

Leveraging Invariant Principle for Heterophilic Graph Structure Distribution Shifts

Anonymous Author(s)

Abstract

Heterophilic Graph Neural Networks (HGNNs) have shown promising results for semi-supervised learning tasks on graphs. Notably, most real-world heterophilic graphs are composed of a mixture of nodes with different neighbor patterns, exhibiting local node-level homophilic and heterophilic structures. However, existing works are only devoted to designing better unified HGNN backbones for node classification tasks on heterophilic and homophilic graph benchmarks simultaneously, and their analyses of HGNN performance concerning nodes are only based on the determined data distribution without exploring the effect caused by the difference of structural pattern between training and testing nodes. How to learn invariant node representations on heterophilic graphs to handle this structure difference or distribution shifts remains unexplored. In this paper, we first discuss the limitations of previous graph-based invariant learning methods in addressing the heterophilic graph structure distribution shifts from the perspective of data augmentation. Then, we propose **HEI**, a framework capable of generating invariant node representations through incorporating Heterophily information, node’s estimated neighbor pattern, to infer latent Environments without augmentation, which are then used for Invariant prediction. We provide detailed theoretical guarantees to clarify the reasonability of HEI. Extensive experiments on various benchmarks and backbones can also demonstrate the effectiveness and robustness of our method compared with existing state-of-the-art baselines. Our codes can be accessed through HEI.

Keywords

Graph Representation Learning, Node Classification, Invariant Learning, Distribution Shifts, Heterophily and Homophily

ACM Reference Format:

Anonymous Author(s). 2018. Leveraging Invariant Principle for Heterophilic Graph Structure Distribution Shifts. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym ’XX)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Graph Neural Networks (GNNs) have emerged as prominent approaches for learning graph-structured representations through the aggregation mechanism that effectively combines feature information from neighboring nodes [41]. Previous GNNs primarily dealt

with *homophilic graphs*, where connected nodes tend to share similar features and labels [42]. However, growing empirical evidence suggests that these GNNs’ performance significantly deteriorates when dealing with *heterophilic graphs*, where most nodes connect with others from different classes, even worse than the traditional neural networks [19]. An appealing way to address this issue is to tailor the heterophily property to GNNs, extending the range of neighborhood aggregation and reorganizing architecture [41], known as the heterophilic GNNs (HGNNs).

Heterophilic Graph Structure distribution Shift (HGSS): A novel data distribution shift perspective to reconsider existing HGNNs works. Despite promising, most previous HGNNs assume the nodes share the determined data distribution [18, 19], we argue that there is data distribution disparity among nodes with different neighbor patterns. As illustrated in Figure 1(a1), heterophilic graphs are composed of a mixture of nodes that exhibit local homophilic and heterophilic structures, *i.e.* the nodes have different neighbor patterns [41]. The node’s neighbor pattern can be measured by node homophily, representing homophily level by comparing the label between the node and its neighbors. Here, we identify their varying neighbor patterns between train and test nodes as the Heterophilic Graph Structure distribution Shift (Figure 1(a2)). This kind of shift was neglected by previous works but truly affected GNN’s performance. As shown in Figure 1(a3), we visualize the HGSS between training and testing nodes on the Squirrel dataset. Compared with test nodes, the train nodes are more prone to be categorized into groups with high homophily, which may yield a test performance degradation. More statistical results on other heterophilic graph datasets can be shown in Figure 5. Notably, though some recent work [27] also discusses homophilic and heterophilic structural patterns, until now they haven’t provided a clear technique solution for this problem. Compared with traditional HGNN works that focus on backbone designs, it’s extremely urgent to seek solutions from a data distribution perspective to address the HGSS issue.

Existing graph-based invariant learning methods perform badly for HGSS due to the augmentation-based environment construction strategy. In the context of general distribution shifts, the technique of invariant learning [30] is increasingly recognized for its efficacy in mitigating these shifts. The foundational approach involves learning node representations to facilitate invariant predictor learning across various constructed environments (Figure 1(b1)), adhering to the Risk Extrapolation (REx) principle [9, 24, 37]. Unfortunately, previous graph-based invariant learning methods may not effectively address the HGSS issue, primarily due to explicit environments that may be ineffective for invariant learning. As illustrated in Figure 1(c1), within HGSS settings, altering the original structure does not consistently affect the node’s neighbor patterns. In essence, obtaining optimal and varied environments pertinent to neighbor patterns is challenging. Our observation (Figure 1(c2)) reveals that EERM [37], a pioneering invariant learning approach utilizing environment augmentation to tackle graph distribution shifts in node-level

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym ’XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXX.XXXXXXX>

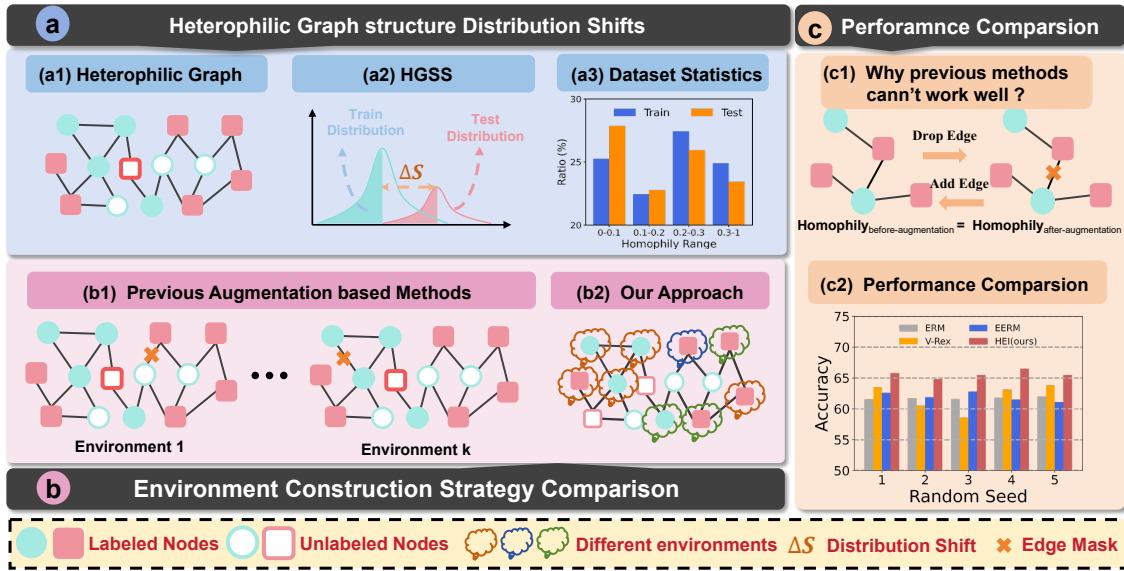


Figure 1: (a) illustrates the heterophilic graph structure distribution shifts (HGSS), where the figure and histogram show the HGSS and neighbor pattern (measured by node homophily) difference between train and test nodes on the Squirrel dataset; (b) displays the comparison of different environment construction strategies between previous invariant learning works and ours from augmentation; (c) shows that the environment construction of previous methods may be ineffective in addressing the HGSS due to the unchanged neighbor pattern distribution. The experimental results between traditional and graph-based invariant learning methods can support our analysis and verify the superiority of our proposed HEI.

tasks, does not perform well under HGSS settings. At times, its enhancements are less effective than simply employing the original V-Rex [16], which involves randomly distributing the train nodes across various environmental groups. We attribute this phenomenon to the irrational environment construction. According to our analysis, EERM is essentially a node environment-augmented version of V-Rex, *i.e.*, the disparity in their performance is solely influenced by the differing strategies in environmental construction. Besides, from the perspective of theory assumption, V-Rex is initially employed to aid model training by calculating the variance of risks introduced by different environments as a form of regularization. The significant improvements by V-Rex also reveal that the nodes of a single input heterophilic graph may reside in distinct environments, considering the variation in neighbor patterns, thus contradicting EERM’s prior assumption that all nodes in a graph share the same environment [37]. Based on this insight, our goal is to break away from previous explicit environment augmentation to learn the latent environment partition, which empowers the invariant learning to address the HGSS better.

HEI: Heterophily-Guided Environment Inference for Invariant Learning. Recent studies explore the effect of prior knowledge on the environment partition [20, 33] and subsequently strengthen the importance of the environment inference and extrapolation for model generalization [36, 39]. Therefore, our initial step should be to quantify the nodes’ neighbor pattern properties related to the HGSS, which are central to the issue at hand. Consequently, a critical question emerges: During the training phase, how can we identify an appropriate metric to estimate the node’s neighbor pattern and leverage it to deduce latent environments to manage this HGSS issue? As previously mentioned, node homophily can assess

the node’s neighbor patterns [19]. Unfortunately, this requires the actual labels of the node and its neighbors, rendering it inapplicable during the training stage due to the potential unlabeled status of neighbor nodes. To cope with this problem, several evaluation metrics pertinent to nodes’ neighbor patterns, including local similarity [7], post-aggregation similarity [26], and SimRank [22], have been introduced. These metrics aim to facilitate node representation learning on heterophilic graphs during the training phase. But these studies primarily concentrate on employing these metrics to help select proper neighbors for improved HGNN architectures, while we aim to introduce a novel invariant learning framework-agnostic backbones to separate the spurious features from selected neighbors, tackling the structure distribution shifts. Therefore, we propose HEI, a framework capable of generating invariant node representations through incorporating heterophily information to infer latent environments, as shown in Figure 1 (b2), which are then used for downstream invariant prediction, under heterophilic graph structure distribution shifts. Extensive experiments on various backbones and benchmarks can verify the effectiveness of our proposed method in addressing this neglected HGSS issue.

Our Contributions: (i) We highlight an important yet often neglected form of heterophilic graph structure distribution shift, which is orthogonal to most HGNN works that focus on backbone designs; (ii) We propose HEI, a novel graph-based invariant learning framework to tackle the HGSS issue. Unlike previous efforts, our method emphasizes leveraging a node’s inherent heterophily information to deduce latent environments without augmentation, thereby significantly improving the generalization and performance of HGNNs; (iii) We demonstrate the effectiveness of HEI on several benchmarks and backbones compared with existing methods.

2 Preliminaries

Notations. Given an input graph $G = (V, X, A)$, we denote $V \in \{v_1, \dots, v_N\}$ as the nodes set, $X \in \mathbb{R}^{N \times D}$ as node features and $A \in \{0, 1\}^{N \times N}$ as an adjacency matrix representing whether the nodes connect, where the N and D denote the number of nodes and features, respectively. The node labels can be defined as $Y \in \{0, 1\}^{N \times C}$, where C represents the number of classes. For each node v , we use A_v and X_v to represent its adjacency matrix and node feature.

Problem Formulation. We first provide the formulation for the general optimized object of node-level OOD (Out of Distribution) problem on graphs, then reclarify the formulation of previous works to help distinguish our work in the next section. From the perspective of data generation, we can get train data (G_{train}, Y_{train}) from train distribution $p(G, Y|e = e)$, the model should handle the test data (G_{test}, Y_{test}) from a different distribution $p(G, Y|e = e')$, varying in different environments e . Thus, the optimized object of node-level OOD problem on graphs can be formulated as follows:

$$\min_{\omega, \Phi} \max_{e \in \mathcal{E}} \mathbb{E}_G \frac{1}{N} \sum_{v \in V} \mathbb{E}_{y \sim p(y|A_v, X_v, e=e)} l(f_\omega(f_\Phi(A_v, X_v)), y_v) \quad (1)$$

where the \mathcal{E} represents the support of environments e , the f_ω and f_Φ refer to GNN's classifier and feature extractor respectively and $l(\cdot)$ is a loss function (e.g. cross entropy). The Eq 1 aims to learn a robust model that minimizes loss across environments as much as possible. Only in this way, can the trained model be likely to adapt to unknown target test distribution well. However, environmental labels for nodes are usually unavailable during the training stage, which inspires many works to seek methods to make use of environmental information to help model training.

Previous Graph-based Invariant Learning. To approximate the optimized object of Eq. 1, previous works [9, 24, 37] mainly construct diverse environments by adopting the masking strategy as Figure 1(b1). Thus, we conclude previous works from the masking strategy ($Mask_\eta(\cdot)$ parameterized with η). Given an input single graph, we can obtain K augmented graphs as Eq. 2, where each graph corresponds to an environment. The K is a pre-defined number of training environments and X^m, A^m , and V^m are corresponding mask versions of feature, adjacency matrix, and node sets.

$$G^{e=k} = Mask_\eta^{e=k}(G) = (X^m, A^m, V^m)_{e=k}, k = 1, 2, \dots, K \quad (2)$$

Then, assisted by these augmented graphs with environment label, the GNN $f(\cdot)$ parameterized by (ω, Φ) can be trained considering environmental information. We can define the ERM (Empirical Risk Minimization) loss in the k -th environment as the Eq.3, which only calculates the loss on the corresponding augmented $G^{e=k}$.

$$R_{e=k}(\omega, \Phi) = \frac{1}{N} \sum_{v \in V^m} l(f_\omega(f_\Phi(A_v, X_v)), y_v) \quad (3)$$

Following the principle of Variance Risk Extrapolation (V-Rex) to reduce the risks from different environments, the final training framework can be defined as Eq. 4. Where the λ controls the effect between reducing the average risk and promoting equality of risks.

$$\min_{\omega, \Phi} \max_{\eta} L(\Phi, \omega, \eta) = \sum_{k=1}^K R_{e=k}(\omega, \Phi) + \lambda Var(R_{e=1}(\omega, \Phi), \dots, R_{e=K}(\omega, \Phi)) \quad (4)$$

The maximization means that we should optimize the masking strategy (parameter η) to construct sufficient and diverse environments, while the minimization aims to reduce the training loss for the model (parameter ω and Φ).

Discussions. Exactly, previous graph-based invariant learning methods introduce extra augmented graphs to construct nodes' environments while our work only infers nodes' environments on a single input graph. Specifically, there exists a latent assumption for previous works that nodes on a single graph belong to the same environment so we need to construct diverse environments by data augmentation. This assumption arises from the insight that nodes on an input graph come from the same *outer domain-related environments* (e.g. Financial graphs or Molecular graphs) [37]. But considering the message-passing mechanism on heterophilic graphs (the ideal aggregation target should be nodes with the same label), the nodes should exist exactly in *inner structure-related environments*. To cope with this issue, as shown in Figure 1(c1), directly utilizing *data augmentation may be ineffective in changing the node's neighbor pattern distribution* to construct diverse environments for invariant prediction. At the same time, the neighbor pattern difference between train and test has verified that even on a single graph, the nodes may belong to different structure-related environments. These simultaneously inspire us to directly infer node environments on a single graph assisted by the node's neighbor pattern, rather than constructing environments from different augmented graphs, for addressing heterophilic graph structure distribution shift.

3 Methodology

In this section, we present the details of the proposed HEI. Firstly, on heterophilic graphs, we verify that the similarity can serve as a neighbor pattern indicator and then review existing similarity-based metrics to estimate the neighbor patterns during training stages. Then, we elaborate the framework to jointly learn environment partition and invariant node representation on heterophilic graphs without augmentation, assisted by the estimated neighbor patterns. Finally, we clarify the overall training process of the algorithm and discuss its complexity. Moreover, we provide a detailed theoretical analysis in Appendix A.2 to clarify the details of HEI.

3.1 Neighbor Patterns Estimation

The node homophily is commonly used to evaluate the node's neighbor patterns, representing the node's structural pattern distribution [19]. Unfortunately, it needs the true labels of the node and its neighbors, which means they can not be used in the training stage because the neighbor nodes may be just the test nodes without labels for node classification tasks when given an input graph. To cope with it, we aim to utilize the similarity between node features to estimate the node's neighbor pattern.

Similarity: An Indicator of Neighbor Patterns. Previous works have shown there exists some relationship between similarity and homophily from the experimental analysis [7], it can not be guaranteed to work well without a theory foundation. Thus, we further investigate its effectiveness from the node cluster view and verify

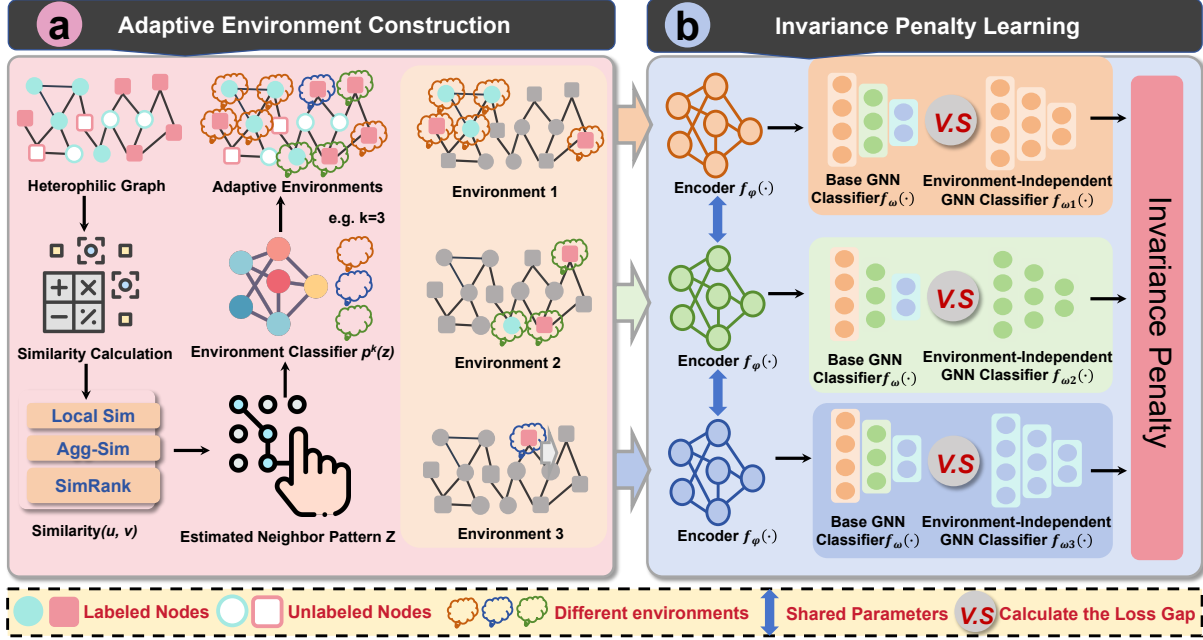


Figure 2: Illustrations of our framework HEI. (a) The neighbor pattern for each train node can be estimated by similarity first and then used for inferring environments without augmentation; (b) Based on the train nodes belonging to different inferred environments, we can train a set of environment-independent GNN classifiers with the shared encoder compared with the base GNN. The shared encoder outputs the representations of nodes in each environment and then forwards them to the base GNN classifier and the environment-independent classifier respectively. By calculating the loss gap between these two different classifiers, an invariance penalty is introduced to improve model generalization.

the similarity between nodes can be exploited to approximate the neighbor pattern without the involvement of label information.

For simplicity, we take K-Means as the cluster algorithm. For two nodes v and u , let v belong to the cluster centroid c_1 and denote the square of the distance between v and u as $\delta = \|u - v\|^2$, we can get $c_1 = \arg \min \|v - c_i\|^2$, where c_i represent the i -th cluster centroid. Then the distance between u and cluster centroid c_1 can be acquired as the Eq. 5. Exactly, the neighbor pattern describes the label relationship between the node and its neighbors. From the Eq 5, we can find the smaller δ , the more likely the v and u belong to the same cluster and own the same label. Therefore, the similarity between nodes can be exploited to serve as a neighbor pattern indicator without using label information.

$$\begin{aligned}
 \|u - c_1\|^2 &= \|(u - v) + (v - c_1)\|^2 \\
 &= \|u - v\|^2 + 2\|u - v\|\|v - c_1\| + \|v - c_1\|^2 \\
 &= \delta + 2\sqrt{\delta}\|v - c_1\| + \|v - c_1\|^2 \\
 &= (\|v - c_1\| + \sqrt{\delta})^2 \geq \delta
 \end{aligned} \tag{5}$$

Existing Similarity-based Metrics. Existing similarity-based metrics on heterophilic graphs can be shown as Eq.6.

$$\text{Similarity}(u, v) = \begin{cases} \text{Sim}(X_v, X_u) & \text{Local Sim} \\ \text{Sim}(\hat{A}_v X_v, \hat{A}_u X_u) & \text{Agg Sim} \\ \frac{1}{|NS(u)||NS(v)|} \sum_{\substack{u' \in NS(u) \\ v' \in NS(v)}} \text{Sim}(X_{u'}, X_{v'}) & \text{SimRank} \end{cases} \tag{6}$$

where the $c \in (0, 1)$ is a decay factor empirically set to 0.6, the $NS(v)$ denotes v 's neighbor set including the nodes connected to v , the \hat{A}_v denotes the aggregation operation on the node v and the Sim denote the similarity calculation between two objects. We can observe that the local similarity (Local Sim [7]) and post-aggregation similarity (Agg-Sim [26]) respectively calculate the similarity of the original and post-aggregation embedding between two nodes. In contrast, the SimRank [22] calculates the similarity between their respective neighbor nodes.

Estimated Node's Neighbor Pattern. Thus, as Eq.7, we can obtain the estimated neighbor patterns z_v for the node v during the training stage by averaging the node's similarity with neighbors.

$$z_v = \frac{1}{|NS(v)|} \sum_{u \in NS(v)} \text{Similarity}(u, v) \tag{7}$$

Notably, we further strengthen our object of using similarity metrics is indeed different from previous HGNN works that utilize the similarity metrics[7, 22, 26] to design backbones. From the perspective of causal analysis shown in Figure 4, when given the neighbors, we aim to separate and weaken the effect of spurious features from full neighbor features by utilizing the estimated neighbor pattern to infer the node's environment for invariant prediction. However, previous HGNN works mainly aim to help the node select proper neighbors and then directly utilize full neighbor features as aggregation targets for better HGNN backbone designs. Our work is exactly *orthogonal* to previous HGNN works.

3.2 HEI: Heterophily-Guided Environment Inference for Invariant Learning

We aim to utilize the estimated neighbor patterns $Z \in R^{G_z}$, which represent the node’s heterophily information, as an auxiliary instrument to jointly learn nodes’ environment partition and invariant node representation without augmentation. Similar techniques can be also shown in [3, 20] for image classification tasks. Specifically, assisted by the estimated neighbor patterns for nodes, we can train an environment classifier $\rho(\cdot) : R^{G_z} \rightarrow R^K$ that softly assigns the train nodes to K environments. The K is a pre-defined number, ρ is a two-layer MLP and the $\rho^{(k)}(\cdot)$ is denoted as the k -th entry of $\rho(\cdot)$, with $\rho(Z) \in [0, 1]^K$ and $\sum_k \rho^{(k)}(Z) = 1$. Denote the ERM loss calculated on all train nodes as $R(\omega, \Phi)$. Then, as shown in Eq. 8, the ERM loss in the k -th inferred environment can be defined as $R_{\rho^{(k)}}(\omega, \Phi)$, which only calculates the loss on the nodes belonging to the k -th environment.

Overall Framework: Based on the above analysis, the training framework of HEI can be defined as follows:

$$R_{\rho^{(k)}}(\omega, \Phi) = \frac{1}{N} \sum_{v \in V} \rho^{(k)}(z_v) l(f_{\omega}(f_{\Phi}(A_v, X_v), y_v)) \quad (8)$$

$$\min_{\omega, \Phi} \max_{\rho, \{\omega_1, \dots, \omega_K\}} L(\Phi, \omega, \omega_1, \dots, \omega_K, \rho) = R(\omega, \Phi) + \lambda \underbrace{\sum_{k=1}^K [R_{\rho^{(k)}}(\omega, \Phi) - R_{\rho^{(k)}}(\omega_k, \Phi)]}_{\text{invariance penalty}} \quad (9)$$

Compared with previous graph-based invariant learning methods shown in Eq. 3 and Eq. 4, our framework mainly differs in the maximization process. Thus, we clarify the effectiveness and reasonability of our framework from two aspects: (i) The invariance penalty learning that introduces a set of environment-dependent GNN classifiers $\{f_{\omega_k}\}_{k=1}^K$, which are only trained on the data belonging to the inferred environments; (ii) The adaptive environment construction through optimizing the environmental classifier $\rho(\cdot)$.

Invariance Penalty Learning. As shown by Eq.1, the ideal GNN classifier f_{ω} is expected to be optimal across all environments. After the environment classifier $\rho^{(k)}(\cdot)$ assigns the train nodes into k inferred environments, we can adopt the following criterion to check if f_{ω} is already optimal in all inferred environments: Take the k -th environment as an example, we can additionally train an environment-dependent classifier f_{ω_k} on the train nodes belonging to the k -th environment. If f_{ω_k} achieves a smaller loss, it indicates that f_{ω} is not optimal in this environment. Moreover, we can further train a set of classifiers $\{f_{\omega_k}\}_{k=1}^K$, each one with a respective individual environment, to assess whether f_{ω} is simultaneously optimal in all environments. Notably, all these classifiers share the same encoder f_{Φ} , if f_{Φ} extracts spurious features that are unstable across the inferred environments, $R_{\rho^{(k)}}(\omega, \Phi)$ will be larger than $R_{\rho^{(k)}}(\omega_k, \Phi)$, resulting in a non-zero invariance penalty, influencing model optimization towards achieving optimality across all environments. In other words, as long as the encoder extracts the invariant feature, the GNN classifier f_{ω} and its related environment-dependent classifier $\{f_{\omega_k}\}_{k=1}^K$ will have the same prediction across

Algorithm 1 HEI: Heterophily-Guided Environment Inference for Invariant Learning

- 1: **Require:** Graph data G and label Y ; Environment classifier ρ ; GNN feature encoder f_{Φ} ; GNN classifier f_{ω} ; Number of training environments: K ; a set of Environment-independent GNN classifiers $\{f_{\omega_k}\}_{k=1}^K$;
 - 2: Estimate neighbor pattern z for each node by Eq. 7;
 - 3: **while** Not converged or maximum epochs not reached **do**
 - 4: Divide the nodes into K environments by $\rho^{(k)}(z)$ and obtain corresponding split graphs $\{G_e = k\}_{k=1}^K$;
 - 5: **for** $k = 1, \dots, K$ **do**
 - 6: Calculate the GNN’s loss on the train nodes belonging to the k -th environment, $R_{\rho^{(k)}}(\omega, \Phi)$, via Eq. 8;
 - 7: Train an additional environment-independent GNN classifier f_{ω_k} with the shared GNN feature encoder f_{Φ} on the train nodes belonging to the k -th inferred environment, calculate its loss $R_{\rho^{(k)}}(\omega_k, \Phi)$;
 - 8: **end for**
 - 9: Calculate invariance penalty and the total loss via Eq. 9;
 - 10: Update ρ via maximizing the invariance penalty;
 - 11: Update f_{Φ}, f_{ω} via minimizing the total loss;
 - 12: **end while**
-

different environments.

Adaptive Environment Construction. As shown in Figure 1(c), the effectiveness of previous methods is only influenced by environmental construction strategy. A natural question arises: What is the ideal environment partition for invariant learning to deal with the HGSS? We investigate it from the optimization of environment classifier $\rho(\cdot)$. Specifically, a good environment partition should construct environments where the spurious features exhibit instability, incurring a large penalty if f_{Φ} extracts spurious features. In this case, we should maximize the invariance penalty to optimize the partition function $\rho(\cdot)$ to generate better environments, which is also consistent with the proposed strategy. Though previous works [9, 24, 37] also adopt the maximization process to construct diverse environments, they just focus on directly optimizing the masking strategy to get augmentation graphs. During the optimization process, these methods lack guidance brought by auxiliary information Z related to environments, ideal or effective environments are often unavailable in this case. That’s why we propose to introduce the environment classifier to infer environments without augmentation, assisted by the Z . Exactly, to make sure the guidance of Z has a positive impact on constructing diverse and effective environments for the invariant node representation learning, there are also two conditions for Z from the causal perspective. We will further clarify it in Appendix A.2.

3.3 Training Process and Complexity Analysis

Training Process: As shown by Algorithm 1: Given a heterophilic graph input, we first estimate the neighbor patterns for each train node by Eq. 7. Then, based on Eq. 8 and Eq. 9, we aim to learn environment partition and invariant node representation, assisted by the estimated neighbor patterns through a min-max alternative

581 optimization. Specifically, maximizing the invariance penalty is
 582 devoted to optimizing the environmental classifier to construct as
 583 diverse environments as possible to enlarge the loss gap between
 584 the base GNN feature encoder and additionally introduced GNN
 585 feature encoders. In contrast, the minimization of total loss aims to
 586 promote the base GNN to learn invariant representation agnostic
 587 neighbor patterns to address the HGSS issues.

589 **Complexity Analysis:** Given a graph with N nodes, the average
 590 degree is d . GNN with l layers calculate embeddings in time and
 591 space $O(Nld^2)$. HEI assigns N nodes into k inferred environments
 592 $N_{e=1} + \dots + N_{e=k} = N$ and executes $k + 1$ classifier computations,
 593 where the k corresponds to the k environment-independent clas-
 594 sifiers, and 1 refers to the basic GNN classifier. Denote N' as the
 595 average number of nodes belonging to an inferred environment,
 596 the overall time complexity is $O(Nld^2 + kN'ld^2)$, which is linear
 597 to the scale of the graph. More detailed efficiency studies compared
 598 with previous methods can be shown in Experiments.

600 4 Experiments

601 In this section, we investigate the effectiveness of HEI to answer
 602 the following questions.

- 603 • **RQ1:** Does HEI outperform state-of-art methods to address
 604 the HGSS issue?
- 605 • **RQ2:** How robust is the proposed method? Can HEI solve
 606 the problem that exists in severe distribution shifts?
- 607 • **RQ3:** How do different similarity-based metrics influence
 608 the neighbor pattern estimation, so as to further influence
 609 the effect of HEI?
- 610 • **RQ4:** What is the sensitivity of HEI concerning the pre-
 611 defined number of training environments?
- 612 • **RQ5:** How efficient is the proposed HEI compared with
 613 previous methods?

614 4.1 Experimental Setup

615 **Datasets.** We adopt six commonly used heterophilic graph datasets
 616 (chameleon, Squirrel, Actor, Penn94, arxiv-year, and twitch-gamer)
 617 and three homophilic graph datasets (Cora, CiteSeer and PubMed)
 618 to verify the effectiveness of HEI [19, 28]. To make sure the eval-
 619 uation is stable and reasonable, we utilize the filtered versions of
 620 existing datasets to avoid data leakage [29]. Notably, considering
 621 that we should further split the test datasets to construct different
 622 evaluation settings. Those excessive small-scale heterophilic graph
 623 datasets, such as Texa, Cornell, and Wisconsin [28], are not fit and
 624 chosen for evaluation due to their unstable outcomes. Moreover,
 625 considering the nodes on homophilic graphs means there exists
 626 a mere structure difference between train and test nodes. We just
 627 provide experiments and discussions in the Appendix A.3.

628 **Settings.** Based on previous dataset splits, we construct two differ-
 629 ent settings to evaluate the effectiveness and robustness of HEI: **(i)**
 630 **Standard Settings:** We sort the test nodes based on their nodes'
 631 homophily values and acquire the median. The part that is higher
 632 than the median is defined as the High Hom Test, while the rest is
 633 defined as the Low Hom Test. The model is trained on the previous
 634 train dataset and evaluated on more fine-grained test groups; **(ii)**

635 Simulation Settings where exists severe distribution shifts.:

636 We sort and split the train and test nodes simultaneously adopting
 637 the same strategy of (i). The model is trained on the Low/High Hom
 638 Train and evaluated on the High/Low Hom Test.

639 **Backbones.** To further verify our framework is orthogonal to pre-
 640 vious HGNN works that focus on backbone designs, we adapt HEI
 641 to two existing SOTA and scalable backbones with different founda-
 642 tions, LINKX (MLP-based) [19] and GloGNN++ (GNN-based) [18].
 643 In this way, our improvements can be attributed to the design that
 644 deals with the neglected heterophilic structure distribution shifts.

645 **Baselines.** Denote the results of the backbone itself as ERM. Our
 646 comparable baselines can be categorized into: (i) Reweight-based
 647 methods considering structure information: Renode [4] and StruRW-
 648 Mixup [23]; (ii) Invariant Learning methods involving environment
 649 inference for node-level distribution shift: SRGNN [43], EERM [37],
 650 BAGNN [9], FLOOD [24], CaNet [36] and IENE [39]; (iii) Prototype-
 651 based methods for structural distribution shift on the special do-
 652 main(e.g. graph anomaly detection): GDN [11]. Notably, though we
 653 can utilize estimated neighbor patterns as auxiliary information to
 654 infer environments related to HGSS, the true environment label is
 655 still unavailable. So we don't compare with those traditional invari-
 656 ant learning methods that rely on the explicit environment labels,
 657 e.g. IRM [2], V-Rex [16] and GroupDRO [31].

658 4.2 Experimental Results and Analysis

659 Handling Distribution Shifts under Standard Settings (RQ1).

660 We first evaluate the effectiveness of HEI under standard settings,
 661 where we follow the previous dataset splits and further evaluate
 662 the model on more fine-grained test groups with low and high
 663 homophily, respectively. The results can be shown in Table 1 and
 664 Table 2. We have the following observations.

665 On the one hand, *the impact brought by the HGSS is still apparent
 666 though we adopted the existing SOTA HGNN backbones.* As shown by
 667 the results of ERM in Table 1 and Table 2, for most datasets, there
 668 are significant performance gaps between the High Hom Test and
 669 Low Hom Test, ranging from 5 to 30 scores. These results further
 670 verify the necessity to seek methods from the perspective of data
 671 distribution rather than backbone designs to deal with this problem.

672 On the other hand, *HEI can outperform previous methods in most
 673 circumstances.* Specifically, compared with invariant learning meth-
 674 ods, though HEI does not augment the training environments, uti-
 675 lizing the estimated neighbor patterns to directly infer latent en-
 676 vironments still benefits invariant prediction and improves model
 677 generalization on different test distributions related to homophily.
 678 In contrast, directly adopting a reweight strategy (Renode and
 679 StruRW) or evaluating the difference between the training domain
 680 and target domain (SRGNN) without environment augmentation
 681 can't acquire superior results than invariant learning methods. This
 682 is because these methods need accurate domain knowledge or struc-
 683 ture information in advance to help model training. However, for
 684 the HGSS issue, the nodes' environments on heterophily graphs
 685 are unknown and difficult to split into the invariant and spuri-
 686 ous domains, like the GOOD dataset [12] which has clear domain
 687 and distribution splits. Simultaneously, the neighbor pattern dis-
 688 tribution represents more fine-grained label relationships between

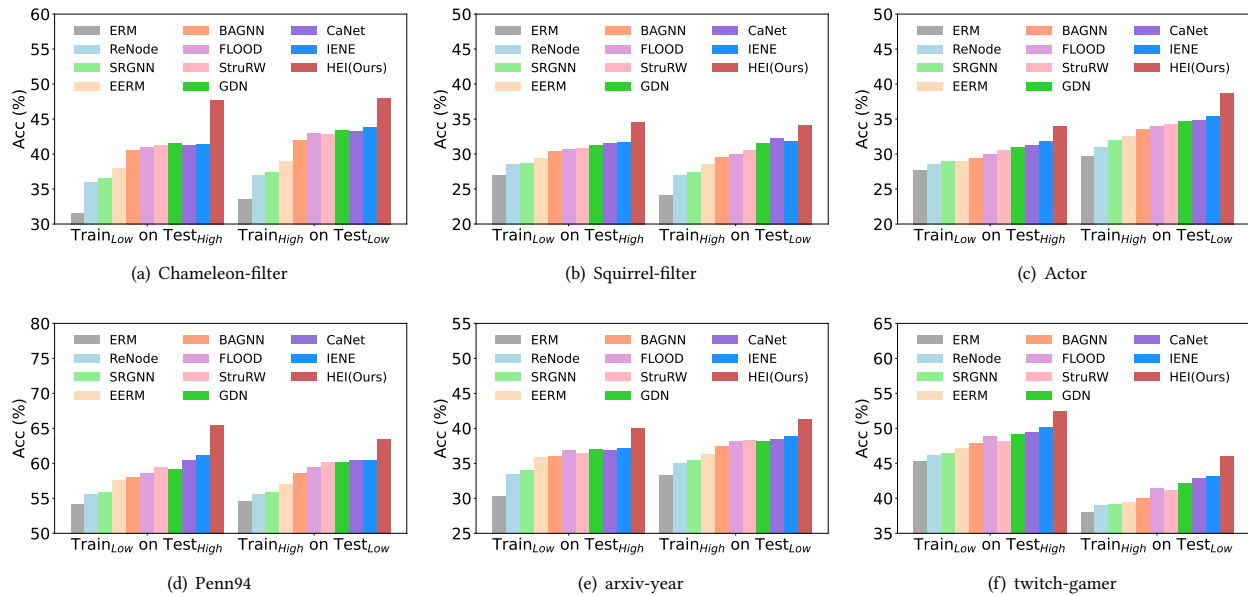


Figure 3: Comparison experiments under Simulation Settings where exists severe distribution shift between train and test, including $Train_{High}$ on $Test_{Low}$ and $Train_{Low}$ on $Test_{High}$. We adopt the LINKX as the backbone there.

Table 3: Comparison experiments on small-scale datasets when we respectively adopt local similarity(Local Sim), post-aggregation similarity (Agg Sim), and SimRank as indicators to estimate nodes' neighbor patterns so as to infer latent environments under Standard settings.

Backbones	Methods	Chameleon-filter			Squirrel-filter			Actor		
		Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test
LINKX	HEI (Local Sim)	24.91 ± 3.68	30.84 ± 3.52	18.95 ± 3.77	38.99 ± 1.85	46.81 ± 1.87	30.15 ± 1.95	36.84 ± 1.75	38.99 ± 2.44	34.51 ± 1.88
	HEI (Agg Sim)	25.11 ± 3.81	30.99 ± 3.92	19.17 ± 3.99	39.15 ± 1.99	47.15 ± 1.97	30.85 ± 2.11	36.87 ± 1.54	39.11 ± 2.48	34.68 ± 1.74
	HEI (SimRank)	25.78 ± 2.23	31.35 ± 2.56	20.21 ± 4.21	40.92 ± 1.31	48.82 ± 2.88	31.21 ± 2.79	37.41 ± 1.17	39.31 ± 1.45	35.42 ± 1.54
GloGNN++	HEI (Local Sim)	27.78 ± 3.51	33.25 ± 3.41	22.51 ± 2.17	37.55 ± 1.55	44.11 ± 1.41	30.11 ± 1.28	38.57 ± 1.59	41.84 ± 1.44	35.41 ± 2.41
	HEI (Agg Sim)	28.24 ± 3.27	33.74 ± 3.11	22.94 ± 2.27	37.99 ± 1.25	44.54 ± 1.28	30.66 ± 1.51	38.81 ± 1.64	41.95 ± 1.55	35.81 ± 2.23
	HEI (SimRank)	29.31 ± 3.68	34.35 ± 3.52	24.25 ± 2.71	39.42 ± 1.51	45.19 ± 1.57	31.45 ± 1.68	39.41 ± 1.51	42.25 ± 1.59	36.12 ± 1.85

The effect of different similarity metrics as neighbor pattern indicators for HEI (RQ3). As depicted in Table 3 and Table 7 in the Appendix, we can draw two conclusions: (i) Applying any similarity-based metrics can outperform the previous SOTA strategy. This verifies the flexibility and effectiveness of HEI and helps distinguish HEI from previous HGNN works that also utilize the similarity for backbone designs; (ii) Applying SimRank to HEI can acquire consistently better performance than other metrics. This can be explained by previous HGNN backbone designs [7, 22, 26], which have verified that SimRank has a better ability to distinguish neighbors patterns compared with Local Sim and Agg-Sim, so as to design a better HGNN backbone, SIMGA[22]. Moreover, from the perspective of definitions as Eq. 6, the SimRank is specifically designed considering structural information, which is more related to our problem considering structure-related distribution shifts.

Sensitivity Analysis of HEI concerning the pre-defined environmental numbers K (RQ4). As shown in Figure 7 in Appendix A.3, we vary the hyper-parameter environment numbers k in Eq. 9 within the range of [2, 12], and keep all other configurations unchanged to explore its impact. We can observe that K has a stable effect on the HEI, especially when $k \geq 6$, which can further verify its effectiveness in addressing HGSS issues.

Efficiency Studies (RQ5). As shown in Table 5 in Appendix A.3, we provide the time to train the model until converges that keep the stable accuracy score on the validation set. Experiments of these large-scale datasets are conducted on a single Tesla V100 GPU with 32G memory and use AdamW as the optimizer following [18]. From the results, we can conclude that the extra time cost can be acceptable compared with the backbone itself and other baselines.

5 Conclusion

In this paper, we emphasize an overlooked yet important variety of graph structure distribution shifts that exist on heterophilic graphs. We verify that previous node-level invariant learning solutions with environment augmentation are ineffective due to the irrationality of constructing environments. To mitigate the effect of this distribution shift, we propose HEI, a framework capable of generating invariant node representations by incorporating the estimated neighbor pattern information to infer latent environments without augmentation, which are then used for downstream invariant learning. Experiments on several benchmarks and backbones demonstrate the effectiveness of our method to cope with this graph structure distribution shift. Finally, we hope this study can draw attention to the structural distribution shift of heterophilic graphs.

References

- [1] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. 2019. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*. PMLR, 21–29.
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893* (2019).
- [3] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. 2020. Invariant rationalization. In *International Conference on Machine Learning*. PMLR, 1448–1458.
- [4] Deli Chen, Yankai Lin, Guangxiang Zhao, Xuancheng Ren, Peng Li, Jie Zhou, and Xu Sun. 2021. Topology-imbalance learning for semi-supervised node classification. *Advances in Neural Information Processing Systems* 34 (2021), 29885–29897.
- [5] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. 2020. Simple and deep graph convolutional networks. In *International conference on machine learning*. PMLR, 1725–1735.
- [6] Yongqiang Chen, Yatao Bian, Kaiwen Zhou, Binghui Xie, Bo Han, and James Cheng. 2023. Does Invariant Graph Learning via Environment Augmentation Learn Invariance? *arXiv preprint arXiv:2310.19035* (2023).
- [7] Yuhan Chen, Yihong Luo, Jing Tang, Liang Yang, Siya Qiu, Chuan Wang, and Xiaochun Cao. 2023. LSGNN: Towards General Graph Neural Network in Node Classification by Local Similarity. *arXiv preprint arXiv:2305.04225* (2023).
- [8] Yongqiang Chen, Yonggang Zhang, Yatao Bian, Han Yang, MA Kaili, Binghui Xie, Tongliang Liu, Bo Han, and James Cheng. 2022. Learning causally invariant representations for out-of-distribution generalization on graphs. *Advances in Neural Information Processing Systems* 35 (2022), 22131–22148.
- [9] Zhengyu Chen, Teng Xiao, and Kun Kuang. 2022. Ba-gnn: On learning bias-aware graph neural network. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 3012–3024.
- [10] Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. 2020. Adaptive universal generalized pagerank graph neural network. *arXiv preprint arXiv:2006.07988* (2020).
- [11] Yuan Gao, Xiang Wang, Xiangnan He, Zhengguang Liu, Huamin Feng, and Yongdong Zhang. 2023. Alleviating structural distribution shift in graph anomaly detection. In *Proceedings of the sixteenth ACM international conference on web search and data mining*. 357–365.
- [12] Shurui Gui, Xiner Li, Limei Wang, and Shuiwang Ji. 2022. Good: A graph out-of-distribution benchmark. *Advances in Neural Information Processing Systems* 35 (2022), 2059–2073.
- [13] Dongxiao He, Chungong Liang, Huixin Liu, Mingxiang Wen, Pengfei Jiao, and Zhiyong Feng. 2022. Block modeling-guided graph convolutional neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 4022–4029.
- [14] Di Jin, Zhizhi Yu, Cuiying Huo, Rui Wang, Xiao Wang, Dongxiao He, and Jiawei Han. 2021. Universal graph convolutional networks. *Advances in Neural Information Processing Systems* 34 (2021), 10654–10664.
- [15] Wei Jin, Tyler Derr, Yiqi Wang, Yao Ma, Zitao Liu, and Jiliang Tang. 2021. Node similarity preserving graph convolutional networks. In *Proceedings of the 14th ACM international conference on web search and data mining*. 148–156.
- [16] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. 2021. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*. PMLR, 5815–5826.
- [17] Haoyang Li, Ziwei Zhang, Xin Wang, and Wenwu Zhu. 2022. Learning invariant graph representations for out-of-distribution generalization. *Advances in Neural Information Processing Systems* 35 (2022), 11828–11841.
- [18] Xiang Li, Renyu Zhu, Yao Cheng, Caihua Shan, Siqiang Luo, Dongsheng Li, and Weining Qian. 2022. Finding global homophily in graph neural networks when meeting heterophily. In *International Conference on Machine Learning*. PMLR, 13242–13256.
- [19] Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and Ser Nam Lim. 2021. Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods. *Advances in Neural Information Processing Systems* 34 (2021), 20887–20902.
- [20] Yong Lin, Shengyu Zhu, Lu Tan, and Peng Cui. 2022. ZIN: When and How to Learn Invariance Without Environment Partition? *Advances in Neural Information Processing Systems* 35 (2022), 24529–24542.
- [21] Gang Liu, Tong Zhao, Jiaxin Xu, Tengfei Luo, and Meng Jiang. 2022. Graph rationalization with environment-based augmentations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1069–1078.
- [22] Haoyu Liu, Ningyi Liao, and Siqiang Luo. 2023. SIMGA: A Simple and Effective Heterophilous Graph Neural Network with Efficient Global Aggregation. *arXiv preprint arXiv:2305.09958* (2023).
- [23] Shikun Liu, Tianchun Li, Yongbin Feng, Nhan Tran, Han Zhao, Qiang Qiu, and Pan Li. 2023. Structural re-weighting improves graph domain adaptation. In *International Conference on Machine Learning*. PMLR, 21778–21793.
- [24] Yang Liu, Xiang Ao, Fuli Feng, Yunshan Ma, Kuan Li, Tat-Seng Chua, and Qing He. 2023. FLOOD: A flexible invariant learning framework for out-of-distribution generalization on graphs. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1548–1558.
- [25] Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. 2021. Is heterophily a real nightmare for graph neural networks to do node classification? *arXiv preprint arXiv:2109.05641* (2021).
- [26] Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. 2022. Revisiting heterophily for graph neural networks. *Advances in neural information processing systems* 35 (2022), 1362–1375.
- [27] Haitao Mao, Zhikai Chen, Wei Jin, Haoyu Han, Yao Ma, Tong Zhao, Neil Shah, and Jiliang Tang. 2023. Demystifying Structural Disparity in Graph Neural Networks: Can One Size Fit All? *arXiv preprint arXiv:2306.01323* (2023).
- [28] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. 2020. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287* (2020).
- [29] Oleg Platonov, Denis Kuznetsov, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. 2023. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? *arXiv preprint arXiv:2302.11640* (2023).
- [30] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2019. Droppedge: Towards deep graph convolutional networks on node classification. *arXiv preprint arXiv:1907.10903* (2019).
- [31] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731* (2019).
- [32] Susheel Suresh, Vinit Budde, Jennifer Neville, Pan Li, and Jianzhu Ma. 2021. Breaking the limit of graph neural networks by improving the assortativity of graphs with local mixing patterns. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1541–1551.
- [33] Xiaoyu Tan, Lin Yong, Shengyu Zhu, Chao Qu, Xihe Qiu, Xu Yinghui, Peng Cui, and Yuan Qi. 2023. Provably invariant learning without domain information. In *International Conference on Machine Learning*. PMLR, 33563–33580.
- [34] Tao Wang, Di Jin, Rui Wang, Dongxiao He, and Yuxiao Huang. 2022. Powerful graph convolutional networks with adaptive propagation mechanism for homophily and heterophily. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 4210–4218.
- [35] Yu Wang and Tyler Derr. 2021. Tree decomposed graph neural network. In *Proceedings of the 30th ACM international conference on information & knowledge management*. 2040–2049.
- [36] Qitian Wu, Fan Nie, Chenxiao Yang, Tianyi Bao, and Junchi Yan. 2024. Graph out-of-distribution generalization via causal intervention. In *Proceedings of the ACM on Web Conference 2024*. 850–860.
- [37] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. 2022. Handling distribution shifts on graphs: An invariance perspective. *arXiv preprint arXiv:2202.02466* (2022).
- [38] Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. 2022. Discovering invariant rationales for graph neural networks. *arXiv preprint arXiv:2201.12872* (2022).
- [39] Haoran Yang, Xiaobing Pei, and Kai Yuan. 2024. IENE: Identifying and Extrapolating the Node Environment for Out-of-Distribution Generalization on Graphs. *arXiv preprint arXiv:2406.00764* (2024).
- [40] Ruihao Zhang, Zhengyu Chen, Teng Xiao, Yueyang Wang, and Kun Kuang. 2024. Discovering Invariant Neighborhood Patterns for Heterophilic Graphs. *arXiv preprint arXiv:2403.10572* (2024).
- [41] Xin Zheng, Yixin Liu, Shirui Pan, Miao Zhang, Di Jin, and Philip S Yu. 2022. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082* (2022).
- [42] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems* 33 (2020), 7793–7804.
- [43] Qi Zhu, Natalia Ponomareva, Jiawei Han, and Bryan Perozzi. 2021. Shift-robust gnn: Overcoming the limitations of localized graph training data. *Advances in Neural Information Processing Systems* 34 (2021), 27965–27977.

1045 A Appendix

1046 In this appendix, we provide the details omitted in the main text due
1047 to the page limit, offering additional experimental results, analyses,
1048 proofs, and discussions.

- 1049 • A.1: We provide a detailed literature review related to our
1050 work and further distinguish our work with these works to
1051 clarify our contribution.
- 1052 • A.2: We provide detailed theoretical analysis from the causal
1053 perspective to clarify the reasonability of HEI: Why we can
1054 utilize the estimated neighbor pattern to infer environments
1055 to address the HGSS issue?
- 1056 • A.3: We provide the full experimental results on small and
1057 large-scale datasets adopting LINKX and GloGNN as back-
1058 bones (4.2 of the main paper).
- 1059 • A.4: We provide detailed implementation details to repro-
1060 duce our experiments.

1061 A.1 Related Work

1062 **Graph Neural Networks with Heterophily.** Existing strategies
1063 for mitigating graph heterophily issues can be categorized into two
1064 groups [41]: (i) Non-Local Neighbor Extension, aimed at identi-
1065 fying suitable neighbors through mixing High-order information
1066 [1, 14, 42] or discovering potential neighbors assisted by various
1067 similarity-based metrics [13–15, 22, 34, 35]; (ii) GNN Architecture
1068 Refinement, focusing on harnessing information derived from the
1069 identified neighbors, through selectively aggregating distinguish-
1070 able and discriminative node representations, including adapting ag-
1071 gregation scheme [18, 25, 32], separating Ego-neighbor [25, 32, 42]
1072 and combining inter-layer [5, 10, 42].

1073 However, these efforts share the common objective of designing
1074 better unified HGNN backbones for node classification tasks on
1075 heterophilic and homophilic graph benchmarks simultaneously.
1076 Moreover, we also found a recent arxiv paper called INPL [40] that
1077 coincides with our work. Though it also mentions the distribution
1078 shifts of neighborhood patterns, it still focuses on proposing an
1079 adaptive Neighborhood Propagation (ANP) module to *optimize the*
1080 *HGNN backbone or architecture without a thorough analysis of pre-*
1081 *vious graph-based invariant-learning methods.* Different from these
1082 works, we instead consider from an identifiable neighbor pattern
1083 distribution perspective and propose a novel invariant learning
1084 framework that can be integrated with most HGNN backbones to
1085 further enhance their performance and generalization.

1086 **Generalization on GNNs.** Many efforts have been devoted to
1087 exploring the generalization ability of GNNs. (i) **For graph-level**
1088 **tasks**, it assumes that every graph can be treated as an instance
1089 for prediction tasks [37]. Many works propose to identify invari-
1090 ant sub-graphs that decide the label Y and spurious sub-graphs
1091 related to environments, such as CIGA [8], GIL [17], GREA [21],
1092 DIR [38] and GALA [6] (ii) **However, for node-level tasks that**
1093 **we focus on in this paper**, the nodes are interconnected in a
1094 graph as instances in a non-iid data generation way, it is not feasi-
1095 ble to transfer graph-level strategies directly. To address this issue,
1096 EERM [37] proposes to regard the node’s ego-graph with corre-
1097 sponding labels as instances and assume that all nodes in a graph
1098 often share the same environment, so it should construct different
1099 environments by data augmentation, *e.g.*, DropEdge [30]. Based
1100 on these findings, BA-GNN [9], FLOOD [24] and IENE [39] inherit
1101 this assumption to improve model generalization. Apart from these
1102 environments-augmentation methods, the SR-GNN [43] and GDN
1103 [11] are two works that address distribution shifts on node-level
1104 tasks from the domain adaption and prototype learning perspec-
1105 tives respectively. Moreover, Renode [4] and StruRW-Mixup [23]
1106 are two reweight-based methods that explore the effect brought by
1107 the structure difference between nodes for node classification tasks.
1108 We also compare them in experiments.

1109 Unlike these works, we highlight a special variety of structure-
1110 related distribution shifts for node classification tasks on heterophilic
1111 graphs and propose a novel invariant learning framework adapted
1112 to heterophilic graphs without dealing with graph augmentation
1113 to address this problem.

1114 **A.2 Theoretical Analysis**

1115 To help understand our framework well, we first provide the com-
1116 parison between our work and previous graph-based invariant
1117 learning works as shown in Figure 4. Specifically, the definitions
1118 of random variables can be defined as follows: We define \mathbf{G} as a
1119 random variable of the input graph, \mathbf{A} as a random variable of
1120 node’s neighbor information, \mathbf{X} as a random variable of node’s fea-
1121 tures, and \mathbf{Y} as a random variable of node’s label vectors. Both node
1122 features \mathbf{X} and node neighbor information \mathbf{A} consist of invariant
1123 predictive information that determines the label Y and the spurious
1124 information influenced by latent environments \mathbf{e} . In this case, we
1125 can denote $\mathbf{X} = [\mathbf{X}^I, \mathbf{X}^S]$ and $\mathbf{A} = [\mathbf{A}^I, \mathbf{A}^S]$.

1126 A.2 Theoretical Analysis

1127 Then, we provide a more detailed theoretical analysis of our
1128 framework from a casual perspective to identify invariant features.
1129 **Casual conditions of Z .** From the casual perspective, some
1130 conditions are also needed for Z to make sure our framework to
1131 address the heterophilic graph structure distribution shifts well [20].
1132 Denote the $H(Y|X, A)$ as the expected loss of an optimal classifier
1133 over $(X, A, \text{ and } Y)$, we can clarify the reasonability of utilizing the
1134 estimated neighbor patterns as Z auxiliary information to infer
1135 environments for invariant prediction based on the following two
1136 conditions.

1137 **CONDITION 1 (INVARIANCE PRESERVING CONDITION).** *Given in-*
1138 *variant feature $(X^I \text{ and } A^I)$ and any function $\rho(\cdot)$, it holds that*

$$1139 H(Y|(X^I, A^I), \rho(Z)) = H(Y|(X^I, A^I)). \quad (10)$$

1140 **CONDITION 2 (NON-INVARIANCE DISTINGUISHING CONDITION).**
1141 *For any feature $X^{Sk} \in X^S$ or $A^{Sk} \in A^S$, there exists a function $\rho(\cdot)$*
1142 *and a constant $C > 0$ satisfy:*

$$1143 H(Y|(X^{Sk}, A^{Sk})) - H(Y|(X^{Sk}, A^{Sk}), \rho(Z)) \geq C. \quad (11)$$

1144 Condition 1 requires that invariant features X^I and A^I should
1145 keep invariant under any environment split obtained by $\rho(Z)$.
1146 Otherwise, if there exists a split where an invariant feature be-
1147 comes non-invariant, then this feature would introduce a positive
1148 penalty as shown in Eq. 9 to further promote the learning of invari-
1149 ant node representation. Exactly, Condition 1 can be met only if
1150 $H(Y|(X^I, A^I), Z) = H(Y|(X^I, A^I))$, which means the auxiliary vari-
1151 able Z should be d-separated by invariant feature X^I and A^I . We
1152 provide a detailed proof in the appendix A.2. Exactly, the estimated
1153 environments by data augmentation, *e.g.*, DropEdge [30]. Based
1154 on these findings, BA-GNN [9], FLOOD [24] and IENE [39] inherit
1155 this assumption to improve model generalization. Apart from these
1156 environments-augmentation methods, the SR-GNN [43] and GDN
1157 [11] are two works that address distribution shifts on node-level
1158 tasks from the domain adaption and prototype learning perspec-
1159 tives respectively. Moreover, Renode [4] and StruRW-Mixup [23]
1160 are two reweight-based methods that explore the effect brought by
1161 the structure difference between nodes for node classification tasks.
1162 We also compare them in experiments.

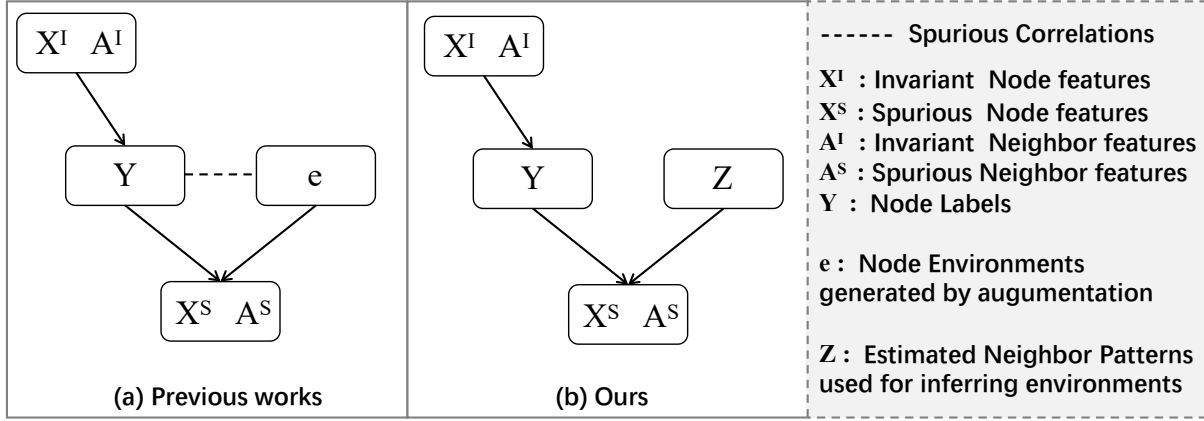


Figure 4: Comparisons between our work and previous graph-based invariant learning works from the causal perspective. Notably, the basic HGNN directly aggregates the selected neighbors’ full features without further separating like the above two types of invariant learning methods.

neighbor pattern just describes the similarity between the node and its neighbors as Eq. 7, while the label Y only has the direct causal relationship with X^I and A^I from the causal perspective. This means the Condition 1 can be met by adopting the estimated neighbor pattern as auxiliary information Z to construct environments.

Condition 2 reveals that for each spurious feature X^S and A^S , there exists at least one environment split where this feature demonstrates non-invariance within the split environment. If a spurious feature doesn’t cause invariance penalties in all environment splits, it can’t be distinguished from true invariant features. As shown in Figure 1(c2), the results of V-Rex are better than ERM, which means even randomly split environments with seeds can have a positive effect on making spurious features produce effective invariance penalty, further promoting the learning of invariant features. It’s more likely to construct comparable or better environments than random seeds under the guidance of the estimated neighbor patterns Z . Thus, Condition 2 can also be guaranteed under our defined heterophilic graph structure distribution shift.

Proof of Meeting Condition 1. We show that for all $\rho(\cdot)$, if $H(Y|(X^I, A^I), Z) = H(Y|(X^I, A^I))$ holds, then there will exist that $H(Y|(X^I, A^I), \rho(Z)) = H(Y|(X^I, A^I))$.

PROOF. On one hand, because $\rho(Z)$ contains less information than Z , we have

$$H(Y|(X^I, A^I), \rho(Z)) \geq H(Y|(X^I, A^I), Z) = H(Y|(X^I, A^I)).$$

On the other hand, (X^I, A^I) and $\rho(Z)$ contain more information than (X^I, A^I) , so we can get

$$H(Y|(X^I, A^I), \rho(Z)) \leq H(Y|(X^I, A^I)).$$

Thus, we conclude $H(Y|(X^I, A^I), \rho(Z)) = H(Y|(X^I, A^I))$.

Assumptions and Theorem to Identify Invariant Features.

In particular, our assumption is consistent with previous invariant learning [20]. So we provide the previous version to support our framework, where $X = [X_v; X_s]$, where the X_v refers to the invariant feature and the X_s refers to the spurious feature.

ASSUMPTION 1. For a given feature mask Φ and any constant $\epsilon > 0$, there exists $f \in F$ such that $E[l(f(\Phi(X)), Y)] \leq H(Y|\Phi(X)) + \epsilon$.

ASSUMPTION 2. If a feature violates the invariance constraint, adding another feature would not make the penalty vanish, i.e., there exists a constant $\delta > 0$ so that for spurious feature $X_1 \subset X_s$ and any feature $X_2 \subset X$,

$$H(Y|X_1, X_2) - H(Y|\rho(Z), X_1, X_2) \geq \delta (H(Y|X_1) - H(Y|\rho(Z), X_1)).$$

ASSUMPTION 3. For any distinct features X_1, X_2 , $H(Y|X_1, X_2) \leq H(Y|X_1) - \gamma$ with fixed $\gamma > 0$.

Exactly, Assumption 1 is a common assumption that requires the function space F be rich enough such that, given Φ , there exists $f \in F$ that can fit $P(Y|\Phi(X))$ well. Assumption 2 aims to ensure a sufficient positive penalty if a spurious feature is included. Assumption 3 indicates that any feature contains some useful information w.r.t. Y , which cannot be explained by other features. Otherwise, we can simply remove such a feature, as it does not affect prediction.

The theorem to identify invariant features can be defined:

THEOREM 1. Depending on the Assumptions 1-3 and Conditions 1-2, if $\epsilon < \frac{C\gamma\delta}{4\gamma+2C\delta H(Y)}$ and $\lambda \in [\frac{H(Y)+1/2\delta C}{\delta C-4\epsilon} - \frac{1}{2}, \frac{\gamma}{4\epsilon} - \frac{1}{2}]$, then we will have $\hat{L}(\Phi_v) < \hat{L}(\Phi)$ for all $\Phi \neq \Phi_v$, where $H(Y)$ denotes the entropy of Y .

A.3 More Experimental Results

We provide more experimental results to further show the effectiveness of the proposed HEI in addressing heterophilic graph structure distribution shifts.

RQ2: Additional Experiments on Simulation Settings using GloGNN++ as backbone. We also conduct experiments under severe distribution shifts using GloGNN++ as the backbone. As shown in Figure 6, our proposed method can acquire superior or comparable results than previous methods to handle graph structure distribution shifts, which further verifies the effectiveness and robustness of our design.

Table 4: Statistics for our used heterophilic graph datasets.

Dataset	Chameleon	Squirrel	Actor	Penn94	arXiv-year	twitch-gamer
Nodes	2277	5201	7600	41554	169343	168114
Edges	36101	216933	29926	1362229	1166243	6797557
Feat	2325	2089	931	5	128	7
Class	5	5	5	2	5	2
Edge hom.	0.23	0.22	0.22	0.47	0.222	0.545

RQ3: Effect of different similarity matrices as neighbor pattern indicators for HEI. We provide large-scale graph experiments as shown in Table 7 to clarify the details of HEI.

RQ4: Sensitive analysis. We provide the experimental results about RQ4 there as shown in Figure 7.

RQ5: Efficiency Studies. As shown in Table 5, referring to [18], we provide the time(seconds) to train the model until it converges which keeps the stable accuracy score on the validation set. From the results, we can conclude that the extra time cost can be acceptable compared with the backbone itself.

Experiments on Homophilic Graph Datasets. Considering the fact that in real-world settings, we can't know whether the input graph is homophilic or heterophilic in advance. Thus, we also provide comparison experiments and discussions for homophilic graphs. As shown in Table 6, from the results, we observe that our method can achieve consistent comparable performance to other baselines. But exactly, the improvements by these methods are all minor compared with the results of ERM. That's because the homophilic graph is not related to our settings. After all, homophilic graph datasets mean the neighbor pattern distribution between the train and test are nearly the same, which is not suitable to clarify our defined distribution shifts. The performance gap between the low home test and the high home test can support our analysis.

Table 5: Efficiency studies of HEI, where the report scores show the time (seconds) to train the model until converge that keeps the stable accuracy score on the validation set. Referring to [18], we adopt the GloGNN++ as the backbone there.

Methods	Penn94	arxiv-year	twitch-gamer
ERM	22.3	7.2	40.5
Renode	23.5	8.5	41.2
SRGNN	24.9	9.1	41.0
EERM	24.7	8.8	41.5
BAGNN	24.8	9.1	42.1
FLOOD	24.5	8.8	41.8
StruRW	23.8	9.8	41.5
GDN	24.7	9.6	42.2
CaNet	25.9	10.8	42.2
IENE	25.4	10.6	42.8
HEI(Ours)	26.8	11.5	44.9

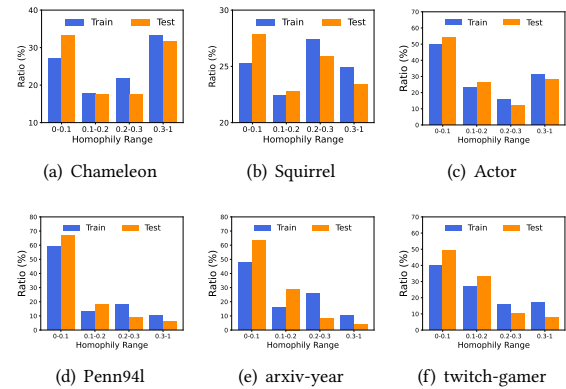


Figure 5: Statistic of homophily ratio for train and test nodes following previous dataset splits. The nodes are categorized into four groups according to the node-level homophily. Compared with test nodes, the train nodes are more prone to be categorized into groups with high homophily. In other words, in the range with high homophily (from 0.2 to 0.3 and from 0.3 to 1), the sub-train ratio in all train nodes is higher than the sub-test in all test nodes. But in the range with low homophily, there exists a contrary phenomenon.

A.4 Implementation Details

We provide detailed implementation details for our experiments.

ERM, It corresponds to the results of the backbone itself [18, 19] under our constructed settings.

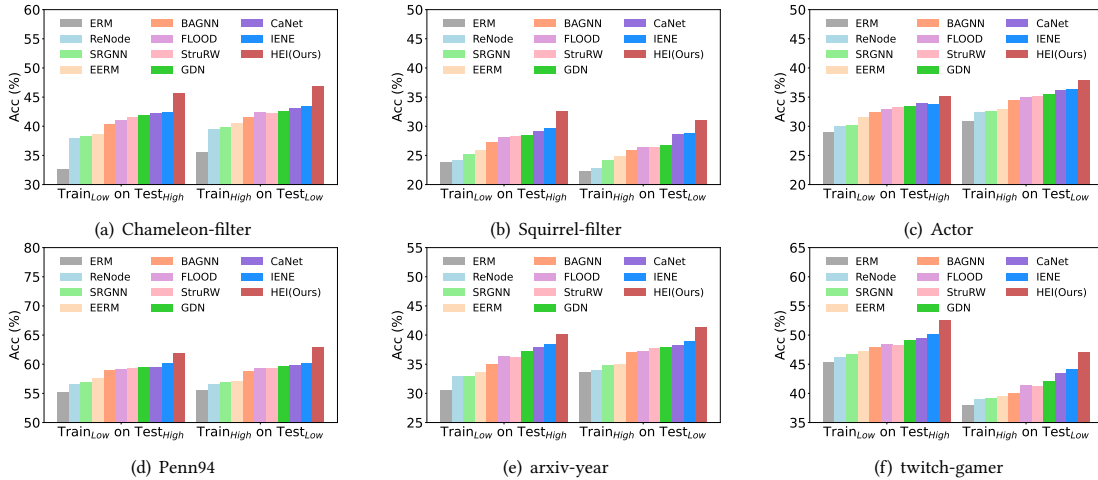
SRGNN[43], Shift-Robust GNN is a framework inspired by domain adaptation, which means it needs prior knowledge from the target domain. Specifically, it strives to adapt a biased sample of labeled nodes to more closely conform to the distributional characteristics present in an IID sample of the graph. In our experiments, we utilize the estimated neighbor distribution information to evaluate the distribution of the source domain and target domain. Apart from this, we entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

Renode[4], in our experiments, it is an extension of the original Renode, which is a model-agnostic training weight schedule mechanism to cope with topology-imbalance problems for semi-supervised node classification. Specifically, they devise a cosine annealing mechanism for the training node weights based on their Totoro values. The Totoro values involve topology information and can also integrate with quantity-imbalanced methods as the paper shown. Therefore, this method can also be seen as a baseline which can copy with agnostic distribution shifts. We adopt the same reweight strategy as the paper stated, and the Totoro values can also describe the neighbor pattern to a certain extent.

EERM[37], Explore-to-Extrapolate Risk Minimization (EERM) is an invariant learning approach that facilitates graph neural networks to leverage invariance principles for prediction. As stated in the paper, EERM resorts to multiple context explorers that are adversarially trained to construct diverse environments with augmentation. Then, following the principle of variance of risks, it

Table 6: Performance comparison on homophilic graph datasets under Standard Settings. The reported scores denote the classification accuracy (%) and error bar (\pm) over 10 trials. We highlight the best score on each dataset in bold and the second score with underline.

Backbones	Methods	CiteSeer			PubMed			Cora		
		Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test
LINKX	ERM	73.19 \pm 0.99	73.89 \pm 1.51	72.79 \pm 1.47	87.86 \pm 0.77	88.49 \pm 1.37	87.19 \pm 1.51	84.64 \pm 1.13	85.13 \pm 1.83	83.85 \pm 1.99
	ReNode	73.25 \pm 0.89	74.00 \pm 1.21	72.89 \pm 1.87	87.91 \pm 0.72	88.58 \pm 1.42	87.21 \pm 1.71	84.70 \pm 1.23	85.28 \pm 1.99	84.24 \pm 2.13
	SRGNN	73.27 \pm 0.99	74.03 \pm 1.11	72.85 \pm 1.87	87.96 \pm 0.81	88.68 \pm 1.72	87.31 \pm 1.51	84.71 \pm 1.25	85.22 \pm 1.87	84.34 \pm 2.03
	StruRW-Mixup	73.29 \pm 0.91	73.93 \pm 1.25	72.99 \pm 1.91	88.12 \pm 0.51	88.71 \pm 1.44	87.58 \pm 1.59	84.67 \pm 1.54	85.33 \pm 1.91	84.34 \pm 2.43
	EERM	73.17 \pm 0.79	73.81 \pm 1.45	73.09 \pm 1.65	87.96 \pm 0.84	88.59 \pm 1.32	87.29 \pm 1.63	84.62 \pm 1.37	85.14 \pm 1.63	83.87 \pm 2.03
	BAGNN	73.33 \pm 0.88	73.99 \pm 1.61	73.19 \pm 1.63	88.01 \pm 0.94	88.78 \pm 1.57	87.69 \pm 1.39	84.60 \pm 1.28	85.24 \pm 1.83	83.84 \pm 2.41
	FLOOD	73.34 \pm 0.91	73.95 \pm 1.55	73.22 \pm 1.67	88.05 \pm 0.95	88.84 \pm 1.62	87.81 \pm 1.59	84.72 \pm 1.41	85.35 \pm 1.63	83.99 \pm 2.51
	CaNet	73.38 \pm 0.95	74.08 \pm 1.54	73.31 \pm 1.55	88.11 \pm 0.98	88.89 \pm 1.67	87.91 \pm 1.54	84.81 \pm 1.31	85.39 \pm 1.57	84.11 \pm 2.58
	IENE	73.43 \pm 0.97	74.15 \pm 1.58	73.32 \pm 1.87	88.12 \pm 0.94	88.90 \pm 1.65	87.90 \pm 1.66	84.92 \pm 1.45	85.41 \pm 1.68	84.45 \pm 2.81
	GDN	73.31 \pm 0.81	73.99 \pm 1.61	73.21 \pm 1.68	88.14 \pm 0.94	88.91 \pm 1.64	87.92 \pm 1.41	84.64 \pm 1.33	85.27 \pm 1.69	83.91 \pm 2.78
	HEI(Ours)	73.51 \pm 0.81	74.18 \pm 1.25	73.42 \pm 1.85	88.50 \pm 0.97	89.01 \pm 1.24	87.99 \pm 1.92	85.17 \pm 1.53	85.44 \pm 1.83	84.84 \pm 1.97
GloGNN++	ERM	77.22 \pm 1.78	78.15 \pm 2.55	76.79 \pm 2.54	89.24 \pm 0.39	90.62 \pm 0.99	88.75 \pm 1.28	88.33 \pm 1.09	90.06 \pm 1.52	87.37 \pm 1.64
	ReNode	77.31 \pm 1.69	78.27 \pm 2.48	76.90 \pm 2.39	89.25 \pm 0.35	90.64 \pm 0.87	88.79 \pm 1.24	88.39 \pm 1.21	90.11 \pm 1.49	87.45 \pm 1.57
	SRGNN	77.33 \pm 1.65	78.24 \pm 2.75	76.91 \pm 2.77	89.33 \pm 0.51	90.81 \pm 1.21	88.99 \pm 1.57	88.53 \pm 1.09	90.46 \pm 1.53	87.58 \pm 1.54
	StruRW-Mixup	77.35 \pm 1.57	78.27 \pm 2.11	76.90 \pm 2.42	89.48 \pm 0.44	90.81 \pm 0.97	89.17 \pm 1.33	88.39 \pm 1.44	90.15 \pm 1.69	87.85 \pm 1.77
	EERM	77.35 \pm 1.81	78.27 \pm 2.45	76.89 \pm 2.81	89.34 \pm 0.39	90.82 \pm 1.09	88.95 \pm 1.38	88.39 \pm 1.21	90.36 \pm 1.42	87.47 \pm 1.74
	BAGNN	77.42 \pm 1.81	78.35 \pm 2.81	76.89 \pm 2.42	89.37 \pm 0.45	90.87 \pm 1.29	88.99 \pm 1.58	88.49 \pm 1.31	90.39 \pm 1.47	87.81 \pm 1.64
	FLOOD	77.43 \pm 1.79	78.39 \pm 2.51	76.95 \pm 2.37	89.41 \pm 0.51	90.91 \pm 1.29	89.08 \pm 1.58	88.51 \pm 1.27	90.40 \pm 1.51	87.88 \pm 1.69
	CaNet	77.42 \pm 1.81	78.32 \pm 2.54	76.75 \pm 2.87	89.51 \pm 0.58	90.97 \pm 1.29	89.19 \pm 1.61	88.54 \pm 1.33	90.51 \pm 1.61	88.11 \pm 1.58
	IENE	77.43 \pm 1.79	78.39 \pm 2.51	76.96 \pm 2.58	89.52 \pm 0.57	91.00 \pm 1.44	89.22 \pm 1.62	88.78 \pm 1.37	90.79 \pm 1.47	88.28 \pm 1.49
	GDN	77.44 \pm 1.51	78.65 \pm 3.75	76.97 \pm 2.53	89.39 \pm 0.51	90.97 \pm 1.01	88.99 \pm 1.21	88.29 \pm 1.34	90.44 \pm 1.91	87.89 \pm 1.22
	HEI(Ours)	77.85 \pm 1.89	79.11 \pm 2.59	77.30 \pm 2.85	89.99 \pm 0.39	91.52 \pm 0.99	89.48 \pm 1.33	88.93 \pm 1.19	90.97 \pm 1.39	88.47 \pm 1.74

**Figure 6: Comparison experiments under Simulation Settings where exists severe distribution shift between train and test, including $Train_{High}$ on $Test_{Low}$ and $Train_{Low}$ on $Test_{High}$. We adopt the GloGNN++ as the backbone there.****Table 7: Comparison experiments on large-scale datasets when we respectively adopt local similarity(Local Sim), post-aggregation similarity (Agg Sim), and SimRank as indicators to estimate nodes' neighbor patterns so as to infer latent environments under Standard settings.**

Backbones	Methods	Penn94			arxiv-year			twitch-gamer		
		Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test	Full Test	High Hom Test	Low Hom Test
LINKX	HEI (Local Sim)	85.12 \pm 0.21	88.28 \pm 0.33	82.15 \pm 0.59	54.41 \pm 0.21	64.23 \pm 0.47	48.29 \pm 0.22	66.18 \pm 0.12	83.75 \pm 0.34	48.12 \pm 0.47
	HEI (Agg Sim)	85.21 \pm 0.17	88.29 \pm 0.38	82.22 \pm 0.54	54.45 \pm 0.23	64.33 \pm 0.49	48.33 \pm 0.32	66.21 \pm 0.15	83.85 \pm 0.39	48.45 \pm 0.57
	HEI (SimRank)	86.22 \pm 0.28	89.24 \pm 0.28	83.22 \pm 0.59	56.05 \pm 0.22	66.53 \pm 0.41	49.33 \pm 0.32	66.79 \pm 0.14	85.33 \pm 0.25	49.21 \pm 0.57
GloGNN++	HEI (Local Sim)	86.08 \pm 0.24	89.70 \pm 0.64	82.18 \pm 0.37	54.42 \pm 0.24	64.48 \pm 1.54	48.55 \pm 0.64	66.30 \pm 0.18	83.21 \pm 0.68	49.00 \pm 0.67
	HEI (Agg Sim)	86.15 \pm 0.25	89.70 \pm 0.69	82.43 \pm 0.38	54.44 \pm 0.25	64.51 \pm 1.54	48.69 \pm 0.81	66.34 \pm 0.21	83.19 \pm 0.78	49.14 \pm 0.57
	HEI (SimRank)	87.18 \pm 0.28	89.99 \pm 0.65	83.59 \pm 0.39	55.71 \pm 0.24	66.29 \pm 1.14	49.52 \pm 0.75	66.99 \pm 0.17	84.37 \pm 0.68	50.40 \pm 0.52




Figure 7: Parameter Sensitivity of environmental numbers k under Standard Settings.

utilizes the variance from multiple virtual environments as regularization to help model training. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

BAGNN[9], Bias-aware(BA) GNN is also a method that deals with agnostic distribution shifts on graphs from the perspective of invariant learning. Specifically, it can be summarized into two steps: the environment clustering module assigns nodes to environments with minimization loss, and the invariant graph learning module learns invariant representation across environments with minimization loss. During the process of environment clustering, it adopts the masking strategy with graph augmentation. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

FLOOD[24], it is also a flexible invariant learning framework for OOD generalization on graphs. Specifically, it includes two key modules, invariant learning, and bootstrapped learning. The invariant learning modules construct multiple environments from graph data augmentation and learn invariant representation under risk extrapolation. Besides, the bootstrapped learning component is inspired by test time adaptation, which proposes to train a shared graph encoder with the invariant learning part according to the test distribution. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

CaNet[36], it is a recently proposed invariant learning framework that integrates an environment estimator with a mixture-of-expert GNN predictor. , aiming to train robust GNNs under node-level distribution shifts. Exactly, it holds the findings that the crux of GNNs' failure in OOD generalization lies in the latent confounding bias from the environment and proposes to estimate the pseudo environments for each layer of the GNN network, assisted causal inference. Their defined environments are different from the environment we clarify in the paper, which is just stated from the perspective of feature separation. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

IENE[39], it is also a recently proposed invariant learning framework that identifies and extrapolates the node environment for Out-of-Distribution Generalization on graphs. However, for extrapolating topological environments, they still adopt graph augmentation techniques to identify structural invariance, which is indeed different from our strategy for inferring environments. We entirely

follow this work to address graph structure distribution shifts on heterophilic graphs.

StruRW[23], it is a structure-reweighting method originally designed for a new type of conditional structure shift (CSS), which the current Graph domain adaptation approaches are provably sub-optimal to deal with. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

GDN[36], it is a prototype learning method originally designed for Graph Anomaly Detection. It teases out the anomaly features and mitigates the effect of heterophilic neighbors by devising a dynamically optimized prototype vector to guide the node representation learning under graph structure distribution shift. We entirely follow this work to address graph structure distribution shifts on heterophilic graphs.

HEI (Ours), our training process can be concluded as follows: Given a heterophilic graph input, we first calculate the SimRank for each node in advance. Then, based on Eq. 9, we collectively learn environment partition and invariant representation on heterophilic graphs, assisted by SimRank, to address graph structure distribution shifts on heterophilic graphs. Therefore, we save the processed SimRank values on nodes on the graph in advance and transfer them into tensors for the training. For training details, we should warm up for some epochs to avoid the learned environments in the initial stage that are not effective, which may influence the optimization of models. So at the beginning warm-up stage, we adopt the ERM strategy. After that, we adopt our proposed framework to learn an invariance penalty to improve model performance. For the range of parameters, we first execute experiments using basic backbones to get the best parameters of num-layers and hidden channels on different datasets. Then, we fix the num-layers and hidden channels to adjust other parameters, penalty weight λ) from $\{1e-, 1e-2, 1e-1, 1, 10, 100\}$, learning rate from $\{1e-2, 5e-3, 1e-3, 5e-4, 1e-4\}$ and weight decay from $\{1e-2, 5e-3, 1e-3\}$. We also provide parameter sensitivity of environment number k in the paper. Moreover, the ρ is a two-layer MLP with the hidden channel from $\{16, 32, 64\}$, and its learning rate should be lower than the backbone in our experiments, within the range from $\{5e-3, 1e-3, 5e-4, 1e-4\}$.