

GraphRARE: Reinforcement Learning Enhanced Graph Neural Network with Relative Entropy

Tianhao Peng[†], Wenjun Wu^{†*}, Haitao Yuan^{‡*}, Zhifeng Bao[§], Zhao Pengrui[†], Xin Yu[†],
Xuetao Lin[†], Yu Liang[¶], Yanjun Pu[†]

[†]Beihang University, China [‡]Nanyang Technological University, Singapore

[§]RMIT University, Australia [¶]Beijing University of Technology, China

[†]{pengtianhao, wwj09315, zhaopengrui, nlsdeyuxin, xtlin, buaapyj}@buaa.edu.cn

[‡]haitao.yuan@ntu.edu.sg, [§]zhifeng.bao@rmit.edu.au, [¶]yuliang@bjut.edu.cn

Abstract—Graph neural networks (GNNs) have shown advantages in graph-based analysis tasks. However, most existing methods have the homogeneity assumption and show poor performance on heterophilic graphs, where the linked nodes have dissimilar features and different class labels, and the semantically related nodes might be multi-hop away. To address this limitation, this paper presents GraphRARE, a general framework built upon node relative entropy and deep reinforcement learning, to strengthen the expressive capability of GNNs. An innovative node relative entropy, which considers node features and structural similarity, is used to measure mutual information between node pairs. In addition, to avoid the sub-optimal solutions caused by mixing useful information and noises of remote nodes, a deep reinforcement learning-based algorithm is developed to optimize the graph topology. This algorithm selects informative nodes and discards noisy nodes based on the defined node relative entropy. Extensive experiments are conducted on seven real-world datasets. The experimental results demonstrate the superiority of GraphRARE in node classification and its capability to optimize the original graph topology.

Index Terms—Graph Neural Networks, Relative Entropy, Deep Reinforcement Learning, Node Classification

I. INTRODUCTION

Graph data structures have been widely used in many real-world scenarios such as database management system [14], [20], [37], knowledge graphs [53], [56], recommendation systems [23], [43], [55], [58], and traffic forecasting [47]–[49]. In recent years, graph neural networks (GNNs) have exhibited advantages in numerous graph-based analytic tasks, including node classification, edge prediction, graph classification, and graph clustering [57].

Most GNNs are based on a message-passing neural network (MPNN) framework, which aggregates features of neighboring nodes. Such a design of node message aggregation is effective under the assumption of homophily, which requires that adjacent nodes contain similar features or belong to the same class labels [57]. Meanwhile, heterophilic graphs have also been widespread in the real world [10], [12], [21], [22], [32], [51], [59]. In these graphs, linked nodes have different features and class labels, whereas semantically related nodes can be multi-hop away. For instance, different types of amino acids are more likely to connect together in protein structures,

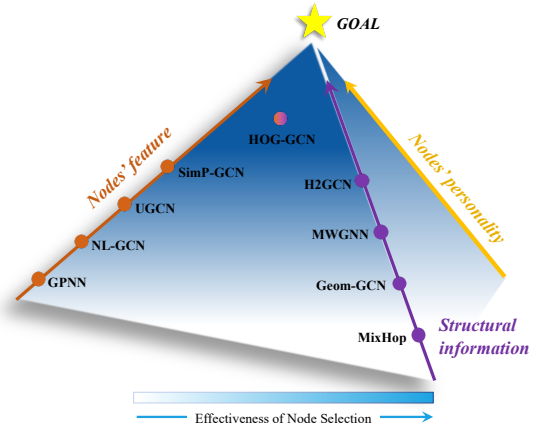


Fig. 1: The visual illustration of related works on heterophilic graph neural networks. The color shades indicate the effectiveness of node selection in reconstructing the original graph topology, with darker shades representing higher effectiveness.

fraudsters are more likely to build connections with customers instead of other fraudsters in online purchasing networks. Therefore, local neighbors of the original graph topology fail to capture informative nodes at a long distance, which may introduce noises, resulting in poor MPNN performance [13]. Thus, capturing and extracting important features from distant but informative nodes can enhance the performance of MPNNs [1], [16], [52], [54], [57], [60].

In recent years, extensive efforts have been made to exploit remote nodes' information under the heterophily setting. The main idea is to optimize the graph topology based on the homophily assumption that message aggregation mechanism works the best in homophilic graphs [4], [6], [7], [17], [31], [44]–[46]. It is widely known that there are two steps to optimize the original topology and reconstruct graphs. The first step is to select an appropriate metric to measure the remote nodes' importance as accurately as possible. The second step is to connect the ego node and a small number of remote nodes based on the importance degree. Existing studies mainly leverage inherent graph features, such as node features and structure similarity, to select relative nodes with a constant number, and then focus on the optimization of the first step.

*Both Wenjun Wu and Haitao Yuan are the corresponding authors.

Unfortunately, the graph reconstruction in the second step is often neglected. As Figure 1 shows, the goal of GNN in heterophilic graphs is to make full use of both the node feature and structural feature of the graph, and consider the personality of each node to improve the accuracy of selecting important nodes. UGCN [16], NL-GNN [25], SimP-GCN [17] and GPNN [45] mainly utilize feature similarity between nodes as a metric to reconstruct the neighbor set. GBK-GNN [4] utilizes kernels to measure the feature similarity between node pairs to guide the aggregation operation. On the contrary, other studies consider the topology of graphs. For example, MixHop [1] and H2GCN [60] utilize the multi-hop node information in each convolution layer. Geom-GCN [31] maps the original graph to a latent space and defines the geometric relationship as a criterion to reconstruct the original graph topology. MWGNN [27] considers the structural similarity to select remote nodes. However, the aforementioned methods rarely consider node features and structural information at the same time, and hence fail to make full use of the rich information of graphs. Hence, Polar-GNN [6] and HOG-GCN [42] consider both node features and structure to guide the aggregation operation of GNNs. Polar-GNN attempts to model pairwise similarities and dissimilarities of nodes, but it fails to capture higher-order structural information or make full use of it. HOG-GCN uses the label propagation technique to estimate the homophily degree matrix, which is affected by the number of labeled nodes, hence making it difficult to fully utilize the structural information. Regarding the optimization for the original graph topology, existing studies do not design appropriate methods for optimization but simply set to hyper-parameters. Examples like UGCN [16], SimP-GCN [17], and MI-GCN [38] select the *top-k* most similar remote nodes to construct new neighboring node sets or remove the *top-d* most dissimilar ones from neighboring sets. These methods rely on prior knowledge or require a huge amount of human efforts to tune the hyper-parameters, highlighting the demand for improved methods in the current approach.

In summary, the aforementioned methods cannot fill the gap in the field of heterophilic graphs, and the major challenges can be summarized as follows:

- *Appropriate Metric.* The key problem in extending higher-hop neighbors is to determine an appropriate metric for measuring the node importance when extending higher-hop neighbors. Most existing studies leverage inherent structural information directly from the original graphs to measure node importance. However, **how to integrate both node features and structural information** to reconstruct the original graph topology and enhance the expressive power of GNNs remains an ongoing challenge.
- *Personality of nodes.* During the process of reconstructing an enhanced graph from the original graph, it is often necessary to set the hyper-parameters, k and d , where we need to select *top-k* important nodes from the higher-order neighbors as the first-order neighbors or delete *top-d* nodes from the first-order neighbors for each node. However, **the values of k**

and d ought to vary across individual nodes, especially in the case of heterogeneous graphs. Therefore, it becomes another challenge to consider the individual characteristics of nodes and set different values of the hyper-parameters.

To address the above challenges, we propose a novel approach, **GraphRARE**¹(Reinforcement leArning enhanced Graph Neural Network with Relative Entropy), specifically designed for heterophilic graphs. GraphRARE aims to adaptively optimize the graph topology while leveraging the advantages of advanced Graph Neural Networks (GNNs), which is critical in the field of database systems with inherently graph-like structures [14], [20], [37], [56].

At first, we define a node-aware relative entropy to measure node importance. In particular, this entropy is composed of node feature entropy and node structural entropy, which are designed to capture node features and structural information, respectively. Therefore, the proposed relative entropy is an appropriate metric, **addressing the first challenge**. Next, for each node, we calculate the node relative entropy between the ego node and its remote neighbors. Subsequently, we construct a descending-order node sequence based on the importance of remote neighbor nodes. **To address the second challenge**, we consider the personality of nodes and employ a Deep Reinforcement Learning (DRL)-based algorithm to determine the values of hyper-parameters k and d for each node. At last, we jointly train the DRL-based algorithm and an existing advanced GNN model, enabling the optimization of the graph topology and facilitating the downstream task (i.e., node classification in this paper).

The main technical contributions of this study can be summarized as follows:

- A node relative entropy is defined to measure the similarity between nodes based on their structure and features, which enhances the application of node entropy theory in the domain of graph data (Sec. IV-A).
- The node personality is captured when enhancing the graph topology, where we use a well-designed reinforcement learning model to set different hyper-parameter values for different nodes (Sec. IV-B).
- The GraphRARE framework is developed for heterophilic graphs based on the proposed node relative entropy and deep reinforcement learning (Sec. III). The combination of DRL and GNNs provides end-to-end training to optimize the original graph topology (Sec. IV-C).
- Extensive experiments have been conducted and the results show that the proposed method represents a general yet useful framework, which can enhance the GNN performance in node classification (Sec. V).

II. PRELIMINARY

In this section, we first introduce some basic concepts and then formulate the problem of heterophilic graph topology optimization. In addition, frequently used symbols are summarized in Table I.

¹Code is available at <https://anonymous.4open.science/r/GraphRARE>.

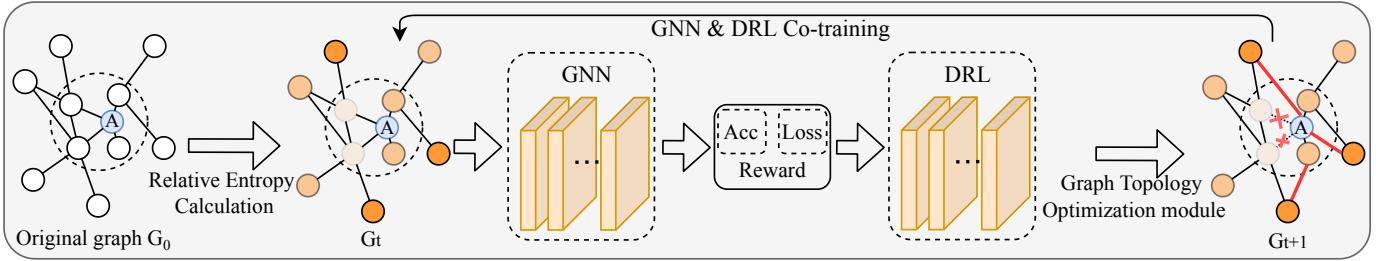


Fig. 2: An illustration of the proposed GraphRARE framework. The GNN and DRL module are trained jointly to optimize the graph topology based on the relative entropy (the depth of the node’s color indicates the value of entropy).

TABLE I: Symbols and Definitions.

Symbol	Definition
$G=(V, E, X, A)$	Graphs
N	Number of nodes in the graph
$V=\{v_1, v_2, \dots, v_N\}$	Set of nodes
$E \subseteq V \times V$	Set of edges
$X \in \mathbb{R}^{N \times d}$	Node feature matrix
$A \in \mathbb{R}^{N \times N}$	Adjacency matrix
\mathcal{H}	Homophily ratio of a graph
$H(v, u)$	Node relative entropy between node v and u
$H_f(v, u)$	Feature Entropy between node v and u
$H_s(v, u)$	Structural entropy between node v and u
λ	The hyper-parameter in node relative entropy
S_t	Multi-discrete state at step t
A_t	Multi-discrete action at step t
$R(S_t)$	Reward of state S_t
y_v	Class of node v
$\mathcal{N}_k(v)$	The k -hop neighbors of node v
$\mathcal{N}_1(v)$	The first-order neighbors of node v

A. Basic Concepts

Graph. A graph can be formulated as $G=(V, E, X, A)$, where $V=\{v_1, v_2, \dots, v_N\}$ denotes a node set, $E \subseteq V \times V$ is an edge set, $X \in \mathbb{R}^{N \times d}$ represents node features, and $A \in \mathbb{R}^{N \times N}$ is the adjacency matrix; N is the number of nodes, and d is the dimension of a node feature. Particularly, $a_{ij}=1$ indicates that an edge between nodes v_i and v_j exists; otherwise, $a_{ij}=0$. Further, x_i represents the feature vector of a node v_i . The degree of a node v is denoted by d_v . The k -hop neighbors of a node v are denoted by $\mathcal{N}_k(v)$. For instance, a directly connected (one-hop) neighbor set of a node v is $\mathcal{N}_1(v)$.

Heterophilic Graph and Homophily Ratio. Different nodes in a graph may correspond to different categories/classes, so we call it a heterophilic graph, where the linked nodes belong to different categories/classes or possess dissimilar features. In addition, the definition of heterophilic graphs can be measured by the homophily ratio, which is computed based on its edge homophily [60]:

$$\mathcal{H} = \frac{|\{(v, u) \in E : y_v = y_u\}|}{|E|} \quad (1)$$

Here, \mathcal{H} is the homophily ratio, y_v is the label of node v . The value range of \mathcal{H} is $[0, 1]$. A large \mathcal{H} indicates a homophilic graph, whereas a low \mathcal{H} indicates a heterophilic graph.

Topology Optimization for Heterophilic Graph Learning.

The majority of Graph Neural Networks (GNNs) follow the homophily assumption, which aggregates messages extracted from local neighbor nodes, but this approach may not effectively generalize to heterophilic graphs. Improving the performance of traditional GNNs (e.g., GCN, GAT, GraphSAGE) in such graphs is a challenging task. One general and effective method is to optimize the original graph topology (i.e., adding edges to link nodes or deleting existing links) using a well-designed metric to increase the homophily ratio. Numerous approaches, such as UGCN [16], NL-GNN [25], SimP-GCN [17], and GPNN [45], have been proposed for this purpose. By optimizing the heterophilic graph topology, the linked nodes in the graphs tend to have more similar features, making them more suitable for traditional GNNs. This enhancement allows the vanilla GNNs to perform better on heterophilic graphs and improve their overall performance and generalization capability.

B. Problem Formulation

The main goal of graph topology optimization for graph learning is to restructure the graph in a way that promotes information exchange and interactions between nodes of different types. To achieve this, a well-designed metric is used to measure the relevance between node pairs. Noisy edges are deleted, and informative node pairs are connected to enhance the flow of information in the graph. Formally, we can define it as follows.

Definition 1 (Topology Optimization for GNNs): Given a heterophilic graph $G = \langle V, E, X, A \rangle$, a graph neural network model (GNN) designed for a downstream task (i.e., node classification in this paper), the objective is to devise an effective transformation approach for creating a refined graph $G' = \langle V, E', X, A' \rangle$. This approach involves the addition or deletion of edges to optimize the performance of the GNN model on the given downstream task.

III. FRAMEWORK

As shown in Figure 2, we first introduce a highly effective relative entropy function designed to quantify the significance of one node with respect to another. In particular, to capture both nodes’ feature and graph structural information, our proposed node relative entropy is comprised of both node

feature entropy and node structural entropy. For example, given a node ‘A’ in the original graph, we leverage the entropy value to measure other nodes’ relative significance. Hence, we can add/delete edges linking ‘A’ and some nodes with high/low entropy values. However, each node exhibits its distinct characteristics or ”personality,” and optimizing the graph topology with a uniform addition or deletion of links may not fully consider these individual traits. Consequently, this approach can lead to sub-optimal results in the optimized graph.

To address this, we propose a novel approach that models the graph topology optimization as a finite horizon Markov Decision Process with multi-discrete states. This allows us to effectively account for the personality of different nodes within the graph. The entire model comprises two neural network modules: the *Graph Neural Network (GNN)*, which handles the downstream task of node classification, and the *Deep Reinforcement Learning (DRL)* module, responsible for determining the number of added or deleted edges for each node. Moreover, we adopt a co-training approach in a loop mode, where the GNN’s training loss or accuracy serves as the reward to train the DRL module. Concurrently, the DRL module generates a reconstructed graph, which then becomes the input to the GNN. The entire pipeline operates as follows: prior to model training, we compute the relative entropy between node pairs. During each training step, the GNN’s performance is evaluated on the optimized graph at step t (G_t). Based on this performance, including accuracy and loss on the training set, the policy network of the DRL is updated. Subsequently, the graph topology optimization module generates the optimized graph G_{t+1} and feeds it into the GNN for the next iteration of co-training.

In summary, by incorporating reinforcement learning into the graph topology optimization process, we aim to better capture and utilize the diverse traits and relationships between nodes, leading to improved performance in various tasks. This framework allows us to tailor the graph topology to suit the specific characteristics of the nodes, enhancing the overall effectiveness of our proposed model.

IV. METHODOLOGY

This section presents our GraphRARE framework for node classification in heterophilic graphs, addressing the main challenges related to metric and personality of nodes, as described in Section I. The framework, depicted in Figure 2, follows a multi-step approach. Initially, GraphRARE computes the relative entropy between nodes, and subsequently trains the GNN and DRL modules jointly in an end-to-end manner to accomplish the node classification task. In particular, we first define the fundamental concepts of node relative entropy, encompassing both node feature entropy and node structural entropy. Subsequently, we elucidate the construction of a node sequence based on the relative entropy theory. In addition, we develop a graph reconstruction mechanism to optimize the original graph topology based on deep reinforcement learning.

Finally, we take advantage of the GNNs in node feature learning to complete the node classification task.

A. Node Relative Entropy Calculation

Entropy is an effective method to measure uncertain information in a graph. Relative entropy (Kullback–Leibler divergence) [19] is one measure to quantify differences in probability distributions, which is defined as follows:

$$D_{KL}(P\|Q) = \sum_{i=1}^n P(i) \log \frac{P(i)}{Q(i)} \quad (2)$$

where P_i and Q_i are two probability distributions of the event i ; n is the number of events; P and Q have the same number of components. Notably, the smaller the entropy is, the smaller the difference between two distributions will be.

However, it is non-trivial to measure the graphs using the original definition of relative entropy where the rich node features and link structures are difficult to quantify. Inspired by [26], [50], our work optimizes the original graph entropy and introduces a node relative entropy that consists of feature entropy and structural entropy to address this issue. To measure the feature similarity between nodes, the graph feature entropy in [26] is optimized in this study by calculating the entropy between any pair of nodes, and the computational complexity is also reduced. Since the feature entropy does not take structural information into account, we introduce structural entropy to measure the structural similarity between nodes. However, the structural entropy in [50] has difficulty in accurately measuring the semantically related node pairs, and we optimize it by introducing Jensen–Shannon (JS) divergence [24].

1) *Node Feature Entropy*: Feature entropy measures the difference between two nodes based on their feature vectors. Inspired by the MinGE model [26], which defines the graph feature entropy, this paper designs a feature entropy for each pair of nodes in graph data. Compared to the graph feature entropy in [26], the node feature entropy requires lower computation complexity and is able to compute entropy at the node level. Based on the assumption of node embedding, considering that nodes in the same class are more similar, the node embedding dot product of node pairs is used as a node feature entropy, which can be expressed as follows:

$$z_v = \phi(x_v) \in \mathbb{R}^h, \forall v \in V \quad (3)$$

$$\begin{aligned} H_f(v, u) &= -P(z_v, z_u) \log P(z_v, z_u) \\ &= -\frac{e^{\langle z_v, z_u \rangle}}{\sum_{i,j} e^{\langle z_i, z_j \rangle}} \log \frac{e^{\langle z_v, z_u \rangle}}{\sum_{i,j} e^{\langle z_i, z_j \rangle}} \end{aligned} \quad (4)$$

where $\phi(\cdot)$ is an embedding function (e.g., multilayer perceptron), which converts original features to vectors; x_v is a node feature of a node v ; h is the dimension of the node feature embedding z_v ; $\langle \cdot, \cdot \rangle$ is the dot product operation; $\sum_{i,j} e^{\langle z_i, z_j \rangle}$ represents the dot product sum of a pair of nodes in a graph G ; $H_f(v, u)$ is the node feature entropy between nodes v and

u . As shown in Eq. (4), the larger the node feature entropy is, the more similar the node features are.

2) *Node Structural Entropy*: Structural entropy measures the difference between two nodes based on their local structure [50]. Node degree has been one of the most important structural features in the graph structure entropy [29]. Therefore, it is reasonable to calculate the node structural entropy by comparing the ordered degree sequences of nodes and their first-order neighbors. However, the local relative entropy [50] is directly measured by KL divergence whose value ranges in $[0, +\infty]$, and thus the entropy has no practical meaning when the value is too large. Therefore, we optimize it by introducing Jensen–Shannon (JS) divergence [24] with values in the range $[0, 1]$ to measure semantically related remote nodes more accurately and more reasonable.

According to the definition of JS divergence, we first compute the KL divergence $D_{KL}(v\|\frac{v+u}{2})$ from a node u to the node v in a graph G , which is calculated as follows:

$$d(v) = [d_{v1}, d_{v2}, \dots, d_{vM}] \in \mathbb{R}^M \quad (5)$$

$$p(v) = [p_{v1}, p_{v2}, \dots, p_{vM}] \in \mathbb{R}^M$$

$$= \left[\frac{d_{v1}}{\sum_{i=1}^M d_{vi}}, \frac{d_{v2}}{\sum_{i=1}^M d_{vi}}, \dots, \frac{d_{vM}}{\sum_{i=1}^M d_{vi}} \right] \quad (6)$$

$$D_{KL}(v\|\frac{v+u}{2}) = \sum_{i=1}^M p_{vi} \log \frac{p_{vi}}{\frac{p_{vi}+p_{ui}}{2}} \quad (7)$$

where $d(v)$ is a descending-order sequence containing degree values d_{vi} of node v and its one-hop neighbors; $M = \max\{degree(v) | v \in G\}$ is the max degree of a node in a graph G ; $d_{vi}=0$ if $i > (degree(v) + 1)$; $\sum_{i=1}^M d_{vi}$ is the sum of the sorted sequences $d(v)$; $p(v)$ is the normalized result of $d(v)$.

Then, the structural entropy between nodes v and u can be calculated based on $D_{KL}(v\|\frac{v+u}{2})$ and $D_{KL}(u\|\frac{v+u}{2})$ as follows:

$$H_s(v, u) = 1 - \frac{1}{2} \left(D_{KL}(v\|\frac{v+u}{2}) + D_{KL}(u\|\frac{v+u}{2}) \right) \quad (8)$$

where $D_{KL}(v\|\frac{v+u}{2})$ represents the relative entropy from a node u to the node v , and $D_{KL}(u\|\frac{v+u}{2})$ is the opposite; $H_s(v, u)$ is symmetric and denotes the structural entropy. As defined in Eq. (8), the larger value of the node structural entropy suggests the higher similarity between the nodes' structures.

3) *Node Relative Entropy*.: Node relative entropy consists of feature entropy and structural entropy. Formally,

$$H(v, u) = H_f(v, u) + \lambda H_s(v, u) \quad (9)$$

where $H(v, u)$ represents node relative entropy between nodes v and u ; $H_f(v, u)$ is symmetric and denotes the feature entropy; $H_s(v, u)$ is symmetric and denotes the structural entropy; λ is a hyper-parameter that controls the ratio.

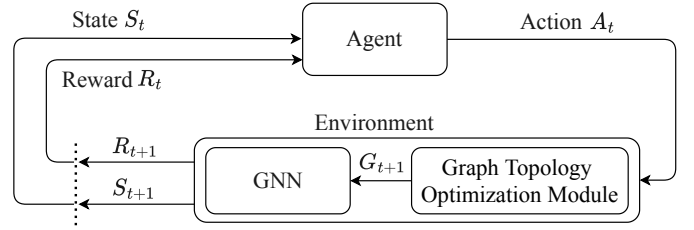


Fig. 3: The agent environment interaction in RL.

4) *Node Entropy Sequence Construction*: The original graph topology might not be appropriate for heterophilic GNNs with a message passing framework, where nodes of the same class have a high degree of structural similarity but may be far away from each other [57]. Considering the node relative entropy, which reflects the relevant relationship between nodes, a sorted sequence is constructed for each node according to the entropy value. Namely, *top-k* new neighbors are selected for each node, and the original graph topology is optimized by connecting the selected new neighbors with the ego node. In this way, potential neighbors that are most similar to the ego node can be discovered and selected. Theoretically, since the node entropy sequence can be constructed flexibly to cover the whole graph, this strategy should be able to capture long-range dependencies even if the two nodes are distant. In this study, we select the *top-k* nodes from the ego node's remote neighbors and delete the *top-d* nodes from the one-hop neighbors based on the relative entropy. The values of k and d can be different on each individual node and they are updated by a deep reinforcement learning-based mechanism during model training.

Time Complexity Analysis. In the worst case, the time complexity of node relative entropy calculation is $O(N^2)$ for any size of the graph due to the matrix multiplication. By using PyTorch [30] functions and matrix operations, we can significantly reduce the complexity of the calculation. For instance, the entropy is symmetric, and calculating the entropy between nodes v and u only needs to be performed once. Additionally, we only need to calculate the entropy once before training the GNN and DRL instead of in each epoch. In practice, graphs are usually sparse, and hence the empirical complexity is much smaller than $O(N^2)$ by using sparse matrix computation techniques.

B. Deep Reinforcement Learning Module

To mitigate the “personality of nodes” challenge in the *top-k* and *top-d* selection, this study proposes a deep reinforcement learning-based mechanism to update the k and d value of a node v in a graph G regardless of the size and topology structure of a graph. If k (or d) is set as a hyper-parameter, meaning that all nodes will select the same number of *top-k* new neighbors (or *top-d* one-hop neighbors deleted), this can either introduce noisy nodes or discard informative nodes. Thus, instead of defining it as a single hyper-parameter, a DRL-based algorithm is used to find the optimal *top-k* and

$top-d$ values for every graph node. The updating process of $top-k$ and $top-d$ is modeled as a finite horizon Markov Decision Process (MDP), which is characterized by a tuple $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $S_t \in \mathcal{S}$, $A_t \in \mathcal{A}$ are state and action observed at step t , P is a state transition function, R is a reward function, and γ is a discount factor. The MDP tuple is defined as follows:

- **State:** The state S is set as $S=[k_1, k_2, \dots, k_N, d_1, d_2, \dots, d_N]$, where k_i represents the number of newly connected neighbors of a node v_i , and d_i is the number of deleted neighbors. The state at step t is S_t ; when $t=0$, $S_0=[0, 0, \dots, 0]$.
- **Action:** The DRL agent updates k_t by taking an action A_t based on S_t . Since state S_t is a multi-discrete state, action $A=[a_1^k, a_2^k, \dots, a_N^k, a_1^d, a_2^d, \dots, a_N^d]$ is defined as a multi-discrete action, and a_i^k (a_i^d) is defined as the number of newly connected neighbors k_i (deleted neighbors d_i) added or subtracted by a fixed value $\Delta k=1$ or kept unchanged.
- **State Transition:** The state transition function is defined by:

$$S_{t+1} = P(A_t|S_t) = S_t + A_t \quad (10)$$

After updating the state S , the original graph topology is reconstructed by adding edges between node v and its entropy sequence's $top-k_v$ nodes.

- **Reward:** The goal of DRL is to find an optimal graph topology for the original graph and GNN model. A discrete reward function is defined based on the GNN performance as follows:

$$R(S_t) = (acc_t - acc_{t-1}) + \lambda_r(loss_{t-1} - loss_t) \quad (11)$$

where acc_t and $loss_t$ are node classification accuracy and loss of the GNN on the training set at step t , respectively; λ_r is a hyper-parameter that controls the ratio. The reward function indicates that if the performance of GNN on graph G_t is better than that on graph G_{t-1} , then the topology of G_t is better than G_{t-1} . It is worth noting that the current reward function can be replaced with alternative ones as long as they are able to guide the DRL towards optimizing the original topology.

- **Graph Topology Optimization Module:** At every step t , the proposed method optimizes the graph topology from G_t to G_{t+1} . The DRL policy network first generates an action A_t based on the state S_t at step t , and the environment updates the state to S_{t+1} according to Eq. (10). Then, the topology optimization module connects node v and the $top-k_v$ nodes in the node entropy sequence seq_v for each node $v \in V$ to construct graph G_{t+1} (see Figure 4). Finally, the GNN model is trained on the new graph G_{t+1} and provides its accuracy and loss value on the training dataset to the DRL module as the reward according to Eq. (11).

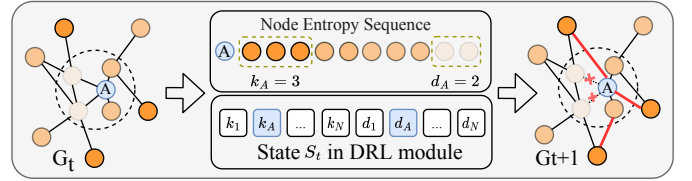


Fig. 4: An illustration of the graph topology optimization module. For any node v in the graph G_t at step t , the graph optimization module generates graph G_{t+1} via adding k_v edges and removing d_v edges for node A based on the state S_t . For node A in this example, $k_A = 3$ edges are added and $d_A = 2$ edges are removed.

Algorithm 1: The overall process of GraphRARE

Require: $G = (V, E, X, A)$
Ensure: Node label y_v

```

// Calculate node relative entropy
1: for  $v, u \in V$  do
2:   Calculate feature entropy  $H_f(v, u)$  via Eq. (4)
3:   Calculate structural entropy  $H_s(v, u)$  via Eq. (8)
4:   Calculate relative entropy  $H(v, u)$  via Eq. (9)
5: end for
6: Construct node entropy sequence based on  $H(v, u)$ 
7: Initialize  $max\_acc = 0$ 
// Train GNN and DRL jointly in an end-to-end manner
8: for iteration=1,2,... do
// Evaluate on training set, no backward
9:    $acc_t, loss_t = GNN(G_t, X)$ 
10:  if  $acc_t > max\_acc$  then
// Additional training on the current graph  $G_t$  as it
// leads to a higher accuracy
11:    $max\_acc \leftarrow acc_t$ 
12:   Train  $GNN(G_t, X)$  for a few more epochs
13:  end if
// DRL process
14:   $reward(S_t) \leftarrow$  Eq. (11)
15:   $A_t \leftarrow$  DRL policy network
16:   $S_{t+1} = P(A_t|S_t)$  via Eq. (10)
17:  Reconstruct graph  $G_{t+1}$  based on  $G_t, S_{t+1}$ , and node
// entropy sequence
18: end for

```

During the training process at step t , the DRL agent and GNN model are trained jointly in an end-to-end manner to optimize the original graph topology dynamically. Figure 3 illustrates the agent-environment interaction in reinforcement learning. Firstly, the GNN model is evaluated on the training set based on the graph topology G_t , and calculate the accuracy acc_t and $loss_t$. Notably, during this evaluation phase, no updates are made to the parameters of GNN. If acc_t surpasses the previous best, we train the GNN model for a few epochs on the graph topology G_t to complete the node classification task. To prevent overfitting on G_t , an early stopping strategy is implemented. Secondly, the DRL agent calculates the reward

according to Eq. 11 and updates the policy network. Thirdly, the DRL agent takes action A_t and updates the state to S_{t+1} according to Eq. 10. Lastly, the graph topology optimization module reconstructs the graph topology by connecting the selected remote nodes based on the state S_{t+1} . The complete process is shown in Lines 7–18 in Algorithm 1.

Lastly, the graph topology optimization module reconstructs the graph topology by connecting the selected remote nodes based on the state S_{t+1} .

As the DRL algorithm, we use the Proximal Policy Optimization algorithm (PPO) [35], which is a classical policy gradient optimization algorithm that can deal with large-scale multi-discrete states and action spaces. It should be noted that in addition to the PPO algorithm, other reinforcement learning algorithms can also be conveniently applied to the proposed framework.

C. Node Feature Learning

The GraphRARE framework can be easily adapted to any existing GNN model to improve the model’s performance. In the proposed framework, the original graph topology is reconstructed by training the DRL and GNN jointly based on the node relative entropy. In order to retain the GNN advantages, we adopt the same mechanism of node feature learning in the existing GNNs for multi-layer local aggregation. A GNN model consists of L local aggregation layers, allowing each node to access the L -hop neighbors’ information. For a node $v \in V$, the calculation formula of the l th layer in a L -layer GNN ($l=1, 2, \dots, L$) is given by:

$$m_v^{(l)} = \text{AGGREGATE}^{(l)} \left(\left\{ h_u^{(l-1)} : u \in \mathcal{N}_1(v) \right\} \right) \quad (12)$$

$$h_v^{(l)} = \text{UPDATE}^{(l)} \left(h_v^{(l-1)}, m_v^{(l)} \right) \quad (13)$$

where $h_v^{(l)}$ is the feature of a node v in the l th layer, and $h_v^{(0)}=x_v$; $\text{AGGREGATE}^{(l)}(\cdot)$ and $\text{UPDATE}^{(l)}(\cdot)$ represent the feature aggregation function (e.g., mean, LSTM, and max pooling) and feature update function (e.g., linear-layer combination and MLP) [57], respectively; $\mathcal{N}_1(v)$ represents the one-hop neighboring node set of a node v .

D. Model Training

In Algorithm 1, we present a step-by-step depiction of the training process. Beginning with a heterophilic graph $G = (V, E, X, A)$, we compute the relative entropy between pairs of nodes, considering both their node features and local structures (lines 1-5). Subsequently, we construct entropy sequences for each node in descending order, based on their relative entropy values (line 6). Lastly, we perform joint training of GNN and DRL module in an end-to-end manner. During iteration t , the GNN evaluates the graph G_t and produces accuracy and loss metrics, which are then utilized to calculate the reward. The DRL policy network is updated based on the reward, and we optimize the graph to G_{t+1} for the subsequent iteration.

TABLE II: Statistics and properties of the seven datasets.

Datasets	#Nodes	#Edges	#Features	#Classes	\mathcal{H}
Chameleon	2,277	36,101	2,325	5	0.23
Squirrel	5,201	217,073	2,089	5	0.22
Cornell	183	295	1,703	5	0.30
Texas	183	309	1,703	5	0.11
Wisconsin	251	499	1,703	5	0.21
Cora	2,708	5,429	1,433	7	0.81
Pubmed	19,717	44,338	500	3	0.80

V. EXPERIMENTS

In this section, we evaluate the performance of the GraphRARE framework on the node classification task on seven real-world datasets. We first introduce the datasets, the baseline models, and hyper-parameter settings. Then, we compare GraphRARE with baseline models and analyze the results. Finally, the contribution of different components to the performance of GraphRARE is evaluated via hyper-parameter analysis, ablation study on the DRL module, convergence analysis for model training, and graph topology optimization analysis.

A. Datasets

The proposed GraphRARE is evaluated on seven widely used real-world graph datasets, including five heterophilic graph datasets and two homophilic graph datasets. The heterophilic graph datasets are Chameleon, Squirrel [34], Cornell, Texas, and Wisconsin [9]. The homophilic graph datasets are Cora [36] and Pubmed [28]. The statistics for all the datasets are summarized in Table II.

B. Compared Methods and Setup

In the experiments, four enhanced GNNs based on the proposed framework are constructed: GCN-RARE, GraphSAGE-RARE, GAT-RARE, and H2GCN-RARE. To be specific, GCN-RARE denotes reinforcement learning enhanced GCN model with relative entropy, and so do the other three models. The models are end-to-end trained using the node classification loss. The proposed four enhanced GNNs are compared with thirteen baseline models, including an attributed-only based MLP, three traditional GNN models (GCN, GraphSAGE, and GAT), and nine GNN-based state-of-the-art (SOTA) methods for heterophilic graphs (MixHop, H2GCN, Geom-GCN, UGCN, SimP-GCN, OTGNet, GBK-GNN, Polar-GNN, and HOG-GCN) for node classification tasks. Specifically, the OTGNet model originally handles the heterophily for temporal graph. Since the underlying mechanisms of OTGNet for managing diverse classes and propagating information are applicable across temporal and static graph structures, we adopt the static graphs as input for a fair comparison.

C. Hyper-parameter setting

Based on the hyper-parameter settings of the original models [11], [18], [40], [60], the final hyper-parameter setting is configured as follows: the dropout rate is set to $p=0.5$, the initial learning rate is 0.05, and the weight decay is set to

TABLE III: Mean accuracy and standard deviation over ten different data splits on the seven real-world graph datasets. The best results are highlighted in **bold**. The improvements of the GraphRARE models compared to their counterparts are calculated, with an improvement denoted by an up arrow \uparrow , and a negative improvement by a down arrow \downarrow .

Method	Chameleon	Squirrel	Cornell	Texas	Wisconsin	Cora	Pubmed	Average
MLP	46.51±2.53	29.29±1.40	80.81±6.91	81.08±5.41	84.12±2.69	74.61±1.97	86.63±0.38	69.01
GCN [18]	59.08±2.47	46.64±1.42	55.73±6.33	52.84±5.43	56.04±5.76	85.16±1.01	87.18±0.42	63.24
GraphSAGE [11]	58.83±2.15	41.44±1.85	72.70±9.16	75.68±6.16	76.08±4.28	84.53±1.38	85.09±0.52	70.62
GAT [40]	54.34±2.46	40.79±3.78	54.22±5.38	56.49±6.89	54.45±5.62	86.02±1.37	86.55±0.47	61.84
MixHop [1]	60.50±2.53	43.80±1.48	73.51±6.34	77.84±7.73	75.88±4.90	83.10±2.03	80.75±2.29	70.77
H2GCN [60]	56.85±1.68	32.20±2.19	78.16±4.05	79.70±5.16	82.08±3.22	86.26±1.08	88.76±0.41	72.00
Geom-GCN [31]	60.90	38.14	60.81	67.57	64.12	85.27	90.05	66.69
UGCN [16]	54.07	34.39	69.77	71.72	69.89	84.00	85.22	67.01
SimP-GCN [17]	62.61	42.57	84.05	81.62	85.49	82.80	81.10	74.33
OTGNet [7]	46.34±5.77	35.39±2.96	58.19±9.61	65.81±10.38	61.23±8.69	73.31±4.48	76.64±2.18	59.56
GBK-GNN [4]	48.46±0.81	36.69±0.33	69.59±2.24	75.59±2.94	78.98±4.19	82.65±0.64	83.48±0.19	67.92
Polar-GNN [6]	64.0±0.6	49.3±0.8	-	-	-	83.1±0.9	80.2±0.4	-
HOG-GCN [42]	54.01±1.28	35.46±1.96	84.32±4.32	85.17±4.40	86.67±3.36	87.04±1.10	88.79±0.40	74.49
GCN-RARE(ours)	68.05±1.87 $\uparrow 8.97$	55.90 ±1.39 $\uparrow 9.26$	64.59±4.95 $\uparrow 8.86$	58.38±6.64 $\uparrow 5.54$	61.76±5.49 $\uparrow 5.72$	87.24 ±1.26 $\uparrow 2.08$	88.41±0.50 $\uparrow 1.23$	69.19 $\uparrow 5.95$
GraphSAGE-RARE(ours)	69.28 ±1.90 $\uparrow 10.45$	52.84±1.33 $\uparrow 11.40$	82.97±5.10 $\uparrow 10.27$	82.16±6.07 $\uparrow 6.48$	85.69±5.29 $\uparrow 9.61$	87.08±1.17 $\uparrow 2.55$	89.03±0.56 $\uparrow 3.94$	78.43 $\uparrow 7.81$
GAT-RARE(ours)	64.56±2.48 $\uparrow 10.22$	49.99±2.79 $\uparrow 9.20$	61.60±3.01 $\uparrow 7.38$	58.11±5.44 $\uparrow 1.62$	61.08±3.80 $\uparrow 6.63$	86.60±1.19 $\uparrow 0.58$	87.41±0.65 $\uparrow 0.86$	67.05 $\uparrow 5.21$
H2GCN-RARE(ours)	58.09±1.91 $\uparrow 1.24$	34.93±1.49 $\uparrow 2.73$	87.84 ±4.05 $\uparrow 9.68$	86.76 ±5.80 $\uparrow 7.06$	90.00 ±2.97 $\uparrow 7.92$	86.82±1.51 $\uparrow 0.56$	90.07 ±0.26 $\uparrow 1.31$	76.36 $\uparrow 4.36$

$\{5E-5, 5E-6\}$. For GCN, GAT, GraphSAGE, and H2GCN, the number of hidden units is selected from the set $\{48, 64, 128\}$, the number of GNN layers is set to two, and the Adam optimizer is adopted. For the sake of fair comparison, in GCN-RARE, GraphSAGE-RARE, GAT-RARE, and H2GCN-RARE, the hyper-parameters are configured to be the same as those in their counterparts (GCN, GraphSAGE, GAT, and H2GCN). We set the λ in Eq.(9) to $\lambda=1.0$ and discuss it in Section V-E. For the PPO algorithm, we choose a MLP network for the policy model.

For all the benchmarks, we adopt the same class labels, feature vectors, and perform 10 random splits on the datasets (60%/20%/20% of nodes per class for training/validation/testing) provided by Pei et al. [31], which is available on their GitHub². We launch the testing procedure when the validation accuracy of the trained model achieves a maximum value for each run, and calculate the average test accuracy over these ten runs.

The GCN, GraphSAGE, GAT, H2GCN and our methods are constructed using Pytorch [30], Deep Graph Library [41], Pytorch Geometric [8], OpenAI Gym [2], and Stable-Baselines3 [33] with one NVIDIA A100-40GB GPU.

D. Node Classification

The experimental results for the node classification task on the seven datasets are illustrated in Table III. The mean accuracy and standard deviation for ten different data splits are presented, and the best results are highlighted in bold.

With the node relative entropy guided higher-order neighbor ranking and the deep reinforcement learning-based graph topology optimization, GraphRARE obtains a better graph topology under the heterophily setting. We apply different enhanced GraphRARE models and make comparisons with corresponding baselines and the GNN-based SOTA models. As displayed in Table III, three patterns can be observed: (1) The enhanced models, GCN-RARE, GraphSAGE-RARE, GAT-RARE, and H2GCN-RARE, demonstrate improvements over their counterparts (GCN, GraphSAGE, GAT, and H2GCN) on all the heterophilic datasets, which indicates that the advantages of the proposed GraphRARE framework are general for common GNNs. Specifically, the proposed framework improves the performance of GCN by 5.95% on average, GraphSAGE by 7.81% on average, GAT by 5.14% on average, and H2GCN by 4.23% on average. (2) On the dataset with strong homophily (Cora and Pubmed), GraphRARE performs better or is comparable to the baselines. To be specific, GraphRARE performs best on Cora, and achieves the second best on Pubmed. The experimental results show that the GraphRARE framework also performs well on homophilic graphs. (3) On the dataset with middle and low homophily, e.g., Chameleon, Squirrel, the four GraphRARE-based models achieve the best results on the five heterophilic datasets, and the overall performance is 6.69% higher than the best SOTA method (SimP-GCN) on average. This is mainly attributed to the fact that GraphRARE adaptively optimizes the graph topology based on the DRL module and the node relative entropy, it effectively ignores the irrelevant nodes and selects the most important remote nodes by considering the node

²<https://github.com/graphdml-uiuc-jlu/geom-gcn/tree/master/splits>

TABLE IV: Hyper-parameter analysis results. The hyper-parameter λ controls the weight of structural entropy in the node relative entropy as defined in Eq.(9). The best results are highlighted in **bold**.

Method	λ	Chameleon	Squirrel	Cornell	Texas	Wisconsin	Cora	Pubmed	Average
GCN-RARE	0.1	67.36±2.25	54.89±1.54	63.92±7.24	57.83±7.81	59.31±5.07	87.34 ±1.07	87.49±0.72	68.31
	0.5	67.56±2.01	54.77±1.46	63.77±5.69	57.78±6.36	58.93±4.81	86.21±1.33	87.62±0.74	68.10
	1.0	68.05 ±1.87	55.90 ±1.39	64.59 ±4.95	58.38 ±6.64	61.76 ±5.49	87.24±1.26	88.41 ±0.50	69.19
	10.0	67.73±1.98	55.45±1.25	63.54±6.23	57.79±7.46	58.82±4.60	86.27±1.48	87.77±0.97	68.20
GraphSAGE-RARE	0.1	68.70±2.66	52.93 ±1.74	80.03±5.70	80.54±5.90	84.16±4.68	86.74±1.37	88.40±0.74	77.36
	0.5	68.95±2.37	52.58±1.61	81.00±8.48	81.46±5.57	84.94±3.73	87.06±0.82	89.12 ±0.38	77.87
	1.0	69.28 ±1.90	52.84±1.33	82.97 ±5.10	82.16 ±6.07	85.69 ±5.29	87.08 ±1.17	89.03±0.56	78.43
	10.0	68.78±2.01	52.67±1.41	81.54±7.23	81.49±5.14	84.96±4.62	86.76±1.33	88.77±0.93	77.85
GAT-RARE	0.1	63.77±2.32	49.54±2.55	60.00±7.03	56.89±6.14	60.12±6.15	86.22±1.26	86.99±0.83	66.22
	0.5	63.18±2.39	48.88±2.30	60.84±4.86	57.08±5.47	60.71±4.59	86.54±1.29	87.38±0.43	66.37
	1.0	64.56 ±2.48	49.99 ±2.79	61.60 ±3.01	58.11 ±5.44	61.08 ±3.80	86.60 ±1.19	87.41 ±0.65	67.05
	10.0	64.30±3.04	49.96±2.25	59.73±4.75	56.24±6.40	60.90±6.74	86.12±1.43	87.16±0.57	66.34
H2GCN-RARE	0.1	58.03±1.64	33.77±2.19	86.86±5.37	86.31±3.60	88.43±4.15	86.40±1.24	89.48±1.24	75.61
	0.5	57.70±1.66	33.65±2.03	87.42±5.92	86.12±3.51	87.65±3.17	86.08±1.76	89.67±0.59	75.47
	1.0	58.09 ±1.91	34.93 ±1.49	87.84 ±4.05	86.76 ±5.80	90.00 ±2.97	86.82 ±1.51	90.07 ±0.26	76.36
	10.0	57.35±1.74	32.90±1.83	86.24±5.51	85.86±3.67	87.65±3.17	86.82 ±1.42	89.35±0.92	75.17

TABLE V: Ablation study on relative entropy and DRL module. “GCN-RE[.]” represents GraphRARE without DRL, instead randomly assigning k and d to each node. “GCN-RA” denotes GraphRARE without the relative entropy component. “GCN-RARE-add” and “GCN-RARE-remove” describe GraphRARE when it only adds edges and only removes edges, respectively.

Method	Chameleon	Squirrel	Cornell	Texas	Wisconsin	Cora	Pubmed	Average
GCN [18]	59.08±2.47	46.64±1.42	55.73±6.33	52.84±5.43	56.04±5.76	85.16±1.01	87.18±0.42	63.24
GCN-RE[0..5](ours)	63.48±2.09	48.03±1.05	59.72±3.72	55.43±9.48	56.17±4.38	84.32±1.60	85.13±0.46	64.61
GCN-RE[0..10](ours)	60.89±2.31	46.04±0.73	61.35±3.64	56.21±5.75	59.49±5.33	83.44±1.59	84.52±0.39	64.56
GCN-RE[0..15](ours)	58.69±2.16	45.19±1.37	61.08±5.01	53.74±6.13	61.25±5.04	83.11±1.63	84.15±0.40	63.89
GCN-RE[0..20](ours)	58.55±1.40	45.09±1.42	61.05±7.57	58.61±7.92	59.62±7.37	83.27±1.36	83.97±0.49	64.31
GCN-RA(ours)	61.48±1.39	47.50±1.71	59.57±6.56	54.57±7.94	59.65±7.83	84.98±1.37	87.42±0.59	63.78
GCN-RARE-add(ours)	66.43±1.36	55.46±1.14	58.11±6.19	58.12±7.74	59.22±4.45	86.58±1.34	88.02±0.52	67.25
GCN-RARE-remove(ours)	67.52±1.68	55.43±1.42	60.95±4.55	55.14±7.66	61.37±7.07	86.88±1.23	87.95±0.62	67.89
GCN-RARE-reward(ours)	66.54±2.12	53.05±1.33	60.64±4.99	54.02±7.26	58.74±6.25	86.72±1.37	87.74±0.57	66.78
GCN-RARE(ours)	68.05 ±1.87	55.90 ±1.39	64.59 ±4.95	58.38 ±6.64	61.76 ±5.49	87.24 ±1.26	88.41 ±0.50	69.19

feature and structural information.

E. Hyper-parameter Analysis

With the calculation of the node relative entropy, the GraphRARE framework optimizes the graph topology adaptively. As defined in Eq.(9), the hyper-parameter λ controls the weight of structural entropy in the node relative entropy. In order to determine the weight of the structure entropy, we conduct extensive experiments to analyze λ on the benchmarks. The experimental results are depicted in Table IV. As shown in the table, it is not hard to see that the structural entropy is often equally important as the feature entropy for the GraphRARE framework. Therefore, given an arbitrary graph, we can simply set λ to the default value of 1.0. As the relative entropy with $\lambda=0.1$ or $\lambda=10.0$ can be considered as feature entropy alone or structural entropy alone respectively, we also observe that the performance is better when considering both feature entropy and structure entropy than when considering either of them separately. The experimental results verify our assumption that both the node features and the local topological structure

contribute to measuring the mutual information between node pairs.

F. Ablation Study

The proposed GraphRARE framework jointly trains the DRL-based module and GNN to optimize the graph topology. An ablation study on the GraphRARE framework is conducted to evaluate the DRL module and the relative entropy’s contribution by (1) setting the same number of new neighbors k and deleted neighbors d for each node to a fixed value. (2) randomly allocating different k and d value for ego nodes. (3) randomly shuffling the node entropy sequence for each node, which is GraphRARE without relative entropy. (4) only adding edges or removing edges in the original graphs. (5) replace the reward function in DRL module.

1) *Under the same k and d value setting:* To fully investigate the performance under different settings, we set the fixed new neighbor number k from one to ten for the four enhanced GNNs on three datasets (Chameleon, Squirrel, and Cora). The experimental results are depicted in Figure 5, and two patterns can be observed: (1) For the four backbone models

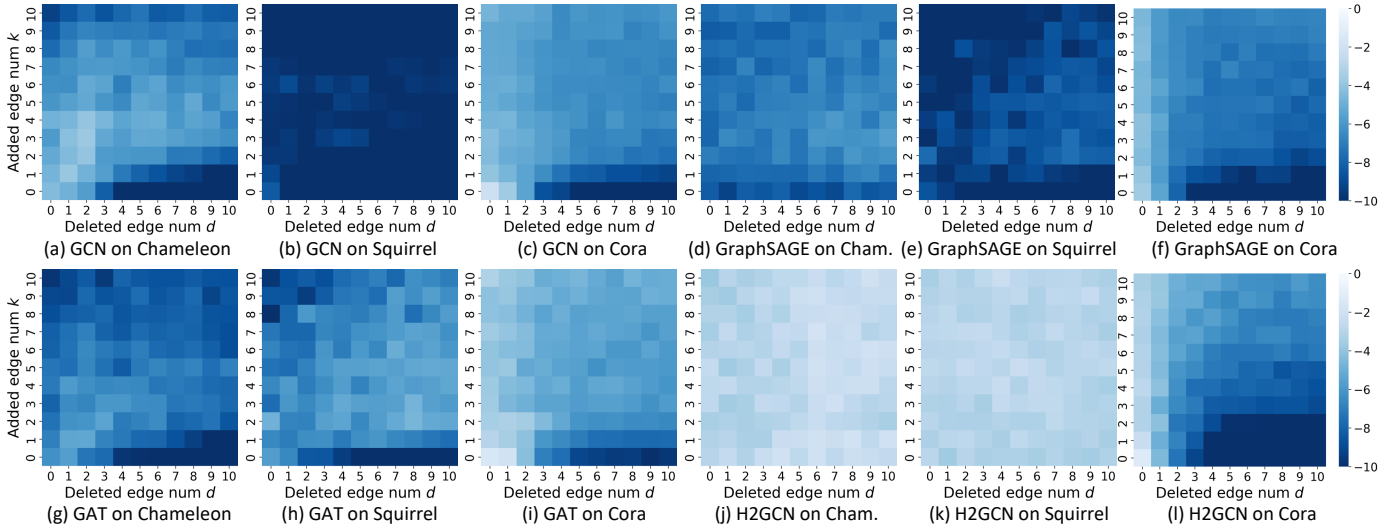


Fig. 5: Ablation study on the DRL module’s contribution. Each heatmap compares GraphRARE’s performance with and without the DRL module. The horizontal axis represents the number of deleted edges (d), while the vertical axis represents the number of added edges (k). Deeper colors indicate more significant performance degradation compared to GraphRARE.

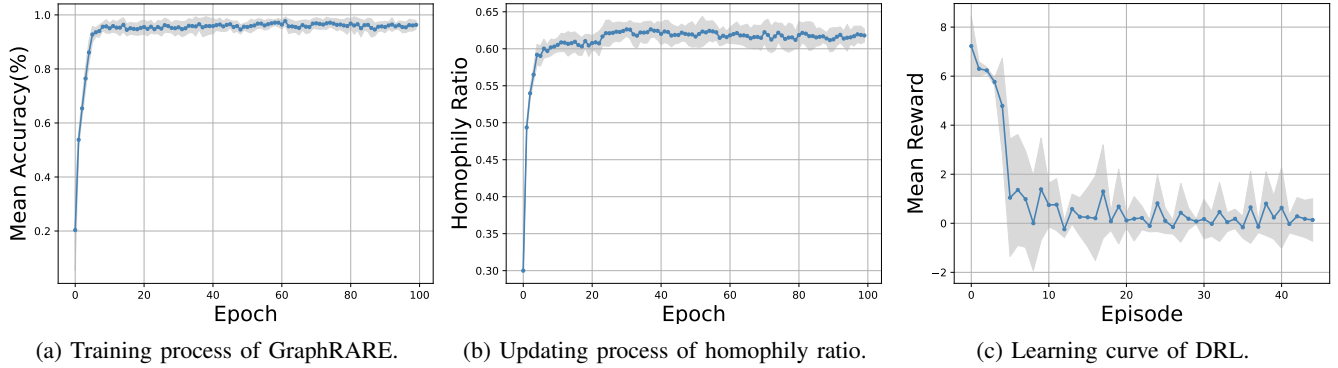


Fig. 6: The training process of GraphRARE. The subfigures illustrate the training performance of GCN-RARE on the Cornell dataset. Subfigure 6a shows node classification accuracy, Subfigure 6b shows graph homophily ratio, and Subfigure 6c shows the mean reward of the DRL module in each episode.

on the three datasets, the proposed framework with the DRL module outperforms all of the models with a fixed new/deleted neighbor number. As discussed above, this could be because the model with a fixed k and d falls into a local sub-optimal solution due to the mixture of useful information and noises of remote nodes. (2) Removing edges has a greater impact on performance than adding edges in most experiments. When edges are removed, it can lead to the disconnection of nodes or subgraphs, reducing the information flow and the ability of GNNs to propagate messages effectively across the graph.

2) *Under the random k and d value setting:* We conduct experiments for GCN on seven datasets that randomly allocate different k -value for ego nodes. The experimental results are illustrated in Table V (see GCN-RE), which shows that randomly allocating different k -values for nodes can improve the performance of the backbone GNN but is inferior to that of using DRL module in some datasets. The results demonstrate the effectiveness of both relative entropy and DRL module.

3) *Without the relative entropy:* The node entropy sequence is built based on the value of node relative entropy. To evaluate the contribution of the relative entropy, we randomly shuffle the node entropy sequence for each node, which equals “GraphRARE without relative entropy”. The results in Table V (see GCN-RA) emphasize the criticality of defining an appropriate metric for measuring the node importance when optimizing the graph topology.

4) *Without adding or removing edge operation:* In our framework, we jointly train the GNN and DRL module to optimize the graph topology by simultaneously adding and removing edges. The results in Table V (see GCN-RARE-add and GCN-RARE-remove) highlight the significance of adding informative new neighbors and removing noisy neighbors are both important.

5) *Alternative reward function:* The reward function is crucial as it guides the policy towards optimized decisions. In order to explore the impact of different reward design choices

TABLE VI: Comparisons in terms of real running time. Those appended by an asterisk [*] denotes SOTA models. The relative entropy is only computed once before model training.

Method	Chameleon	Squirrel	Cornell	Texas	Wisconsin
GCN	11.36	13.3	9.00	9.32	9.32
GAT	34.10	57.16	21.52	20.68	21.90
GraphSAGE	12.68	13	11.04	11.16	12.70
H2GCN	25.52	57.46	13.58	16.18	15.62
SimP-GCN*	35.70	44.86	19.68	18.64	20.68
HOG-GCN*	77.28	246.60	56.46	55.05	53.34
GCN-RARE (ours)	57.44	186.12	16.40	19.38	16.58
GAT-RARE (ours)	66.34	209.88	33.70	26.98	25.77
GraphSAGE-RARE (ours)	41.06	95.04	24.17	28.72	26.11
H2GCN-RARE (ours)	70.61	229.07	22.04	25.09	31.29
Entropy Computation	28.67	266.48	0.0596	0.0615	0.1974

on the final performance of our model, we replace the reward function in Eq. 11 with AUC score. The results in Table V (see GCN-RARE-reward) show that an appropriate reward function is beneficial for guiding the optimization process.

G. Efficiency Study

To further show the efficiency of our GraphRARE, we compare our approach with baselines on the same device as shown in Table VI. Each model is trained for 500 epochs on each dataset and the average training time per epoch is reported. Additionally, we report the computation time required for the node relative entropy calculation, which is computed only once before the model training. The computational cost is comparable to or more efficient than SOTA models (Simp-GCN and HOG-GCN).

H. Convergence of GraphRARE

Since the DRL algorithm and the GNN are trained jointly in the proposed GraphRARE framework, the updating and convergence process is indeed important. We visualize the node classification accuracy during model training, the average reward of reinforcement learning, and the homophily ratio of the graph topology as shown in Figure 6. In Figure 6a, the shadowed area is enclosed by the min and max value of ten cross-validation training runs, and the solid line in the middle is the mean value of each epoch. The result shows that the GNN model converges well when training with the DRL module under the GraphRARE framework. As shown in Figure 6c, the mean episode reward of the DRL module does not update very steadily at the beginning because the accuracy and loss of GNN change a lot. When the framework gradually converges, we can observe that the RL algorithm converges with a stable learning curve. Figure 6b shows the homophily ratio of the graph topology during the training process, the ratio convergences to the mean value of 0.63. The results show that the DRL module outputs a stable number of homogeneous adjacent nodes at the later stage of model training, and GNN also gives a stable classification result. In summary, in the early stage of model training, the accuracy of GNN rises quickly, the homophily ratio of the reconstructed graph varies greatly, and the DRL module receives a large reward value in each episode according to Eq. 11. In the later stage of

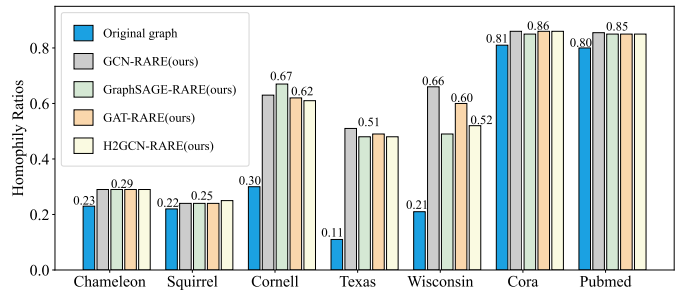


Fig. 7: Comparisons of the homophily ratios between the original graph and the optimized graphs given by the GraphRARE models on the seven datasets.

model training, GNN gives stable classification results, and DRL outputs a stable number of homogeneous adjacent nodes. Since the accuracy and loss of GNN converge to a stable stage, the DRL module converges to the mean reward zero. This suggests that the proposed GraphRARE framework can optimize the graph topology adaptively.

I. Graph Topology Optimization Analysis

The GNNs based on a node message aggregation mechanism are effective under the homophily assumption. Intuitively, the GNN’s performance can be improved by increasing the homophily ratio of a graph. The proposed GraphRARE framework can achieve this goal by training the GNNs and DRL jointly based on the relative entropy to optimize the graph topology. The effectiveness of the proposed GraphRARE framework is demonstrated through the homophily ratio comparison between the original and optimized graph topologies. The GCN-RARE, GraphSAGE-RARE, GAT-RARE, and H2GCN-RARE are evaluated on all seven datasets. The results are presented in Figure 7. The subdued enhancements in the homophily ratio for the Chameleon and Squirrel datasets could be attributed to their intricate graph topology. Notably, these datasets exhibit denser connectivity compared to the other five datasets, with nodes, on average, having about eight times more connections. This increased density complicates the task of identifying clear patterns and understanding node similarities, thereby posing challenges in achieving notable enhancements in the homophily ratio. The four enhanced GraphRARE models all increase the homophily ratio over the original graph by an average of 0.20 (GCN-RARE), 0.17 (GraphSAGE-RARE and GAT-RARE), and 0.18 (H2GCN-RARE). One can see that the optimized graph topologies of the GCN-RARE, GraphSAGE-RARE, GAT-RARE, and H2GCN-RARE improve the homophily ratio, indicating the effectiveness and generalization ability of the proposed framework.

J. Relative Entropy Analysis

Since the node relative entropy computation is an essential module in our framework, we conduct a simple but intuitive visualization experiment to demonstrate its effectiveness. As shown in the Figure 8, we visualize the node relative entropy between node pairs, where deep colors indicate higher entropy

values. Notably, pairs of nodes that share the same label exhibit higher node relative entropy. Consequently, our GraphRARE would connect these high-entropy node pairs, aligning with the homophily assumption.

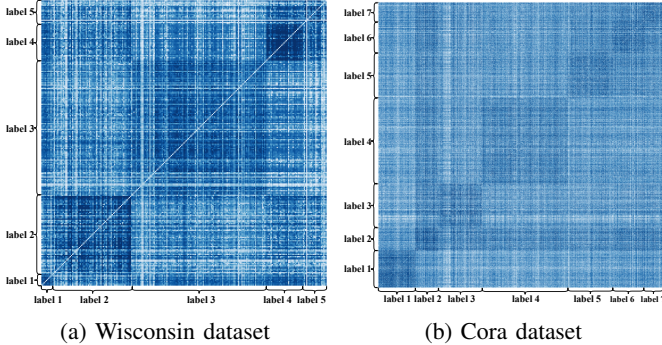


Fig. 8: The visualization of node relative entropy on Wisconsin and Cora datasets. Deeper colors indicate higher entropy value between node pairs.

VI. RELATED WORK

Graph Neural Networks. Most existing GNNs adopt the message passing framework and use permutation invariant local aggregation schemes to update node representations. For instance, GCN [18] averages features of all neighboring nodes. GAT [40] uses an attention mechanism to assign different weights to neighboring nodes. GraphSAGE [11] samples fixed-size neighbors of a node and aggregates their features for realizing fast and scalable GNN training. However, traditional GNNs suffer from a severe performance reduction under the heterophily setting.

Relative Entropy on Graph Data. Relative entropy on graph data is a measure of distance between nodes, which is calculated based on the probability distribution of a node set [3]. This metric can be used to rank nodes by importance [29]. The measurement of relative entropy between nodes has been mainly based on structural parameters, such as the ego node and its neighboring nodes’ degrees. Tian et al. [38] define the node joint information entropy and node mutual information between a node and its remote neighbors. Zhang et al. [50] measure the similarity between nodes based on the relative entropy of node degree. In addition to structural information of nodes, node features are also important metrics in measuring node information as they contain rich semantic information (e.g., the research titles in citation networks). However, fewer studies [26], [42] have considered node features and a local topological structure jointly to measure mutual information between node pairs. Considering node features and structural information, the proposed framework introduces a node relative entropy to measure the similarity between node pairs.

Graph Topology Optimization. In recent years, research efforts have been conducted on GNNs with heterophily. Their main idea is to optimize the original graph topology by extending local neighbors to higher-order neighbors. UGCN [16] connects the $top-k$ similar node pairs to reconstruct the graph

topology by measuring the feature-level cosine similarity between the ego node and remote nodes using the kNN algorithm. In addition, through the node similarity analysis, SimP-GCN [17] selects the $top-k$ similar node pairs in terms of feature similarity to construct new neighboring node sets. Based on the structural information, MI-GCN [38] reconstructs the graph topology by calculating mutual information between node pairs and setting a fixed number of $top-k$ new neighbors and $top-d$ deleted neighbors. However, the $top-k$ selection methods have mostly been based on a heuristic that the value of k is diverse in different graphs and has to be chosen adaptively, which will cost many human efforts and can result in a mixture of useful and irrelevant information of the multi-hop nodes. Additionally, it is crucial to address the presence of noisy edges in the original graph as they can impact the performance of GNNs. The original graph topology optimization methods strongly rely on prior knowledge or require conducting massive experiments to tune the k value. However, these methods fail to achieve optimal results through end-to-end training. To address this shortcoming, the proposed framework adopts a deep reinforcement learning-based algorithm to choose a proper value of k and d for each node in a graph, which can be trained jointly with the GNN in an end-to-end manner.

Deep neural networks is typically modeled as computational graphs, and graph optimization techniques offer structured and efficient computational graph representations that enhance reasoning and elevate training efficiency within the expansive field of database technology. MetaFlow [15] utilizes a rule-based approach to uncover additional optimization opportunities in computation graphs. Tensor Comprehensions [39] uses black-box auto-tuning and polyhedral optimizations. Additionally, Fang et al. [5] propose a pruning algorithm and a dynamic programming approach to refine computation graph optimization further. Thus, the optimized graph topology not only enhances the performance of GNN models on downstream tasks but also contributes to the broader field of database technology by offering a more organized and efficient data representation.

VII. CONCLUSIONS

This work focuses on heterophilic graphs and designs an innovative framework named by GraphRARE. In heterophilic graphs, linked nodes have different features and class labels, whereas the semantically related nodes might be multi-hop away, which can lead to poor GNN’s performance. The main objective of the GraphRARE framework is to discover remote but informative nodes based on the defined node relative entropy. In the framework, the DRL module and GNN are combined to optimize the graph topology in an end-to-end manner. Experimental results demonstrate that the proposed framework can be easily adapted to the existing GNN models and can improve their performance on the benchmarks. In future works, a number of extensions and potential improvements are possible, such as extending GraphRARE to incorporate multi-modal graphs or spatial-temporal graphs.

VIII. ACKNOWLEDGMENTS

This work is supported in part by the National Key R&D Program of China (Grant No.2021ZD0112901), National Natural Science Foundation of China (Grant No.6227073648) and the Beijing Postdoctoral Research Foundation (No.2023-22-97). Zhifeng Bao is supported in part by ARC DP240101211 and DP220101434.

REFERENCES

- [1] S. Abu-El-Haija, B. Perozzi, A. Kapoor, N. Alipourfard, K. Lerman, H. Harutyunyan, G. Ver Steeg, and A. Galstyan, "Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing," in *ICML*, 2019, pp. 21–29.
- [2] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016.
- [3] T. M. Cover, J. A. Thomas *et al.*, "Entropy, relative entropy and mutual information," *Elements of information theory*, vol. 2, no. 1, pp. 12–13, 1991.
- [4] L. Du, X. Shi, Q. Fu, X. Ma, H. Liu, S. Han, and D. Zhang, "GBK-GNN: gated bi-kernel graph neural networks for modeling both homophily and heterophily," in *WWW*, 2022, pp. 1550–1558.
- [5] J. Fang, Y. Shen, Y. Wang, and L. Chen, "Optimizing DNN computation graph using graph substitutions," *Proc. VLDB Endow.*, vol. 13, no. 11, pp. 2734–2746, 2020.
- [6] Z. Fang, L. Xu, G. Song, Q. Long, and Y. Zhang, "Polarized graph neural networks," in *WWW*, 2022, pp. 1404–1413.
- [7] K. Feng, C. Li, X. Zhang, and J. Zhou, "Towards open temporal graph neural networks," *arXiv preprint arXiv:2303.15015*, 2023.
- [8] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019, pp. 1–12.
- [9] A. P. García-Plaza, V. Fresno-Fernández, R. Martínez-Unanue, and A. Zubiaga, "Using fuzzy logic to leverage HTML markup for web page representation," *IEEE Trans. Fuzzy Syst.*, vol. 25, no. 4, pp. 919–933, 2017.
- [10] Z. Gu, K. Zhang, G. Bai, L. Chen, L. Zhao, and C. Yang, "Dynamic activation of clients and parameters for federated learning over heterogeneous graphs," in *ICDE*, 2023, pp. 1597–1610.
- [11] W. L. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *NeurIPS*, vol. 30, 2017, pp. 1024–1034.
- [12] Y. Han, G. Li, H. Yuan, and J. Sun, "An autonomous materialized view management system with deep reinforcement learning," in *ICDE*, 2021, pp. 2159–2164.
- [13] D. He, C. Liang, H. Liu, M. Wen, P. Jiao, and Z. Feng, "Block modeling-guided graph convolutional neural networks," in *AAAI*, 2022, pp. 4022–4029.
- [14] S. Horchidan, "Query optimization for inference-based graph databases," in *VLDB*, ser. CEUR Workshop Proceedings, V. Efthymiou and X. Hu, Eds., vol. 3452, 2023, pp. 33–36.
- [15] Z. Jia, J. Thomas, T. Warszawski, M. Gao, M. Zaharia, and A. Aiken, "Optimizing DNN computation with relaxed graph substitutions," in *Proceedings of Machine Learning and Systems*, 2019.
- [16] D. Jin, Z. Yu, C. Huo, R. Wang, X. Wang, D. He, and J. Han, "Universal graph convolutional networks," in *NeurIPS*, vol. 34, 2021, pp. 10654–10664.
- [17] W. Jin, T. Derr, Y. Wang, Y. Ma, Z. Liu, and J. Tang, "Node similarity preserving graph convolutional networks," in *WSDM*, 2021, pp. 148–156.
- [18] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *ICLR*, 2017, pp. 1–14.
- [19] S. Kullback and R. A. Leibler, "On information and sufficiency," *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [20] C. Li, Y. Tsai, and J. C. Liao, "Graph neural networks for tabular data learning," in *ICDE*, 2023, pp. 3589–3592.
- [21] Y. Li, H. Yuan, Z. Fu, X. Ma, M. Xu, and S. Wang, "ELASTIC: edge workload forecasting based on collaborative cloud-edge deep learning," in *WWW*, Y. Ding, J. Tang, J. F. Sequeda, L. Aroyo, C. Castillo, and G. Houben, Eds., 2023, pp. 3056–3066.
- [22] Y. Li, H. Xiong, L. Kong, S. Wang, Z. Sun, H. Chen, G. Chen, and D. Yin, "Ltrgcn: Large-scale graph convolutional networks-based learning to rank for web search," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2023, pp. 635–651.
- [23] Y. Li, H. Xiong, L. Kong, R. Zhang, F. Xu, G. Chen, and M. Li, "Mhrr: Moocs recommender service with meta hierarchical reinforced ranking," *IEEE Transactions on Services Computing*, 2023.
- [24] J. Lin, "Divergence measures based on the shannon entropy," *IEEE Transactions on Information theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [25] M. Liu, Z. Wang, and S. Ji, "Non-local graph neural networks," 2021.
- [26] G. Luo, J. Li, H. Peng, C. Yang, L. Sun, P. S. Yu, and L. He, "Graph entropy guided node embedding dimension selection for graph neural networks," *IJCAI*, pp. 2767–2774, 2021.
- [27] X. Ma, Q. Chen, Y. Ren, G. Song, and L. Wang, "Meta-weight graph neural network: Push the limits beyond global homophily," in *WWW*, 2022, pp. 1270–1280.
- [28] G. Namata, B. London, L. Getoor, B. Huang, and U. Edu, "Query-driven active surveying for collective classification," in *10th International Workshop on Mining and Learning with Graphs*, vol. 8, 2012, p. 1.
- [29] Y. M. Omar and P. Plapper, "A survey of information entropy metrics for complex networks," *Entropy*, vol. 22, no. 12, p. 1417, 2020.
- [30] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," *NeurIPS*, 2017.
- [31] H. Pei, B. Wei, K. C. Chang, Y. Lei, and B. Yang, "Geom-gcn: Geometric graph convolutional networks," in *ICLR*, 2020, pp. 1–12.
- [32] T. Peng, Y. Liang, W. Wu, J. Ren, Z. Pengrui, and Y. Pu, "Clgt: A graph transformer for student performance prediction in collaborative learning," in *AAAI*, vol. 37, no. 13, 2023, pp. 15947–15954.
- [33] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dornmann, "Stable-baselines3: Reliable reinforcement learning implementations," *JMLR*, vol. 22, no. 268, pp. 1–8, 2021.
- [34] B. Rozemberczki, C. Allen, and R. Sarkar, "Multi-scale attributed node embedding," *Journal of Complex Networks*, vol. 9, no. 2, p. cnab014, 2021.
- [35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017.
- [36] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad, "Collective classification in network data," *AI magazine*, vol. 29, no. 3, pp. 93–93, 2008.
- [37] Z. Shen, C. Hu, and Z. Zhao, "Lynx: A graph query framework for multiple heterogeneous data sources," *Proc. VLDB Endow.*, vol. 16, no. 12, pp. 3926–3929, 2023.
- [38] L. Tian and H. Wu, "MI-GCN: node mutual information-based graph convolutional network," in *WWW*, 2022, pp. 996–1003.
- [39] N. Vasilache, O. Zinenko, T. Theodoridis, P. Goyal, Z. DeVito, W. S. Moses, S. Verdoolaege, A. Adams, and A. Cohen, "Tensor comprehensions: Framework-agnostic high-performance machine learning abstractions," *CoRR*, vol. abs/1802.04730, 2018.
- [40] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *ICLR*, 2018, pp. 1–12.
- [41] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai *et al.*, "Deep graph library: A graph-centric, highly-performant package for graph neural networks," 2019.
- [42] T. Wang, D. Jin, R. Wang, D. He, and Y. Huang, "Powerful graph convolutional networks with adaptive propagation mechanism for homophily and heterophily," in *AAAI*, 2022, pp. 4210–4218.
- [43] C. Wu, C. Wang, J. Xu, Z. Fang, T. Gu, C. Wang, Y. Song, K. Zheng, X. Wang, and G. Zhou, "Instant representation learning for recommendation over large dynamic graphs," in *ICDE*, 2023, pp. 82–95.
- [44] Y. Yan, M. Hashemi, K. Swersky, Y. Yang, and D. Koutra, "Two sides of the same coin: Heterophily and oversmoothing in graph convolutional neural networks," in *ICDM*, 2022, pp. 1287–1292.
- [45] T. Yang, Y. Wang, Z. Yue, Y. Yang, Y. Tong, and J. Bai, "Graph pointer neural networks," in *AAAI*, 2022, pp. 8832–8839.
- [46] Y. Yang, Z. Guan, J. Li, W. Zhao, J. Cui, and Q. Wang, "Interpretable and efficient heterogeneous graph convolutional network," *TKDE*, vol. 35, no. 2, pp. 1637–1650, 2023.
- [47] H. Yuan and G. Li, "A survey of traffic prediction: from spatio-temporal data to intelligent transportation," *Data Sci. Eng.*, vol. 6, no. 1, pp. 63–85, 2021.
- [48] H. Yuan, G. Li, and Z. Bao, "Route travel time estimation on A road network revisited: Heterogeneity, proximity, periodicity and dynamicity," *Proc. VLDB Endow.*, vol. 16, no. 3, pp. 393–405, 2022.
- [49] H. Yuan, G. Li, Z. Bao, and L. Feng, "An effective joint prediction model for travel demands and traffic flows," in *ICDE*, 2021, pp. 348–359.

- [50] Q. Zhang, M. Li, and Y. Deng, "Measure the structure similarity of nodes in complex networks based on relative entropy," *Physica A: Statistical Mechanics and its Applications*, vol. 491, pp. 749–763, 2018.
- [51] Y. Zhang, Q. Yao, L. Yue, X. Wu, Z. Zhang, Z. Lin, and Y. Zheng, "Emerging drug interaction prediction enabled by flow-based graph neural network with biomedical network," *CoRR*, vol. abs/2311.09261, 2023.
- [52] Y. Zhang, Z. Zhou, Q. Yao, X. Chu, and B. Han, "Adaprop: Learning adaptive propagation for graph neural network based knowledge graph reasoning," in *SIGKDD*. ACM, 2023, pp. 3446–3457.
- [53] Y. Zhang, W. Wang, H. Yin, P. Zhao, W. Chen, and L. Zhao, "Disconnected emerging knowledge graph oriented inductive link prediction," in *ICDE*, 2023, pp. 381–393.
- [54] H. Zhao, Q. Yao, and W. Tu, "Search to aggregate neighborhood for graph neural network," in *ICDE*, 2021, pp. 552–563.
- [55] J. Zhao, Z. Zhou, Z. Guan, W. Zhao, W. Ning, G. Qiu, and X. He, "Intentgc: A scalable graph convolution framework fusing heterogeneous information for recommendation," in *SIGKDD*, 2019, pp. 2347–2357.
- [56] S. Zheng, W. Wang, J. Qu, H. Yin, W. Chen, and L. Zhao, "MMKGR: multi-hop multi-modal knowledge graph reasoning," in *ICDE*, 2023, pp. 96–109.
- [57] X. Zheng, Y. Liu, S. Pan, M. Zhang, D. Jin, and P. S. Yu, "Graph neural networks for graphs with heterophily: A survey," 2022.
- [58] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, 2020.
- [59] G. Zhu, Z. Zhu, W. Wang, Z. Xu, C. Yuan, and Y. Huang, "Autoac: Towards automated attribute completion for heterogeneous graph neural network," in *ICDE*, 2023, pp. 2808–2821.
- [60] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, "Beyond homophily in graph neural networks: Current limitations and effective designs," in *NeurIPS*, vol. 33, 2020, pp. 7793–7804.