

Enhancing Distribution and Label Consistency for Graph Out-of-Distribution Generalization

Song Wang[†], Xiaodong Yang[‡], Rashidul Islam[‡], Huiyuan Chen[‡], Minghua Xu[‡], Jundong Li[†], Yiwei Cai[‡]

[†]*University of Virginia, Charlottesville, USA, {sw3wv, jundong}@virginia.edu*

[‡]*Visa Research, USA, {xiaodyan, raislam, hchen, mixu, yicai}@visa.com*

Abstract—To deal with distribution shifts in graph data, various graph out-of-distribution (OOD) generalization techniques have been recently proposed. These methods often employ a two-step strategy that first creates augmented environments and subsequently identifies invariant subgraphs to improve generalizability. Nevertheless, this approach could be suboptimal from the perspective of *consistency*. First, the process of augmenting environments by altering the graphs while preserving labels may lead to graphs that are not realistic or meaningfully related to the origin distribution, thus lacking *distribution consistency*. Second, the extracted subgraphs are obtained from directly modifying graphs, and may not necessarily maintain a consistent predictive relationship with their labels, thereby impacting *label consistency*. In response to these challenges, we introduce an innovative approach that aims to enhance these two types of consistency for graph OOD generalization. We propose a modifier to obtain both augmented and invariant graphs in a unified manner. With the augmented graphs, we enrich the training data without compromising the integrity of label-graph relationships. The label consistency enhancement in our framework further preserves the supervision information in the invariant graph. We conduct extensive experiments on real-world datasets to demonstrate the superiority of our framework over other state-of-the-art baselines.

Index Terms—Graph Neural Networks, Distribution Shifts, Out-of-Distribution (OOD) Generalization

I. INTRODUCTION

Graph data is prevalent in various applications, such as molecular property predictions and financial analysis [1], [2], [3], [4], [5], [6]. Recently, techniques in graph machine learning have significantly advanced. For example, Graph Neural Networks (GNNs) have been developed to effectively handle graph data [7], [8], [9]. GNNs learn representations of individual nodes by aggregating information from their neighboring nodes [10], [11], [12], [13]. Despite the notable success of graph machine learning, a significant challenge remains. GNNs typically assume that the training data follows the same distribution as the test data. However, this assumption often does not hold true in real-world scenarios [14], [15], [16], [17]. Consequently, when the test data distribution differs significantly, the performance of GNNs trained on the training data deteriorates. This issue is commonly referred to as distribution shift [18], [19], [20].

Although numerous solutions to distribution shifts have been proposed for Euclidean data (e.g., images) [21], [22], [23], [24], [25], [26], these techniques are challenging to apply to graphs due to their complex node connections [27], [28], [29]. To address distribution shifts in graphs, recent

methods have proposed generating augmented environments through data augmentation [30], [31]. By training on this augmented data, GNNs can learn the decisive information for classification while ignoring the spurious information of each environment, thereby achieving the robustness necessary to perform well on test data [32], [33].

However, such a strategy can be suboptimal, facing two crucial challenges: **1 Distribution Consistency**: When generating augmented environments, the graphs are obtained via modification, which might diverge significantly from real-world examples. For example, in molecular property prediction, minor alterations to molecular structures can result in significant changes in chemical properties, thus disrupting the distribution consistency of augmented graphs [34]. This inconsistency arises because the augmented graphs may not accurately reflect realistic or probable molecular configurations.

2 Label Consistency: The process of extracting invariant subgraphs, without careful consideration of how these changes impact label relevance, risks impairing the learned relationship between graph structures and labels. This is because the identified invariant subgraphs may not accurately reflect the labels they should belong to, which diminishes the preserved supervision information.

To deal with these challenges, we propose a novel framework named **DLG**, which aims to enhance **D**istribution- and **L**abel-consistency for **G**raph OOD generalization. Specifically, we design a novel graph modifier to produce both augmented and invariant graphs by sampling edges from the learned edge masks. Using such a strategy enables us to unify the processes of augmenting existing graphs and extracting invariant graphs by formulating them both as modifications, in order to enhance consistency from perspectives of distributions and labels. **1 Distribution Consistency Enhancement**. To obtain augmented graphs during training while preserving distribution consistency, we propose to maximize the information between augmented graphs and existing graphs. In this manner, we could ensure the augmented graphs are informative and align with existing graphs to preserve distribution consistency. **2 Label Consistency Enhancement**. To ensure the generated invariant graphs remain consistent across different domains regarding their labels, we explicitly enhance label consistency by ensuring that the extracted invariant subgraphs share maximal supervision information with the original graph G . In summary, our contributions are as follows:

* **Investigation.** We investigate the challenges of invariant

learning in the problem of OOD generalization on graphs, with a particular emphasis on *Distribution Consistency* and *Label Consistency*.

- ★ **Design.** We develop a novel OOD generalization framework with two essential strategies: (1) a distribution consistency enhancement module that aims to achieve augmented graphs that exhibit better cross-domain consistency; (2) a label consistency enhancement module that ensures that the extracted invariant subgraph contains maximal decisive information for classification;
- ★ **Effectiveness.** We conduct extensive experiments on a variety of both graph-level and node-level graph OOD generalization datasets that cover both synthetic and real-world graph data. The results further demonstrate the superiority of our proposed framework.

II. RELATED WORKS

The challenge of out-of-distribution (OOD) generalization is to enable models to generalize to previously unseen test distributions after training on similar but distinct data [35], [36]. Recent research has focused on invariant learning [37], [38], which aims to maintain a stable relationship between inputs and outputs across different distributions. Current approaches typically achieve this by learning invariant representations [39] or invariant causal predictions [40], [41]. Additionally, some methods seek OOD generalization by optimizing the worst-case group performance without predefined partitions in the training data [42], [43], [44].

The issue of OOD generalization in graph data has gained increased attention recently [45], [46], driven by the prevalence of distribution shifts in various real-world scenarios [47]. Among these approaches, EERM [18] was proposed for node classification and uses invariance principles to generate domains by maximizing loss variance across domains in an adversarial manner. For graph classification tasks, DIR [34] employs graph representations as causal rationales and generates additional distributions via interventional augmentations. CIGA [48] and DisC [49] both propose to capture the causal information in subgraphs. To address the limitations of CIGA and DisC in scenarios where the prior knowledge about the variance of causal and spurious information is unavailable, GALA [50] uses an assistant model to detect distribution shifts based on this variance.

III. PROBLEM FORMULATION

We first denote the set of training graphs as $\mathcal{D}_{tr} = \{(G_i, Y_i) | i \in [1, |\mathcal{D}_{tr}|]\}$. Here $G \in \mathcal{G}$ and $Y \in \mathcal{Y}$ are graph samples and their labels. We further denote each graph as $G = (\mathbf{A}, \mathbf{X})$, where $\mathbf{A} \in \mathbb{R}^{N \times N}$ and $\mathbf{X} \in \mathbb{R}^{N \times d_x}$ represents the adjacency matrix and feature matrix, respectively. Here N is the number of nodes in \mathcal{G} , and the j -th row vector of \mathbf{X} is the d_x -dimensional attributes of the j -th node. Following existing works [48], [47], [34], the goal of OOD graph generalization is to learn a classification function $f_c(\cdot) : \mathcal{G} \rightarrow \mathcal{Y}$ that takes a graph as input and predicts its label Y . The ultimate goal of OOD graph generalization is to learn the optimal classification

function f^* that could precisely predict the labels of graphs in the test set, i.e., $\mathcal{D}_{te} = \{(G_i^*, Y_i^*) | i \in [1, |\mathcal{D}_{te}|]\}$

IV. METHODOLOGY

We elaborate on our framework DLG for graph OOD generalization problem via enhancing distribution and label consistency. As illustrated in Fig. 1, our framework leverages a classifier $f_c(\cdot)$ and two graph modifiers $g_a(\cdot)$ and $g_v(\cdot)$, which generate augmented and invariant graphs, respectively. The encoder and modifiers are optimized based on our designed objectives for the two types of consistency. Note that our modifier only considers the edges in each graph, i.e., only modifying the adjacency matrix. That being said, we could potentially create edges that do not exist in the original graph.

A. Modification

Considering an input graph $G = (\mathbf{A}, \mathbf{X})$ from the training dataset \mathcal{D}_{tr} , we first aim to learn its invariant graph $G_v = (\mathbf{A}_v, \mathbf{X}_v)$ that preserves label consistency for classification. Additionally, during training, we also generate an augmented graph $G_a = (\mathbf{A}_a, \mathbf{X}_a)$ that preserves distribution consistency. The augmented graph G_a is also classified by the classifier $f(\cdot)$ with the same label as G .

In the following, we describe the process of obtaining the invariant graph G_v and the augmented graph G_a in detail. Particularly, we first learn the representations for modification:

$$\mathbf{H}_v = \text{GNN}_v(\mathbf{A}, \mathbf{X}) \quad \text{and} \quad \mathbf{H}_a = \text{GNN}_a(\mathbf{A}, \mathbf{X} \oplus \mathbf{y}). \quad (1)$$

Here GNN_v and GNN_a are two GNNs used in $g_v(\cdot)$ and $g_a(\cdot)$, respectively. \mathbf{H}_v (or \mathbf{H}_a) represents the learned embedding matrix of G_v (or G_a). Notably, for the input to GNN_a , we additionally incorporate the information from label Y of G . This is because the augmented graphs are only used during training, which means its label information is always available. Moreover, with label information added, it also benefits our label consistency enhancement when used for training, as introduced in the following section. During the inference stage, the label information is unavailable, and thus we do not incorporate label Y in $g_v(\cdot)$.

With the learned representations, we illustrate the detailed process of modification. As the process applies to both invariant and augmented graphs, here we use \mathbf{H} for simplicity. Particularly, given \mathbf{H} , our target is to learn from edge masks $\{e_{i,j} | i, j = 1, 2, \dots, N\}$. Here each mask is in $[0, 1]$, denoting the sampling probability of a specific edge. In general, a larger mask leads to a higher sampling probability, indicating that the edge is more likely to appear in the invariant graphs or augmented graphs.

To learn the edge masks, we apply a Multi-Layer Perceptron (MLP) and calculate the dot product:

$$e_{i,j} = \text{Sigmoid} \left((\text{MLP}_e(\mathbf{h}_i))^T \cdot \text{MLP}_e(\mathbf{h}_j) \right), \quad (2)$$

where \mathbf{h}_v denotes the representation of the i -th node, i.e., the i -th row vector in \mathbf{H} . With the learned masks, we could perform sampling for edges. Nevertheless, such a sampling

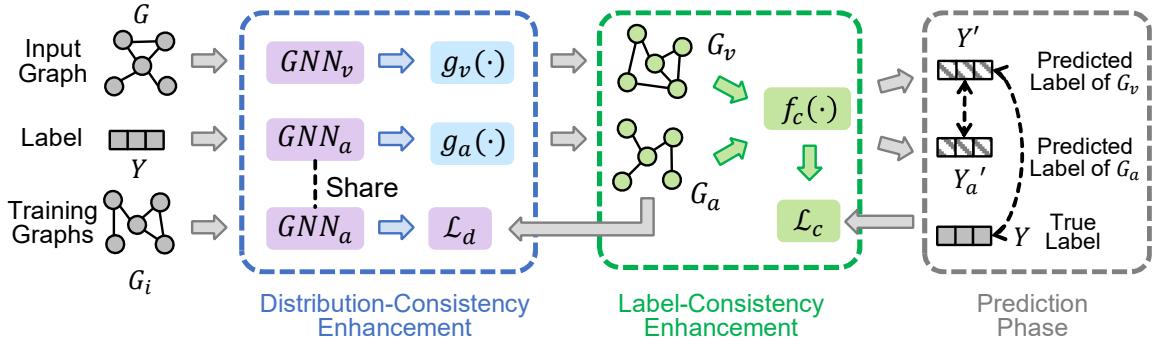


Fig. 1: The overall framework of DLG. Given a graph $G = (\mathbf{A}, \mathbf{X})$, we feed it into GNN_v to learn a graph representation, which will be used to generate an invariant graph G_v for classification. Meanwhile, the input graph is also processed by GNN_a to generate an augmented graph G_a , which shares the same label as G_v but with different structures and edges. To enhance label consistency, we propose the classification loss \mathcal{L}_c for optimization. To enhance distribution consistency, we design a loss \mathcal{L}_d that leverages randomly sampled graphs from existing data to ensure G_a aligns with distributions of existing data.

process is non-differentiable. To enable the optimization of our framework via gradient descent, we utilize the sigmoid variant of gumbel-softmax [51] to sample edges. By factorizing the modification process into sampling individual edges, we could avoid the extensive computational cost caused by the large graph space [52].

B. Distribution Consistency Enhancement

Although we generate augmented graphs to expand the training data distribution, the strategy cannot guarantee that the generated graphs preserve consistency with existing graphs. In other words, the generated graphs may largely deviate from the distribution of existing data. Moreover, the labels of augmented graphs may also be different from the graph before augmentation.

To deal with this, we propose to maximize the mutual information between augmented graphs G_a and other existing graphs with the same label. Specifically, we first randomly sample a batch of graphs from all training data as the support set $\mathcal{S} = \{G_i | i = 1, 2, \dots, C\}$, containing one graph for each class. Here C is the number of classes. Moreover, the label of G_i is y_i , i.e., the i -th label in the label space \mathcal{Y} . In this manner, the objective of maximizing the mutual information between G_a and \mathcal{S} can be represented as follows:

$$\max I(G_a; \mathcal{S}) = \max \sum_{i=1}^C p(G_a, G_i) \log \frac{p(G_a, G_i)}{p(G_a)p(G_i)}. \quad (3)$$

Since the mutual information $I(G_a; \mathcal{S})$ is intractable and thus infeasible to maximize [53], we transform the objective into a more accessible form:

$$I(G_a; \mathcal{S}) = \sum_{i=1}^C p(G_a|G_i)p(G_i) \log \frac{p(G_i|G_a)}{p(G_i)}. \quad (4)$$

This form is easier to optimize as it contains the term $p(G_i)$. As we randomly sample G_i from each class, we assume that the prior probability, i.e., $p(G_i)$, follows a uniform distribution

and set it $p(G_i) = 1/C$. According to Bayes' theorem, the objective then becomes:

$$I(G_a; \mathcal{S}) = \frac{1}{C} \sum_{i=1}^C p(G_a|G_i) \left(\log p(G_i|G_a) - \log \frac{1}{C} \right). \quad (5)$$

To achieve the value of $p(G_a|G_i)$, we estimate it as $p(G_a|G_i) = \mathbb{1}(Y = y_i)$, where Y and y_i are the labels of G and G_i , respectively. Here $\mathbb{1}(Y = y_i) = 1$ if G and G_i share the same label; otherwise $\mathbb{1}(Y = y_i) = 0$. In this manner, considering that $\log(1/C)$ is a constant, the above objective can be further simplified as follows:

$$\max I(G_a; \mathcal{S}) = \max \frac{1}{C} \sum_{i=1}^C \mathbb{1}(Y = y_i) \log p(G_i|G_a). \quad (6)$$

To further estimate $\log p(G_i|G_a)$, we can define the probability of G_i sharing the same label as G (i.e., $y_i = Y$) according to the squared ℓ_2 norm of the embedding distance. Specifically, we first apply the same GNN_a on G_i to obtain its representation:

$$\mathbf{h}_i = \frac{1}{|\mathcal{V}_i|} \sum_{j=1}^{|\mathcal{V}_i|} \mathbf{H}_i(j), \text{ where } \mathbf{H}_i = GNN_a(\mathbf{A}_i, \mathbf{X}_i \oplus \mathbf{y}_i), \quad (7)$$

where \mathbf{A}_i and \mathbf{X}_i are adjacency matrix and feature matrix of G_i , respectively. \mathbf{y}_i is the corresponding one-hot vector of the label of G_i . Based on the representation, the term $p(G_i|G_a)$ can be formulated as below, after normalization with a softmax function:

$$p(G_i|G_a) = \frac{\exp(-(\mathbf{h}_a - \mathbf{h}_i)^2)}{\sum_{j=1}^C \exp(-(\mathbf{h}_a - \mathbf{h}_j)^2)}, \quad (8)$$

where \mathbf{h}_a and \mathbf{h}_i are graph-level representations of G and G_i , respectively. Note that as G_a is generated from G , here we use the graph representation of G as \mathbf{h}_a to obtain the objective. If we further apply the ℓ_2 normalization to both \mathbf{h}_a and \mathbf{h}_i , we obtain $(\mathbf{h}_a - \mathbf{h}_i)^2 = 2 - 2\mathbf{h}_a \cdot \mathbf{h}_i$. Additionally, to enhance diversity, we still need to ensure the obtained G_a is

not excessively similar to G_v , even though they are of the same label. Thus, we modify the denominator in Eq. (8) to involve the graph-level representation of G_v , i.e., \mathbf{h}_v . The objective becomes:

$$\mathcal{L}_d = - \sum_{i=1}^C \mathbb{1}(Y = y_i) 2\mathbf{h}_a \cdot \mathbf{h}_i - \log(\exp(2\mathbf{h}_a \cdot \mathbf{h}_v) + \sum_{j=1}^C \exp(2\mathbf{h}_a \cdot \mathbf{h}_j)). \quad (9)$$

With this objective, we aim to ensure that the representation \mathbf{H}_a used for generating G_a should be similar to existing graphs with the same label while distinct from other graphs with different labels, and more importantly, the original graph G . Consequently, we are able to broaden the distribution of the class to which G belongs, through the use of the generated G_a , while also preserving the informativeness of G_a .

C. Label Consistency Enhancement

As our goal for the extracted invariant graph is to ensure label consistency, we strive to preserve as much useful information from the original graph G as possible. Therefore, we formulate our objective as maximizing the conditional likelihood of $P(G_v|G, Y)$, representing the probability of the generated consistent graph G_v conditioned given the original graph G . That being said, we aim to ensure that the generated G_v share maximal supervision information with G while ignoring unrelated information. Therefore, we obtain the classification losses: $\mathbf{q}_v = f(G_v)$, $\mathbf{q}_a = f(G_a)$, and

$$\mathcal{L}_c = - \sum_{i=1}^C \left(\mathbb{1}(Y = y_i) \log(q_v(i)) + q_v(i) \log(q_a(i)) \right), \quad (10)$$

where $\mathbf{q}_v \in \mathbb{R}^C$ (or $\mathbf{q}_a \in \mathbb{R}^C$) is the probability that G_v (or G_a) belongs to each class in the label space \mathbb{Y} , and C is the number of classes. Moreover, $q_v(i)$ (or $q_a(i)$) is the i -th element in \mathbf{q}_v (or \mathbf{q}_a). Here $\mathbb{1}(Y = y_i) = 1$ if the label Y of G is y_i ; otherwise $\mathbb{1}(Y = y_i) = 0$. In this way, we enforce the class probability distributions of G_v and G_a to be close to each other while utilizing label information of G .

D. Optimization

In this subsection, we introduce the detailed optimization steps of DLG. As described in Sec. IV-C and Sec. IV-B, we propose two losses to optimize our framework, i.e., \mathcal{L}_d and \mathcal{L}_c . However, directly applying all two losses can be detrimental to the overall performance, as the modules in our framework maintain different objectives. Let us denote the parameters of GNN_v and GNN_a as θ_v and θ_a , respectively. We apply both losses for optimization:

$$\{\theta_v^*, \theta_a^*\} = \arg \min_{\{\theta_v, \theta_a\}} \alpha \mathcal{L}_d + (1 - \alpha) \mathcal{L}_c, \quad (11)$$

where θ_v^* and θ_a^* are the optimal parameters for GNN_v and GNN_a , respectively. $\alpha \in [0, 1]$ is a hyper-parameter that controls the importance of \mathcal{L}_d .

Moreover, as the classifier is only involved in the classification loss, we optimize it with \mathcal{L}_c :

$$\theta_f^* = \arg \min_{\theta_f} \mathcal{L}_c, \quad (12)$$

where θ_f^* denotes the optimal parameters for the classifier $f_c(\cdot)$, respectively. In this manner, we can jointly optimize all modules in our framework in an end-to-end manner.

V. EXPERIMENTS

In this section, we evaluate our framework DLG on a variety of synthetic and real-world graph datasets for out-of-distribution generalization.

A. Datasets and Settings

Graph-Level OOD Datasets. For experiments, we utilize a combination of synthetic and real datasets for OOD generalization on graph classification tasks. Specifically, we first consider graph-level datasets used in GALA [50]. The included datasets are as follows: ① TPG (Two-Piece Graph) datasets constructed based on the BA-2motifs [54], resulting in four variants of 3-class two-piece graph datasets. Each variant is controlled by a hyperparameter that adjusts the correlations between the label and the spurious information in the graph, ranging from $+0.2$ to -0.2 . ② Six Datasets from the DrugOOD benchmark [55], which is centered on the challenging real-world task of AI-aided drug affinity prediction. The DrugOOD datasets include splits based on Assay, Scaffold, and Size from two categories: EC50 (denoted as EC50-) and Ki (denoted as Ki-). These datasets provide a diverse range of scenarios to evaluate the effectiveness of various methods in predicting drug affinities across different biological and chemical contexts. We further consider four graph-level datasets used in DIR [34]: ③ **SP-Motif** [56] is a synthetic dataset, in which the degree of bias can be manually controlled. Each graph consists of a base (with structures from Tree, Ladder, and Wheel) and a motif (with structures from Cycle, House, and Crane). The label is entirely decided by the structure of the motif in each graph. We consider the average results of four bias degrees: $1/3$, 0.5 , 0.7 , and 0.9 . ④ **MNIST-75sp** [57] is converted from the MNIST image dataset, and random noises are added to create distribution shifts. There exist 10 labels in the dataset. Moreover, the node features contain random noises in the test set. ⑤ **Molhiv** (OGBG-Molhiv) [58], [59], [60] is a molecular property prediction dataset, and the distribution shift is based on different structures of molecules. ⑥ **Graph-SST2** contains graphs constructed from sentimental sentences, and the distribution shift originates from different partitions of graphs. Note that due to the different properties of graph structures in these datasets, we follow DIR [34] to use different GNN backbones for various datasets. Among these four datasets, the SP-Motif dataset is synthetic, while the other three datasets are obtained from real-world data. Our code is provided at <https://github.com/SongW-SW/DLG>.

Baselines. To evaluate our proposed framework DLG while comparing other state-of-the-art methods, we consider the following baselines: ① OOD methods proposed for Euclidean data (e.g. images), which include IRM [35] and V-REx [30]. ② OOD methods proposed for graph data, which include DIR [34], GIL [47], CIGA [48], and GALA [50].

TABLE I: The performance of various methods on TPG datasets with different degrees of bias. We report the results for test accuracy (in %), and the best results are in **bold**.

Dataset	TPG	EC50	Ki	SP-Motif	MNIST-75sp	Graph-SST2	Molhiv
ERM	61.77 \pm 1.78	67.53 \pm 1.45	73.54 \pm 1.76	38.81 \pm 1.61	12.71 \pm 1.43	81.44 \pm 0.59	76.20 \pm 1.14
IRM	61.38 \pm 1.63	67.93 \pm 1.28	73.56 \pm 1.98	39.71 \pm 2.11	18.62 \pm 1.22	81.01 \pm 1.13	74.46 \pm 2.74
V-REx	60.94 \pm 1.98	67.70 \pm 1.20	74.08 \pm 1.86	39.04 \pm 1.97	18.92 \pm 1.41	81.76 \pm 0.08	75.62 \pm 0.79
DIR	—	—	—	43.94 \pm 1.70	20.36 \pm 1.78	83.29 \pm 0.53	77.05 \pm 0.57
GIL	64.50 \pm 2.13	64.20 \pm 3.18	73.69 \pm 2.56	52.29 \pm 2.96	21.94 \pm 0.38	83.44 \pm 0.37	79.08 \pm 0.54
CIGA	67.18 \pm 1.85	68.18 \pm 1.57	73.43 \pm 2.71	69.86 \pm 3.41	24.81 \pm 2.08	81.43 \pm 1.26	79.70 \pm 1.55
GALA	79.21 \pm 1.34	69.36 \pm 1.51	76.16 \pm 1.80	—	—	—	—
DLG	79.69 \pm 1.65	69.85 \pm 1.84	76.55 \pm 2.04	73.34 \pm 2.84	31.22 \pm 0.97	84.01 \pm 0.89	79.80 \pm 0.60

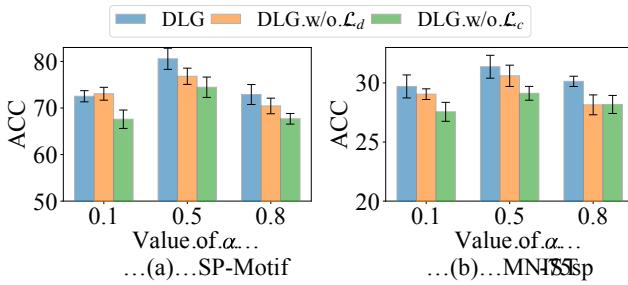


Fig. 2: The performance of our framework DLG with different modules removed.

B. Results of Graph-Level OOD Generalization

We conduct experiments for out-of-distribution generalization on two sets of datasets used in GALA [50] and DIR [18], respectively. We present the results of TPG (Two-Piece Graph), DrugOOD, and the four datasets used in DIR in Table I. From the results, we have the following observations:

❶ DLG achieves the best result across most synthetic and real-world datasets. We observe that the performance of IRM and V-REx is less competitive, demonstrating that directly applying OOD generalization methods to graph data is insufficient. In concrete, these results indicate that DLG maintains superior generalizability compared to other baselines across various datasets. **❷ DLG exhibits superior performance on datasets with node features of fewer dimensions.** Among the datasets evaluated, the Graph-SST2 dataset uniquely preserves node features with a substantial dimension size of 768. In contrast, all other datasets maintain node features with sizes less than 10. This characteristic emphasizes the need to learn consistent representations while concurrently preserving their informativeness. Owing to the proposed informative loss integrated into our framework, DLG effectively retains informativeness within the learned representations, preventing the omission of valuable information when applied to datasets with limited node features.

C. Ablation Study of DLG

In this subsection, we further present several ablation studies on various datasets to validate the significance of the two

key designs within DLG. First, the diversity enhancement component is disabled by removing the loss \mathcal{L}_d , causing the framework unable to preserve distribution consistency. This variant is denoted as DLG w/o \mathcal{L}_d . Second, we discard the consistency enhancement by removing \mathcal{L}_c . This means that the augmented graphs and invariant graphs are not restricted by label consistency enhancement. We denote this version as DLG w/o \mathcal{L}_c . The results from our ablation study are illustrated in Fig. 2. From these results, we observe that DLG outperforms all other variants, which validates the effectiveness of the two modules in our framework. First, integrating the distribution consistency enhancement yields a significant performance boost, especially when the weight α is increased. More significantly, without the label consistency enhancement, i.e., loss \mathcal{L}_c , the performance decreases rapidly on SP-Motif and MNIST-75sp. Those results demonstrate the importance of label consistency enhancement in the presence of a larger distribution shift.

VI. CONCLUSION

In this work, we have investigated the problem of graph out-of-distribution (OOD) generalization on both graph classification and node classification. We demonstrate that the two traditional steps: augmentations and invariant subgraph extraction, may fail to achieve distribution consistency and label consistency, respectively. To deal with this, we propose a novel framework DLG to enhance these two types of consistency. Based on learning edge masks with our designed modifiers, we are able to preserve the consistency in augmented graphs and invariant graphs. We conduct extensive experiments on a variety of both graph-level and node-level OOD generalization datasets. The results demonstrate that our framework DLG consistently outperforms other state-of-the-art graph OOD generalization methods.

ACKNOWLEDGEMENTS

This work is supported in part by the National Science Foundation under grants (IIS-2006844, IIS-2144209, IIS-2223769, CNS-2154962, BCS-2228534, and CMMI-2411248), the Commonwealth Cyber Initiative Awards under grants (VV-1Q24-011, VV-1Q25-004), and the research gift funding from Netflix and Snap.

REFERENCES

- [1] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su, “Arnetminer: extraction and mining of academic social networks,” in *SIGKDD*, 2008. 1
- [2] J. McAuley, R. Pandey, and J. Leskovec, “Inferring networks of substitutable and complementary products,” in *SIGKDD*, 2015. 1
- [3] F. Zhou, C. Cao, K. Zhang, G. Trajcevski, T. Zhong, and J. Geng, “Metagnn: On few-shot node classification in graph meta-learning,” in *CIKM*, 2019. 1
- [4] A. Bojchevski and S. Günnemann, “Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking,” in *ICLR*, 2018. 1
- [5] Z. Tan, K. Ding, R. Guo, and H. Liu, “Graph few-shot class-incremental learning,” in *WSDM*, 2022. 1
- [6] S. Wang, Z. Tan, H. Liu, and J. Li, “Contrastive meta-learning for few-shot node classification,” in *SIGKDD*, 2023. 1
- [7] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *ICLR*, 2017. 1
- [8] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, “Graph attention networks,” in *ICLR*, 2018. 1
- [9] 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*, 2020. 1
- [10] S. Chang, W. Han, J. Tang, G.-J. Qi, C. C. Aggarwal, and T. S. Huang, “Heterogeneous network embedding via deep architectures,” in *SIGKDD*, 2015. 1
- [11] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *NeurIPS*, 2017. 1
- [12] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?” in *ICLR*, 2019. 1
- [13] S. Wang, C. Chen, and J. Li, “Graph few-shot learning with task-specific structures,” in *NeurIPS*, 2022. 1
- [14] S. Cao, W. Lu, and Q. Xu, “Deep neural networks for learning graph representations,” in *AAAI*, 2016. 1
- [15] S. Miao, M. Liu, and P. Li, “Interpretable and generalizable graph learning via stochastic attention mechanism,” in *ICML*, 2022. 1
- [16] J. Yu, J. Liang, and R. He, “Mind the label shift of augmentation-based graph ood generalization,” in *CVPR*, 2023. 1
- [17] S. Wang, Y. Dong, B. Zhang, Z. Chen, X. Fu, Y. He, C. Shen, C. Zhang, N. V. Chawla, and J. Li, “Safety in graph machine learning: Threats and safeguards,” *arXiv:2405.11034*, 2024. 1
- [18] Q. Wu, H. Zhang, J. Yan, and D. Wipf, “Handling distribution shifts on graphs: An invariance perspective,” in *ICLR*, 2022. 1, 2, 5
- [19] J. You, R. Ying, X. Ren, W. Hamilton, and J. Leskovec, “Graphrn: Generating realistic graphs with deep auto-regressive models,” in *ICML*, 2018. 1
- [20] F. M. Bianchi, D. Grattarola, L. Livi, and C. Alippi, “Graph neural networks with convolutional arma filters,” *IEEE TPMI*, 2021. 1
- [21] Y. Mansour, M. Mohri, and A. Rostamizadeh, “Domain adaptation: Learning bounds and algorithms,” in *COLT*, 2009. 1
- [22] G. Blanchard, G. Lee, and C. Scott, “Generalizing from several related classification tasks to a new unlabeled sample,” in *NeurIPS*, 2011. 1
- [23] K. Muandet, D. Balduzzi, and B. Schölkopf, “Domain generalization via invariant feature representation,” in *ICML*, 2013. 1
- [24] S. Beery, G. Van Horn, and P. Perona, “Recognition in terra incognita,” in *ECCV*, 2018. 1
- [25] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do imagenet classifiers generalize to imagenet?” in *ICML*, 2019. 1
- [26] J. Su, D. V. Vargas, and K. Sakurai, “One pixel attack for fooling deep neural networks,” *IEEE TEVC*, 2019. 1
- [27] S. Fakhraei, J. Foulds, M. Shashanka, and L. Getoor, “Collective spammer detection in evolving multi-relational social networks,” in *SIGKDD*, 2015. 1
- [28] S. Gui, X. Li, L. Wang, and S. Ji, “Good: A graph out-of-distribution benchmark,” in *NeurIPS Datasets and Benchmarks Track*, 2022. 1
- [29] S. Wang, Y. Dong, K. Ding, C. Chen, and J. Li, “Few-shot node classification with extremely weak supervision,” in *WSDM*, 2023. 1
- [30] D. Krueger, E. Caballero, J.-H. Jacobsen, A. Zhang, J. Binns, D. Zhang, R. Le Priol, and A. Courville, “Out-of-distribution generalization via risk extrapolation (rex),” in *ICML*, 2021. 1, 4
- [31] S. Chang, Y. Zhang, M. Yu, and T. Jaakkola, “Invariant rationalization,” in *ICML*, 2020. 1
- [32] E. A. AlBadawy, A. Saha, and M. A. Mazurowski, “Deep learning for segmentation of brain tumors: Impact of cross-institutional training and testing,” *Medical Physics*, 2018. 1
- [33] D. Dai and L. Van Gool, “Dark model adaptation: Semantic image segmentation from daytime to nighttime,” in *ITSC*, 2018. 1
- [34] Y. Wu, X. Wang, A. Zhang, X. He, and T.-S. Chua, “Discovering invariant rationales for graph neural networks,” in *ICLR*, 2022. 1, 2, 4
- [35] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, “Invariant risk minimization,” *arXiv*, 2019. 2, 4
- [36] Y. Ganin and V. Lempitsky, “Unsupervised domain adaptation by backpropagation,” in *ICML*, 2015. 2
- [37] H. Li, S. J. Pan, S. Wang, and A. C. Kot, “Domain generalization with adversarial feature learning,” in *CVPR*, 2018. 2
- [38] K. Ding, J. Wang, J. Li, K. Shu, C. Liu, and H. Liu, “Graph prototypical networks for few-shot learning on attributed networks,” in *CIKM*, 2020. 2
- [39] B. Sun, J. Feng, and K. Saenko, “Return of frustratingly easy domain adaptation,” in *AAAI*, 2016. 2
- [40] P. Bühlmann, “Invariance, causality and robustness,” 2020. 2
- [41] C. Heinze-Deml, J. Peters, and N. Meinshausen, “Invariant causal prediction for nonlinear models,” *JCI*, 2018. 2
- [42] W. Hu, G. Niu, I. Sato, and M. Sugiyama, “Does distributionally robust supervised learning give robust classifiers?” in *ICML*, 2018. 2
- [43] Q. Qian, S. Zhu, J. Tang, R. Jin, B. Sun, and H. Li, “Robust optimization over multiple domains,” in *AAAI*, 2019. 2
- [44] S. Sagawa, P. W. Koh, T. B. Hashimoto, and P. Liang, “Distributionally robust neural networks,” in *ICLR*, 2020. 2
- [45] Y. Wu, A. Bojchevski, and H. Huang, “Adversarial weight perturbation improves generalization in graph neural networks,” in *AAAI*, 2023. 2
- [46] Y. Sui, X. Wang, J. Wu, A. Zhang, and X. He, “Adversarial causal augmentation for graph covariate shift,” *arXiv:2211.02843*, 2022. 2
- [47] H. Li, Z. Zhang, X. Wang, and W. Zhu, “Learning invariant graph representations for out-of-distribution generalization,” in *NeurIPS*, 2022. 2, 4
- [48] Y. Chen, Y. Zhang, Y. Bian, H. Yang, M. Kaili, B. Xie, T. Liu, B. Han, and J. Cheng, “Learning causally invariant representations for out-of-distribution generalization on graphs,” *NeurIPS*, 2022. 2, 4
- [49] S. Fan, X. Wang, Y. Mo, C. Shi, and J. Tang, “Debiasing graph neural networks via learning disentangled causal substructure,” *NeurIPS*, 2022. 2
- [50] Y. Chen, Y. Bian, K. Zhou, B. Xie, B. Han, and J. Cheng, “Does invariant graph learning via environment augmentation learn invariance?” *Advances in Neural Information Processing Systems*, vol. 36, 2023. 2, 4, 5
- [51] E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with gumbel-softmax,” in *ICLR*, 2017. 3
- [52] Y. Sui, X. Wang, J. Wu, M. Lin, X. He, and T.-S. Chua, “Causal attention for interpretable and generalizable graph classification,” in *SIGKDD*, 2022. 3
- [53] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv:1807.03748*, 2018. 3
- [54] D. Luo, W. Cheng, D. Xu, W. Yu, B. Zong, H. Chen, and X. Zhang, “Parameterized explainer for graph neural network,” *Advances in neural information processing systems*, vol. 33, pp. 19 620–19 631, 2020. 4
- [55] Y. Ji, L. Zhang, J. Wu, B. Wu, L.-K. Huang, T. Xu, Y. Rong, L. Li, J. Ren, D. Xue *et al.*, “Drugood: Out-of-distribution (ood) dataset curator and benchmark for ai-aided drug discovery—a focus on affinity prediction problems with noise annotations,” *arXiv:2201.09637*, 2022. 4
- [56] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, “Gnnexplainer: Generating explanations for graph neural networks,” in *NeurIPS*, 2019. 4
- [57] B. Knyazev, G. W. Taylor, and M. Amer, “Understanding attention and generalization in graph neural networks,” in *NeurIPS*, 2019. 4
- [58] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. Pande, “Moleculenet: a benchmark for molecular machine learning,” *Chemical Science*, 2018. 4
- [59] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec, “Open graph benchmark: Datasets for machine learning on graphs,” in *NeurIPS*, 2020. 4
- [60] W. Hu, M. Fey, H. Ren, M. Nakata, Y. Dong, and J. Leskovec, “Ogb-lsc: A large-scale challenge for machine learning on graphs,” *arXiv:2103.09430*, 2021. 4