan 430070, China

² Hubei Key Laboratory of Agricultural Bioinformatics, Huazhong Agricultural University, Wuhan 430070, China

{kexinzhang, fhuang233, luotaoliu, xiongzk}@webmail.hzau.edu.cn,
{zhy630, quanyuan, zhangwen}@mail.hzau.edu.cn

Abstract

The recent focus on microbes in human medicine highlights their potential role in the genetic framework of diseases. To decode the complex interactions among genes, microbes, and diseases, computational predictions of gene-microbe-disease (GMD) associations are crucial. Existing methods primarily address gene-disease and microbe-disease associations, but the more intricate triple-wise GMD associations remain less explored. In this paper, we propose a **Heterogeneous Causal Metapath Graph Neural Network (HCMGNN)** to predict GMD associations. HCMGNN constructs a heterogeneous graph linking genes, microbes, and diseases through their pairwise associations, and utilizes six predefined causal metapaths to extract directed causal subgraphs, which facilitate the multi-view analysis of causal relations among three entity types. Within each subgraph, we employ a causal semantic sharing message passing network for node representation learning, coupled with an attentive fusion method to integrate these representations for predicting GMD associations. Our extensive experiments show that HCMGNN effectively predicts GMD associations and addresses association sparsity issue by enhancing the graph's semantics and structure.

1 Introduction

Contemporary human genetics focuses on unraveling the links between genotypes and diseases [Lappalainen and MacArthur, 2021]. Recent research increasingly connects genetic variant-related diseases with human microbes [Hall *et al.*, 2017; Manor *et al.*, 2020; Quan *et al.*, 2023]. Furthermore, there's growing evidence of causal relations among genes, microbes, and diseases, offering insights into complex disease pathogenesis from a multifactorial angle [Luca *et al.*, 2018]. Classical bioinformatics approaches, like multi-omics integration, while used for gene-microbe-disease (GMD) association discovery [Priya *et al.*, 2022; Markowitz *et al.*,

2022], lack efficiency at scale. Machine learning-based methods, proposed as an efficient alternative, typically focus only on pairwise associations, such as gene-disease or microbe-disease associations [Ata *et al.*, 2021; Wen *et al.*, 2021]. Their limitation is not considering the third entity's role in these associations. In contrast, exploring triple-wise GMD associations provides a more comprehensive understanding of biological mechanisms.

Several approaches are applicable for triple-wise association prediction, including tensor decomposition [Wang *et al.*, 2018; Chen and Li, 2019; Chen and Li, 2020], hypergraph neural networks [Cheng *et al.*, 2022a; Liu *et al.*, 2022], and n-ary relational knowledge graph embedding methods [Wen *et al.*, 2016; Fatemi *et al.*, 2020; Liu *et al.*, 2021]. Models based on tensor decomposition, like those proposed by [Wang *et al.*, 2018] and [Chen and Li, 2020], have been used for predicting drug-protein-disease associations. [Liu *et al.*, 2022] introduced a method for anti-cancer drug synergy prediction using a multi-way relation-enhanced hypergraph, while [Cheng *et al.*, 2022a] developed an interactive hypergraph neural network for personalized product search. [Wen *et al.*, 2016] focused on representation learning in n-ary relational knowledge graphs. However, these methods typically treat triple-wise associations as instances of three entities without adequately considering the directional relations that indicate causal interactions among them.

With the advancement of microbiome genome-wide association studies, bidirectional interactions between genes and microbes have garnered attention [Nichols and Davenport, 2021]. Furthermore, causal relations between microbes and diseases have also been validated through causal inference methods such as Mendelian Randomization [Lv *et al.*, 2021]. Consequently, there exist intricate causal modes among genes, microbes, and diseases [Hughes *et al.*, 2020; Xu *et al.*, 2020], as summarized by [Luca *et al.*, 2018]. Illustratively, *NOD2* gene heightens the susceptibility to inflammatory bowel disease through the modulation of *Enterobacteriaceae* abundance [Hall *et al.*, 2017]. Despite the insightful triple-wise associations revealed by these causal modes, modeling such interactions presents inherent challenges. Heterogeneous graphs, key in modeling complex systems, use metapaths to reveal specific connections and hidden patterns

among entities [Shi *et al.*, 2018]. While traditional metapath-based methods on these graphs typically model relations between two nodes, how to extend these methods to encompass triple-wise relations is underexplored. Such an extension has significant potential in GMD associations prediction, as the intricate causal modes among genes, microbes, and diseases can be comprehensively mapped using metapaths, termed causal metapaths, which can be used for learning interactive causal semantic information (more details of the causal metapaths description are presented in the Appendix Section 1).

In our study, we propose a **Heterogeneous Causal Metapath Graph Neural Network (HCMGNN)** to predict GMD associations. HCMGNN constructs a heterogeneous graph that connects genes, microbes, and diseases through their pairwise associations, and utilizes six specific causal metapaths to extract directed causal subgraphs from the heterogeneous graph, which facilitate the multi-view analysis of causal relations among three types of entities (genes, microbes, and diseases). Within each subgraph, we apply a causal semantic sharing message passing network for learning node representations, further augmented by an attentive fusion method for integrating these representations in predicting GMD associations. Extensive experiments demonstrate that HCMGNN achieves satisfactory results. In summary, the main contributions of this work are described as follows:

- To our knowledge, HCMGNN is the first metapath-based heterogeneous graph framework for triple-wise association prediction, which can efficiently capture structural and semantic information for large-scale discovery of potential associations among genes, microbes, and diseases.
- To model the triple-wise relations among genes, microbes, and diseases, we consider six distinct metapaths to represent their interaction modes, which have been observed in the real world. These metapaths serve as a foundation for generating various subgraphs, enabling us to capture a rich array of semantic information.
- Within each subgraph, we employ a causal semantic sharing message passing network for node representation learning, coupled with an attentive fusion method to integrate multi-view node representations for predicting GMD associations.

2 Related Work

2.1 Biomedical Association Mining

Within the biomedical domain, research on disease-related association mining primarily highlights the functional information of genes [Ata *et al.*, 2021]. Recently, several studies have shifted their focus towards elucidating the role of microbes in complex diseases, thereby discovering novel associations between microbes and diseases [Wen *et al.*, 2021]. Consequently, the exploration of associations between diseases and multiple entity types has emerged as a burgeoning field in biomedicine. Here, we summarize the triple-wise association prediction methods in bioinformatics in recent years. Examples include tensor decomposition mod-

els, applied by [Wang *et al.*, 2018] and [Chen and Li, 2020] for predicting drug-protein-disease associations; Hypergraph neural networks, employed by [Liu *et al.*, 2022] to predict anti-cancer drug synergy; N-ary relational knowledge graph embedding methods, utilized by [Wen *et al.*, 2016] to learn representation for predicting n-ary facts. Compared to the above methods, we model the triple-wise associations based on heterogeneous graph framework, and capture higher-order interactions by causal modes among different types of entities.

2.2 Metapath-based HGNNs

Heterogeneous graphs, characterized by diverse node and edge types, exhibit pivotal advancement in modeling complex systems. Recently, heterogeneous graph neural networks (HGNNs) have attracted extensive attention from different research fields due to their great potential in learning and generalizing patterns within heterogeneous structures and have been applied to various downstream tasks, such as node classification [Wang *et al.*, 2019; Wang *et al.*, 2021], recommender systems [Yan *et al.*, 2021; Yang *et al.*, 2021] and bioinformatics [Li *et al.*, 2022; Zhang *et al.*, 2023; Li *et al.*, 2023]. To distill diverse semantic information within a specific context from heterogeneous graph, most existing HGNNs commonly employ metapaths, which refers to a sequence of relations connecting different types of nodes in heterogeneous graph.

Metapath-based HGNNs can be broadly divided into two categories according to their metapath-induced message passing formation: (1) aggregating metapath-induced neighbors [Wang *et al.*, 2019; Liu *et al.*, 2023; Wang *et al.*, 2021; Yan *et al.*, 2023]. Usually, the target nodes are head entities in metapath instances and their metapath-induced neighbors are tail entities. For instance, HAN generates homogeneous graphs based on several symmetric metapaths, where graph attention layers are adopted to update node embeddings [Wang *et al.*, 2019]. (2) aggregating metapath instances [Fu *et al.*, 2020; Zheng *et al.*, 2022]. For example, MAGNN update target nodes by aggregating all symmetric metapath instances with consistent head and tail node types involving the target node [Fu *et al.*, 2020]. Almost all abovementioned traditional metapath-based HGNNs predominantly concentrate on capturing high-order relations within homogeneous node pairs, while neglecting intricate interactions among entities of distinct types.

3 Methodology

3.1 Problem Formulation

Let $\mathcal{N}$, $\mathcal{M}$ and $\mathcal{D}$ denote the set of genes, microbes, and diseases, respectively. Based on the verified pairwise associations among them, i.e. gene-microbe, gene-disease and microbe-disease associations, we can construct a GMD heterogeneous graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, in which $\mathcal{V} = \mathcal{N} \cup \mathcal{M} \cup \mathcal{D}$ is the node set and $\mathcal{E}$ is a set of edges representing the three types of pairwise associations. Each node $v \in \mathcal{V}$ is characterised by a feature vector denoted as $\mathbf{x}_v$. All feature vectors of gene nodes, microbe nodes and disease nodes are compiled into $\mathbf{X}_{\mathcal{N}} \in \mathbb{R}^{|\mathcal{N}| \times F_{\mathcal{N}}}$, $\mathbf{X}_{\mathcal{M}} \in \mathbb{R}^{|\mathcal{M}| \times F_{\mathcal{M}}}$ and $\mathbf{X}_{\mathcal{D}} \in$

$\mathbb{R}^{|\mathcal{D}| \times F_{\mathcal{D}}}$. We consider a triplet $(n, m, d) \in \mathcal{N} \times \mathcal{M} \times \mathcal{D}$ has a triple-wise association only when all three entities have associations with each other. Such a triplet is assigned a label $y = 1$ representing the existence of the triple-wise association. Other triplets are labeled as $y = 0$, which means potential associations undiscovered before. Our goal is to build an effective in-silico model to prioritizes potential triple-wise associations based on the GMD heterogeneous graph.

3.2 Model Architecture

The overall framework of the HCMGNN model is shown in Figure 1. HCMGNN mainly includes: (1) feature transformation; (2) causal subgraph generation; (3) intra-subgraph message passing; (4) inter-subgraph attentive fusion; (5) model training.

3.3 Feature Transformation

In the GMD heterogeneous graph, different node types represent distinct biological entity categories. Different types of node feature vectors are compiled by diverse means of data collection and preprocessing, and in result they have unequal dimension and scale. Therefore, we simply operate a type-specific linear transformation on each type of node features to map all nodes into the same latent space. Let $\mathcal{K}$ be an option from $\mathcal{M}$, $\mathcal{N}$ and $\mathcal{D}$. The type-specific linear transformation is formulated by:

$$\mathbf{h}_v = W_{\mathcal{K}} \cdot \mathbf{x}_v \quad (1)$$

where $W_{\mathcal{K}} \in \mathbb{R}^{F' \times F_{\mathcal{K}}}$ is the learnable weight in the linear transformation for genes, microbes or diseases, $\mathbf{x}_v \in \mathbb{R}^{F_{\mathcal{K}}}$ is node feature vector of node $v \in \mathcal{K}$, and $\mathbf{h}_v \in \mathbb{R}^{F'}$ is the projected feature vector (supplementary details about node features are presented in the Appendix Section 2).

3.4 Causal Subgraph Generation

As mentioned above, the causal modes among genes, microbes, and diseases can be effectively represented using metapaths. To ensure that each metapath conveys the interaction direction among the three entity types, we predefine six metapaths with causal modes, named causal metapaths, encompassing $G-M-D$, $G-D-M$, $D-M-G$, $D-G-M$, $M-D-G$, and $M-G-D$. For example, $G-M-D$ indicates that genes have the potential to influence complex diseases through the modulation of microbial changes [Luca *et al.*, 2018] (more details and underlying implications are presented in Appendix Section 1). Next, we extract six directed subgraphs from the heterogeneous graph $\mathcal{G}$ under the guidance of metapaths to enable the model to learn informative causal context.

Specifically, we consider all edges in $\mathcal{E}$ are bi-directional such that $\mathcal{E}$ contains six types of directed edges corresponding to six kinds of relations. Note that each causal metapath is composed of two relations. Let $\mathcal{E}_r \subset \mathcal{E}$ is the set of edges belonging to relation r . Given a metapath p that consists of two relations r and r' , we can generate a directed subgraph $\mathcal{G}_p = \{\mathcal{V}, \mathcal{E}_p\}$ of $\mathcal{G}$, where $\mathcal{E}_p = \{\mathcal{E}_r, \mathcal{E}_{r'}\}$. Then, the causal subgraph comprises all metapath instances of p and each instance implies causal semantic information that the head node exerts impact on the tail node through the intermediate node.

3.5 Intra-Subgraph Message Passing

In this section, we present a message passing network with causal semantic sharing mechanism to update node embeddings in each subgraph. The key idea is to share the causal semantic information in each metapath instance across its head, intermediate and tail nodes. Formally, we define that $\mathcal{S}_p$ is the set of instances of metapath p , and $\mathcal{S}_p(v) \subset \mathcal{S}_p$ is the set of instances involving node v . Then, the node embedding $\mathbf{h}_v^p$ of v in the view of subgraph $\mathcal{G}_p$ is given by:

$$\begin{aligned} \mathbf{m}_s^p &= \Phi(\{\mathbf{h}_u \mid u \in \mathcal{V}_s^p\}) \\ \mathbf{h}_v^p &= \Psi(\mathbf{h}_v, \{\mathbf{m}_s^p \mid s \in \mathcal{S}_p(v)\}) \end{aligned} \quad (2)$$

where $\mathbf{m}_s^p$ is the embedding of the metapath instance $s \in \mathcal{S}_p$ and $\mathcal{V}_s^p$ contains three nodes in $s \in \mathcal{S}_p$. Obviously, the embedding of metapath instance s serves as the message shared with every node of it. What follows, we introduce detailed implementation for the causal semantic encoder $\Phi(\cdot)$ and the message aggregation function $\Psi(\cdot)$ in our work.

Causal Semantic Encoder

Recall that all metapath instances contain three nodes respectively from $\mathcal{N}$, $\mathcal{M}$ and $\mathcal{D}$, connected by two kinds of relations. To embed causal semantic information of the metapath instance s into $\mathbf{m}_s^p$, $\Phi(\cdot)$ requires not only the fusion of node embeddings but also capability in capturing relational patterns. We adopt a approach proposed by [Sun *et al.*, 2019; Fu *et al.*, 2020] to model the causal relations. Suppose that a metapath instance $s \in \mathcal{S}_p$ is defined by $h \xrightarrow{r_{he}} e \xrightarrow{r_{et}} t$, where h , e and t are the head, intermediate and tail node; and r_{he} and r_{et} are the relations. Then, $\mathbf{m}_s^p$ is obtained by:

$$\mathbf{m}_s^p = \frac{1}{3}((\mathbf{h}_h \odot \mathbf{r}_{he} + \mathbf{h}_e) \odot \mathbf{r}_{et} + \mathbf{h}_t) \quad (3)$$

where $\mathbf{r}_{he}$ and $\mathbf{r}_{et}$ are both learnable parameters for embedding relations, and $\odot$ is Hadamard product.

Message Aggregation

Considering different metapath instances will contribute to the representation of the node in varying extents, here we utilize the graph attention layer [Veličković *et al.*, 2018] to aggregate causal semantics by taking all the metapath instances involving node v as its neighbors. Specifically, we have:

$$\mathbf{h}_v^p = \sigma \left(\sum_{s \in \mathcal{S}_p(v)} \alpha_{sv}^p \cdot \mathbf{m}_s^p \right) \quad (4)$$

where σ is an activation function and α_{sv}^p is the attention value indicating the contribution of semantic message $\mathbf{m}_s^p$ from metapath instance $s \in \mathcal{S}_p(v)$ to node v in the subgraph $\mathcal{G}_p$. The attention value α_{sv}^p for $s \in \mathcal{S}_p(v)$ can be calculated by:

$$\alpha_{sv}^p = \frac{\exp(\text{LeakyReLU}(\mathbf{a}_p^\top \cdot [\mathbf{h}_v \parallel \mathbf{m}_s^p]))}{\sum_{s' \in \mathcal{S}_p(v)} \exp(\text{LeakyReLU}(\mathbf{a}_p^\top \cdot [\mathbf{h}_v \parallel \mathbf{m}_{s'}^p]))} \quad (5)$$

where $\mathbf{a}_p^\top$ is the trainable parameterized attention vector for metapath p , and $\parallel$ represents concatenation operation of vector. To stabilize the learning process and reduce the high variance of the attention coefficient caused by graph heterogeneity, we employ multi-head attention similarly to [Veličković




Figure 1: Overall architecture of HCMGNN. HCMGNN mainly includes: (1) feature transformation; (2) causal subgraph generation; (3) intra-subgraph message passing; (4) inter-subgraph attentive fusion; (5) model training.

et al., 2018] that concatenates multi-view node embeddings from multiple attention heads.

3.6 Inter-Subgraph Attentive Fusion

Now that we obtain multiple views of node embeddings from six subgraphs generated before, a type-specific attentive fusion operator is exploited to merge the multi-view information. Concretely, suppose that $\mathcal{P}$ is the set of our predefined metapaths and $p \in \mathcal{P}$, the attention coefficient $e_{\mathcal{K}}^p$ assigned for each node in $\mathcal{K}$ in the subgraph $\mathcal{G}_p$ is computed by:

$$e_{\mathcal{K}}^p = \frac{1}{|\mathcal{K}|} \sum_{v \in \mathcal{K}} \mathbf{q}_{\mathcal{K}}^{\top} \cdot \tanh(\mathbf{W}'_{\mathcal{K}} \cdot \mathbf{h}_v^p + \mathbf{b}_{\mathcal{K}}) \quad (6)$$

where $\mathbf{q}_{\mathcal{K}}^{\top}$ is a parameterized query vector for node type $\mathcal{K}$, and $\mathbf{W}'_{\mathcal{K}}$ and $\mathbf{b}_{\mathcal{K}}$ are the learnable weight matrix and bias vector respectively.

Note that the attention coefficient is shared across all nodes in $\mathcal{K}$, which measures an equal contribution of causal semantic information behind p to all node embeddings of node type $\mathcal{K}$. Then we further normalize these attention coefficients with a softmax layer.

$$\beta_{\mathcal{K}}^p = \frac{\exp(e_{\mathcal{K}}^p)}{\sum_{p' \in \mathcal{P}} \exp(e_{\mathcal{K}}^{p'})} \quad (7)$$

Finally, we weightedly sum all views of embeddings of node $v \in \mathcal{K}$ using the weights given by the normalized attentions to obtain the final node embedding $\mathbf{z}_v$.

$$\mathbf{z}_v = \sum_{p \in \mathcal{P}} \beta_{\mathcal{K}}^p \cdot \mathbf{h}_v^p \quad (8)$$

3.7 Model Training

Through the modules introduced above, we obtain highly-abstracted node representations for genes, microbes, and diseases, denoted as $\mathbf{z}_n$ for node $n \in \mathcal{N}$, respectively $\mathbf{z}_m$ for node $m \in \mathcal{M}$ and $\mathbf{z}_d$ for node $d \in \mathcal{D}$. After that, we feed the concatenation of $\mathbf{z}_n$, $\mathbf{z}_m$ and $\mathbf{z}_d$ into a MLP layer followed by a Sigmoid layer for calculating the probability of the existence of GMD association among the triplet (n, m, d) .

$$\hat{y} = \text{Sigmoid}(\text{MLP}(\mathbf{z}_n \parallel \mathbf{z}_m \parallel \mathbf{z}_d)) \quad (9)$$

Then, we adopt the loss function used in [Wang *et al.*, 2022] to optimize our model, which is formulated as:

$$L = (1 - \gamma) \|\mathbf{y} \odot (\mathbf{y} - \hat{\mathbf{y}})\|^2 + \gamma \|(1 - \mathbf{y}) \odot (\mathbf{y} - \hat{\mathbf{y}})\|^2 \quad (10)$$

where γ is a balance coefficient, $\|\cdot\|^2$ is the L_2 norm, $\hat{\mathbf{y}}$ is a vector storing prediction scores for all triplets in training set and $\mathbf{y}$ represents the corresponding ground truth labels.

4 Results

4.1 Experimental Settings

Dataset

We construct a dataset of gene-microbe-disease associations using several public databases. The gene-microbe associations are collected from GIMICA [Tang *et al.*, 2021] and gut-MGene [Cheng *et al.*, 2022b]. The microbe-disease associations are derived from MicroPhenoDB [Yao *et al.*, 2020], GutMDdisorder [Cheng *et al.*, 2020] and Peryton [Skoufos *et al.*, 2021]. And the gene-disease associations are obtained from the DisGeNET [Piñero *et al.*, 2016]. Then we merge the gene-microbe associations, microbe-disease associations,

Type	Method	Hit@1	Hit@3	Hit@5	NDCG@1	NDCG@3	NDCG@5	MRR
Machine Learning	RF	0.4880	0.7996	0.8964	0.4880	0.6713	0.7115	0.6614
	MLP	0.5321	0.8277	0.9122	0.5321	0.7060	0.7409	0.6937
	XGBoost	0.5719	0.8339	0.9132	0.5719	0.7269	0.7598	0.7184
Tensor Decomposition	Tucker	0.7500	0.8021	0.8271	0.7500	0.7805	0.7907	0.7890
	CP	0.7914	0.8403	0.8578	0.7914	0.8208	0.8278	0.8271
	NeurTN	0.5424	0.8391	0.9168	0.5424	0.7169	0.7493	0.7017
Knowledge Graph Embedding	mTransH	0.6943	0.8326	0.8779	0.6943	0.7757	0.7944	0.7771
	RAM	0.7019	0.8357	0.8844	0.7019	0.7806	0.8007	0.7830
Graph Neural Network	HAN	0.6771	0.8731	0.9122	0.6771	0.7934	0.8096	0.7819
	MAGNN	0.7432	0.8863	0.9362	0.7432	0.8262	0.8486	0.8256
	HypergraphSynergy	0.7853	0.9326	0.9563	0.7853	0.8773	0.8863	0.8656
	HCMGNN	0.7947	0.9417	0.9641	0.7947	0.8827	0.8923	0.8712

Table 1: Performances of HCMGNN and baselines evaluated by Hit@n, NDCG@n and MRR in 5-fold CV.

and gene-disease associations into triplets $\langle \text{gene-microbe-disease} \rangle$. After filtering, we obtain a total of 3431 triplets, involving 301 genes, 176 microbes, and 153 diseases.

Baselines

In order to verify the effectiveness of our model, we compare HCMGNN with baselines as follows (more details of baselines can be found in Appendix Section 3):

- Random Forest (RF), multilayer perception (MLP) and XGBoost are classic machine learning methods that utilize concatenated features of entities for prediction.
- Tucker and CP are tensor decomposition models, which utilize multi-dimensional tensors to represent and reconstruct the high-order relations [Kolda and Bader, 2009].
- NeurTN is a neural tensor model combining tensor algebra and deep neural network that captures associations among drugs, targets, and diseases [Chen and Li, 2020].
- mTransH and RAM are n-ary relational knowledge graph embedding methods. mTransH is a translational distance model to handle multiple relation types [Wen *et al.*, 2016]. RAM introduce a role-aware modeling to learn the semantic relatedness among roles for completing n-ary facts in knowledge graph [Liu *et al.*, 2021].
- HAN and MAGNN are metapath-based heterogeneous graph neural network models for node classification and link prediction [Wang *et al.*, 2019; Fu *et al.*, 2020].
- HypergraphSynergy is a multi-way relation-enhanced hypergraph representation learning method for predicting drug-drug-cell line associations [Liu *et al.*, 2022].

Implementation Details

We randomly split the dataset into 90% cross-validation (CV) set and 10% independent test set. Then we perform the 5-fold CV on the CV set to train the model and optimize the hyper-parameters, and evaluate model performance on the independent test set. In each CV, we randomly select four out of five positive samples in the CV set and generate an equal number of negative samples as the training set. When verifying and testing the model, we generate 30 corresponding

negative samples for each test triplet, so as to sort the prediction scores of all samples. To evaluate the performance of the model, we focus on three evaluation metrics commonly used in recommender systems and knowledge graph embeddings model [Chen and Li, 2020]: hit ratio (Hit@n) and normalized discounted cumulative gain (NDCG@n) with $n=1, 3, 5$ and mean reciprocal rank (MRR). For training our model, we set the balance coefficient of the loss function to 0.7, employ the Adam optimizer with a learning rate set to 0.005 to optimize the model, and adopt early stopping mechanism with a patience of 50 to terminate the training early (the hyper-parameter sensitivity analysis are provided in Appendix Section 4). Source code and dataset of HCMGNN are available online at <https://github.com/zkxinxin/HCMGNN>.

4.2 Performance Comparison

In this section, We compare HCMGNN with baseline methods in 5-fold CV and independent test experiments. As shown in Table 1, HCMGNN outperforms other baselines under all evaluation metrics on the CV set, demonstrating its effectiveness in GMD association prediction. We also have the following observations: (1) Compared with RF, MLP, and XGBoost, which only consider entities feature information, HCMGNN achieves a significant performance improvement, which indicates that the rich structural and semantic information learned by heterogeneous graph enhances the capability of node representation. (2) Compared with Tucker, CP, and NeurTN, which model higher-order associations through multi-dimensional tensors, HCMGNN considers the relational directionality modeled by causal metapaths, making it more conducive to comprehensively learning triple-wise interaction information than tensors. (3) mTransH and RAM, which model n-ary relational knowledge graphs by translation distance and role semantics, respectively, surpass the above methods on most metrics. HCMGNN achieves further improvement based on them, which may be due to that HCMGNN jointly considers the topological neighborhood structure of heterogeneous graph and triple-wise interaction information, which is beneficial to aggregate more comprehensive information for nodes. (4) MAGNN, which learn node embeddings by aggregating metapath instances




Figure 2: The t-SNE visualization of four models.

in heterogeneous graph, performs superior to HAN. However, HCMGNN still make improvement over them, which implies that compared with the traditional metapath-based HGNNs, HCMGNN is better able to model the interaction of multi-view triple-wise interactions by utilizing causal metapaths and update node representations in a shared semantics way comprehensively. (5) HypergraphSynergy, which utilizes hypergraphs to represent triple-wise associations and exhibits better performance than other baselines. However, HCMGNN still outperforms it across all metrics, suggesting that metapaths, as an impactful means to connect high-order relations, can effectively learn representations of nodes from multiple causal views. These results provide valuable implications for addressing the triple-wise association problems.

In addition, according to the performance of HCMGNN and baselines in independent test set (experiment result can be found in Appendix Section 5), we find that the general trend of the results is consistent with the CV set, HCMGNN outperforms all baselines under every evaluation metrics and exceeds a state-of-the-art method HypergraphSynergy by 2.0%, 1.8%, and 1.6% in Hit@1, Hit@3 and MRR, which verifies its superiority in predicting the GMD triple-wise association.

To intuitively verify the powerful ability of HCMGNN to learn GMD representation, we performed t-SNE visualization of triplets embeddings learned by HCMGNN and several state-of-the-art baselines from independent test set. Figure 2 shows the embedding visualization results of positive and negative samples of HCMGNN, hypergraph neural network (HypergraphSynergy) and tensor decomposition with deep neural networks (NeurTN) and n-ary relational knowl-




Figure 3: Performances of HCMGNN and variants in ablation study.

edge graph embedding model (mTransH). It can be seen that HCMGNN can better distinguish positive samples from negative samples compared with the comparison methods and achieves the best silhouette score of 0.5557, which indicates that the GMD triplets embeddings learned by HCMGNN are sufficiently discriminative and effective for prediction task.

4.3 Ablation Study

To validate the necessity of each component in our model, we further compared several different variants of HCMGNN in independent test experiment.

- **HCMGNN without the intra-subgraph sharing message passing (w/o MP) (w/o MP-i)** updates node embeddings by aggregating metapath-induced neighbors information instead of triple-wise semantic information represented by metapath instances. **(w/o MP-ii)** updates node embeddings by aggregating the causal semantic information of each metapath instance only to its head/tail node, instead of sharing information to each node of the instance. **(w/o MP-iii)** defines non-causal metapaths (e.g., $G-M-D-M-G$) with symmetric structure to embody triple-wise information interaction, and updates node embeddings by aggregating non-causal semantics to the head/tail node of each instance.
- **HCMGNN without defined triple-wise interactive metapaths (w/o TM)** defines metapaths of length 2 containing two node types that do not have triple-wise information interactions (e.g., $G-M$).
- **HCMGNN without inter-subgraph attentive fusion (w/o AF)** adopts the average strategy instead of the attention mechanism to achieve node aggregation between subgraphs.
- **HCMGNN without biological features of nodes (w/o BF)** employs one-hot encoding to replace biological attributes as initial features of nodes.

As shown in Figure 3, the performances of all variants of HCMGNN significantly decline, which indicates that these components all contribute to the GMD associations prediction. Specifically, we have the following observations: (1) Compared with HCMGNN (w/o MP-ii) and HCMGNN (w/o MP-iii), HCMGNN (w/o MP-i) achieves a large performance drop, which shows that only aggregating metapath-induced neighbors information is insufficient for fully capture higher-order relations in heterogeneous graph. Next, the differences result between HCMGNN (w/o MP-ii) and HCMGNN (w/o MP-iii) suggests that causal metapaths are more helpful in representing triple-wise relations from multiple views than non-causal metapaths with symmetric structures. In addition, HCMGNN outperforms HCMGNN (w/o MP-ii) and HCMGNN (w/o MP-iii), which indicates that our designed intra-subgraph sharing message passing module is effective for the GMD association prediction. The triple-wise semantic information represented by metapaths cannot simply be aggregated only to the head/tail nodes connecting metapath instances. (2) HCMGNN (w/o TM) has the worst predictive effect across almost all evaluation metrics, demonstrating that metapaths containing only two types of entities are difficult to reflect the entangled interaction of the three entities, which further reveals the capability of our defined causal metapaths to mine high-order relations from heterogeneous graph. (3) HCMGNN (w/o AF) shows performance decline on all evaluation metrics, indicating that different subgraphs divided by causal metapath contribute differently to the model predictions, and inter-subgraph attentive fusion helps to learn multi-view nodes representation comprehensively. (4) HCMGNN (w/o BF) also exhibits a drop in all metrics, suggesting that biological features of genes, microbes, and diseases provide useful domain knowledge for GMD association prediction.

4.4 Why Do Causal Metapaths Improve Predictive Performance?

In this section, we further explore why causal metapaths contribute to the improvement of triple-wise associations predictive performance. Given that metapaths are able to capture high-order relations in heterogeneous graph, we speculate that triple-wise semantic information mined by causal metapaths may have a better impact on low-degree nodes with fewer neighbors (called sparse nodes), which helps to alleviate association sparsity issue caused by sparse nodes. Thus, we calculate the degrees of nodes in heterogeneous graph, and average the degrees for the nodes contained in each triplet to obtain the average node degree (N) for each triplet in independent test set. Next, we divide the test sets into twelve groups by counting triplet numbers under different average node degree interval and evaluate the performances of HCMGNN and HCMGNN without the intra-subgraph sharing message passing (w/o MP) by Hit@1 in the same group.

As shown in Figure 4, We can observe a consistent performance decline across all models as the average node degree of triplets decreased. This trend indicates that most models exhibit limited capability in aggregating information for sparse nodes. In addition, we also have the following observations: (1) Compared with HCMGNN (w/o MP-i), the predictive performance of HCMGNN is significantly improved




Figure 4: Performances of HCMGNN and w/o MP in test triplets with different average node degree (N) distributions. Each bar shows the result on the triplets in the average node degree interval $(0, N]$.

in all degree distributions, suggesting that aggregating triple-wise semantic information from multiple causal views is more conducive to learning sparse nodes representations than metapath-induced neighbors. (2) Compared with HCMGNN (w/o MP-ii) and HCMGNN (w/o MP-iii), the performance gap between HCMGNN and them has notably widened with the decrease of the degree value (see the red area change in Figure 4), which indicates that HCMGNN has superior predictive performance in low-degree triplets, suggesting that causal semantic sharing mechanism enhances the aggregation of information for sparse nodes, effectively alleviating the association sparsity issue in triple-wise association prediction (more details and additional experiments can be found in Appendix Section 6). Based on this result, we infer that causal metapaths have the potential to augment the original semantics and structure of heterogeneous graph, thereby improving the overall predictive performance of triple-wise associations by enhancing the predictive capability for sparse associations.

5 Conclusion

In this paper, we propose a heterogeneous causal metapath graph neural network for GMD association prediction, named HCMGNN. HCMGNN constructs a heterogeneous graph consisting of genes, microbes, and diseases through their pairwise associations, and then extracts multiple subgraphs according to the patterns of causal metapaths for data augmentation. Next, intra-subgraph message passing with causal semantic sharing mechanism and inter-subgraph attentive fusion are implemented to learn multi-view node representations. Experimental results show that the superior performances of HCMGNN over baselines, and causal metapaths can indeed augment the inherent semantics and structure of heterogeneous graph to alleviate association sparsity issue.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62072206, 62372204, 61772381, 62102158, 32300545); National Key R&D Program of China (2022YFA1304104); Fundamental Research Funds for the Central Universities (2662021JC008, 2662022JC004); Huazhong Agricultural University Scientific & Technological Self-innovation Foundation. The funders have no role in study design, data collection, data analysis, data interpretation, or writing of the manuscript.

References

- [Ata *et al.*, 2021] Sezin Kircali Ata, Min Wu, Yuan Fang, Le Ou-Yang, Chee Keong Kwoh, and Xiao-Li Li. Recent advances in network-based methods for disease gene prediction. *Briefings in bioinformatics*, 22(4):bbaa303, 2021.
- [Chen and Li, 2019] Huiyuan Chen and Jing Li. Modeling relational drug-target-disease interactions via tensor factorization with multiple web sources. In *The World Wide Web Conference*, pages 218–227, 2019.
- [Chen and Li, 2020] Huiyuan Chen and Jing Li. Learning data-driven drug-target-disease interaction via neural tensor network. In *International joint conference on artificial intelligence (IJCAI)*, 2020.
- [Cheng *et al.*, 2020] Liang Cheng, Changlu Qi, He Zhuang, Tongze Fu, and Xue Zhang. gutmdisorder: a comprehensive database for dysbiosis of the gut microbiota in disorders and interventions. *Nucleic Acids Research*, 48(D1):D554–D560, 2020.
- [Cheng *et al.*, 2022a] Dian Cheng, Jiawei Chen, Wenjun Peng, Wenqin Ye, Fuyu Lv, Tao Zhuang, Xiaoyi Zeng, and Xiangnan He. Ihgnn: Interactive hypergraph neural network for personalized product search. In *Proceedings of the ACM Web Conference 2022*, pages 256–265, 2022.
- [Cheng *et al.*, 2022b] Liang Cheng, Changlu Qi, Haixiu Yang, Minke Lu, Yiting Cai, Tongze Fu, Jialiang Ren, Qu Jin, and Xue Zhang. gutmgene: a comprehensive database for target genes of gut microbes and microbial metabolites. *Nucleic acids research*, 50(D1):D795–D800, 2022.
- [Fatemi *et al.*, 2020] Bahare Fatemi, Perouz Taslakian, David Vazquez, and David Poole. Knowledge hypergraphs: Prediction beyond binary relations. In *International joint conference on artificial intelligence (IJCAI)*, 2020.
- [Fu *et al.*, 2020] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020*, pages 2331–2341, 2020.
- [Hall *et al.*, 2017] Andrew Brantley Hall, Andrew C Tolonen, and Ramnik J Xavier. Human genetic variation and the gut microbiome in disease. *Nature Reviews Genetics*, 18(11):690–699, 2017.
- [Hughes *et al.*, 2020] David A Hughes, Rodrigo Bacigalupe, Jun Wang, Malte C Rühlemann, Raul Y Tito, Gwen Falony, Marie Joossens, Sara Vieira-Silva, Liesbet Henckjaerts, Leen Rymenans, et al. Genome-wide associations of human gut microbiome variation and implications for causal inference analyses. *Nature microbiology*, 5(9):1079–1087, 2020.
- [Kolda and Bader, 2009] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [Lappalainen and MacArthur, 2021] Tuuli Lappalainen and Daniel G MacArthur. From variant to function in human disease genetics. *Science*, 373(6562):1464–1468, 2021.
- [Li *et al.*, 2022] Zhengwei Li, Tangbo Zhong, Deshuang Huang, Zhu-Hong You, and Ru Nie. Hierarchical graph attention network for mirna-disease association prediction. *Molecular Therapy*, 30(4):1775–1786, 2022.
- [Li *et al.*, 2023] Mei Li, Xiangrui Cai, Sihan Xu, and Hua Ji. Metapath-aggregated heterogeneous graph neural network for drug–target interaction prediction. *Briefings in Bioinformatics*, 24(1):bbac578, 2023.
- [Liu *et al.*, 2021] Yu Liu, Quanming Yao, and Yong Li. Role-aware modeling for n-ary relational knowledge bases. In *Proceedings of the Web Conference 2021*, pages 2660–2671, 2021.
- [Liu *et al.*, 2022] Xuan Liu, Congzhi Song, Shichao Liu, Menglu Li, Xionghui Zhou, and Wen Zhang. Multi-way relation-enhanced hypergraph representation learning for anti-cancer drug synergy prediction. *Bioinformatics*, 38(20):4782–4789, 2022.
- [Liu *et al.*, 2023] Yanbei Liu, Lianxi Fan, Xiao Wang, Zhi-tao Xiao, Shuai Ma, Yanwei Pang, and Jerry Chun-Wei Lin. Hgber: Heterogeneous graph neural network with bidirectional encoding representation. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–12, 2023.
- [Luca *et al.*, 2018] Francesca Luca, Sonia S Kupfer, Dan Knights, Alexander Khoruts, and Ran Blekhanman. Functional genomics of host–microbiome interactions in humans. *Trends in Genetics*, 34(1):30–40, 2018.
- [Lv *et al.*, 2021] Bo-Min Lv, Yuan Quan, and Hong-Yu Zhang. Causal inference in microbiome medicine: principles and applications. *Trends in microbiology*, 29(8):736–746, 2021.
- [Manor *et al.*, 2020] Ohad Manor, Chengzhen L Dai, Sergey A Kornilov, Brett Smith, Nathan D Price, Jennifer C Lovejoy, Sean M Gibbons, and Andrew T Magis. Health and disease markers correlate with gut microbiome composition across thousands of people. *Nature communications*, 11(1):5206, 2020.
- [Markowitz *et al.*, 2022] Robert H George Markowitz, Abigail Leavitt LaBella, Mingjian Shi, Antonis Rokas, John A Capra, Jane F Ferguson, Jonathan D Mosley, and Seth R Bordenstein. Microbiome-associated human genetic variants impact phenome-wide disease risk. *Proceedings of*

- the National Academy of Sciences*, 119(26):e2200551119, 2022.
- [Nichols and Davenport, 2021] Robert G Nichols and Emily R Davenport. The relationship between the gut microbiome and host gene expression: a review. *Human genetics*, 140(5):747–760, 2021.
- [Piñero et al., 2016] Janet Piñero, Àlex Bravo, Núria Queralt-Rosinach, Alba Gutiérrez-Sacristán, Jordi Deu-Pons, Emilio Centeno, Javier García-García, Ferran Sanz, and Laura I Furlong. Disgenet: a comprehensive platform integrating information on human disease-associated genes and variants. *Nucleic acids research*, page gkw943, 2016.
- [Priya et al., 2022] Sambhawa Priya, Michael B Burns, Tonya Ward, Ruben AT Mars, Beth Adamowicz, Eric F Lock, Purna C Kashyap, Dan Knights, and Ran Blekhan. Identification of shared and disease-specific host gene-microbiome associations across human diseases using multi-omic integration. *Nature microbiology*, 7(6):780–795, 2022.
- [Quan et al., 2023] Yuan Quan, Ke-Xin Zhang, and Hong-Yu Zhang. The gut microbiota links disease to human genome evolution. *Trends in Genetics*, 39(6):451–461, 2023.
- [Shi et al., 2018] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and S Yu Philip. Heterogeneous information network embedding for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 31(2):357–370, 2018.
- [Skoufos et al., 2021] Giorgos Skoufos, Filippos S Kardaras, Athanasios Alexiou, Ioannis Kavakiotis, Anastasia Lambropoulou, Vasiliki Kotsira, Spyros Tastsoglou, and Artemis G Hatzigeorgiou. Peryton: a manual collection of experimentally supported microbe-disease associations. *Nucleic acids research*, 49(D1):D1328–D1333, 2021.
- [Sun et al., 2019] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*, 2019.
- [Tang et al., 2021] Jing Tang, Xianglu Wu, Minjie Mou, Chuan Wang, Lidan Wang, Fengcheng Li, Maiyuan Guo, Jiayi Yin, Wenqin Xie, Xiaona Wang, et al. Gimica: host genetic and immune factors shaping human microbiota. *Nucleic Acids Research*, 49(D1):D715–D722, 2021.
- [Veličković et al., 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2018.
- [Wang et al., 2018] Ran Wang, Shuai Li, Man Hon Wong, and Kwong Sak Leung. Drug-protein-disease association prediction and drug repositioning based on tensor decomposition. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 305–312, 2018.
- [Wang et al., 2019] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The world wide web conference*, pages 2022–2032, 2019.
- [Wang et al., 2021] Xiao Wang, Nian Liu, Hui Han, and Chuan Shi. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 1726–1736, 2021.
- [Wang et al., 2022] Hongzhun Wang, Feng Huang, Zhankun Xiong, and Wen Zhang. A heterogeneous network-based method with attentive meta-path extraction for predicting drug-target interactions. *Briefings in Bioinformatics*, 23(4):bbac184, 2022.
- [Wen et al., 2016] Jianfeng Wen, Jianxin Li, Yongyi Mao, Shini Chen, and Richong Zhang. On the representation and embedding of knowledge bases beyond binary relations. In *International joint conference on artificial intelligence (IJCAI)*, page 1300–1307, 2016.
- [Wen et al., 2021] Zhongqi Wen, Cheng Yan, Guihua Duan, Suning Li, Fang-Xiang Wu, and Jianxin Wang. A survey on predicting microbe-disease associations: biological data and computational methods. *Briefings in Bioinformatics*, 22(3):bbaa157, 2021.
- [Xu et al., 2020] Fengzhe Xu, Yuanqing Fu, Ting-yu Sun, Zengliang Jiang, Zelei Miao, Menglei Shuai, Wanglong Gou, Chu-wen Ling, Jian Yang, Jun Wang, et al. The interplay between host genetics and the gut microbiome reveals common and distinct microbiome features for complex human diseases. *Microbiome*, 8(1):1–14, 2020.
- [Yan et al., 2021] Surong Yan, Haosen Wang, Yixiao Li, Yuan Zheng, and Long Han. Attention-aware metapath-based network embedding for hin based recommendation. *Expert Systems with Applications*, 174:114601, 2021.
- [Yan et al., 2023] Yeyu Yan, Chao Li, Yanwei Yu, Xiangju Li, and Zhongying Zhao. Osgnn: Original graph and sub-graph aggregated graph neural network. *Expert Systems with Applications*, 225:120115, 2023.
- [Yang et al., 2021] Haoran Yang, Hongxu Chen, Lin Li, S Yu Philip, and Guandong Xu. Hyper meta-path contrastive learning for multi-behavior recommendation. In *2021 IEEE International Conference on Data Mining (ICDM)*, pages 787–796, 2021.
- [Yao et al., 2020] Guocai Yao, Wenliang Zhang, Minglei Yang, Huan Yang, Jianbo Wang, Haiyue Zhang, Lai Wei, Zhi Xie, and Weizhong Li. Microphenodb associates metagenomic data with pathogenic microbes, microbial core genes, and human disease phenotypes. *Genomics, proteomics & bioinformatics*, 18(6):760–772, 2020.
- [Zhang et al., 2023] Ran Zhang, Zhanjie Wang, Xuezhi Wang, Zhen Meng, and Wenjuan Cui. Mhtan-dti: Metapath-based hierarchical transformer and attention network for drug-target interaction prediction. *Briefings in Bioinformatics*, 24(2):bbad079, 2023.
- [Zheng et al., 2022] Jiayin Zheng, Juanyun Mai, and Yanlong Wen. Explainable session-based recommendation with meta-path guided instances and self-attention mechanism. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2555–2559, 2022.