

# Interpolating Graph Pair to Regularize Graph Classification

Hongyu Guo<sup>1,2</sup>, Yongyi Mao<sup>2</sup>

<sup>1</sup>National Research Council Canada

<sup>2</sup>University of Ottawa

{hongyu.guo, ymao}@uottawa.ca

## Abstract

We present a simple and yet effective interpolation-based regularization technique, aiming to improve the generalization of Graph Neural Networks (GNNs) on supervised graph classification. We leverage Mixup (Zhang et al. 2018a), an effective regularizer for vision, where random sample pairs and their labels are interpolated to create synthetic images for training. Unlike images with grid-like coordinates, graphs have arbitrary structure and topology, which can be very sensitive to any modification that alters the graph’s semantic meanings. This posts two unanswered questions for Mixup-like regularization schemes: Can we directly mix up a pair of graph inputs? If so, how well does such mixing strategy regularize the learning of GNNs? To answer these two questions, we propose *ifMixup*, which first adds dummy nodes to make two graphs have the same input size and then simultaneously performs linear interpolation between the aligned node feature vectors and the aligned edge representations of the two graphs. We empirically show that such simple mixing schema can effectively regularize the classification learning, resulting in superior predictive accuracy to popular graph augmentation and GNN methods.

## Introduction

Graph Neural Networks (GNNs) (Kipf and Welling 2017) have recently shown promising performance in many challenging applications, including predicting molecule property (Wu et al. 2018), forecasting protein activation (Jiang et al. 2017), and estimating circuit functionality (Zhang, He, and Katabi 2019). Nevertheless, like other successfully deployed deep networks, GNNs also suffer from the data-hungry issue due to their over-parameterized learning paradigm. Consequently, regularization techniques have been actively proposed, aiming to empower the learning of GNNs while avoiding over-smoothing (Li, Han, and Wu 2018), over-squashing (Alon and Yahav 2021) and over-fitting (Zhang et al. 2017). Amongst these generalization strategies, data augmentation schemes have been proven to be very effective (Rong et al. 2020; Zhou, Shen, and Xuan 2020). Existing graph data augmentation methods, nevertheless, mostly involve graph manipulations on a *single graph* (Rong et al. 2020; Zhou, Shen, and Xuan 2020; You et al. 2020).

In this paper, we look into a very successful *pairwise data augmentation* technique for image recognition (Zhang et al.

2018a; Verma et al. 2018; Guo, Mao, and Zhang 2019b; Kim, Choo, and Song 2020), called Mixup. Mixup was originally introduced by (Zhang et al. 2018a) as an interpolation-based regularizer for image classification. It regularizes the learning of deep classification models by training with synthetic samples, which are created by linearly interpolating a pair of randomly selected training samples, naturally well-aligned, as well as their training targets.

Motivated by the effectiveness and simplicity of Mixup in regularizing image classification models, we are motivated to design a similar, straight-forward, “Mixup” scheme for graph data. Nevertheless, unlike images, which are supported on grid coordinates, graph data have arbitrary structure and topology, such as having different numbers of nodes and these nodes from different graphs are typically not readily aligned. Furthermore, such irregular graph topologies typically play a critical role in the graph semantics, and consequently even simply modifying one node or edge can dramatically change the semantic meaning of a graph. Hence, the following two questions naturally arise:

*Can we directly mix up a pair of graph inputs? If so, how well does such mixing strategy regularize the learning of GNNs?*

To answer these two questions, we propose a simple input mixing strategy for Mixup on graph, coined *ifMixup* for graph-level classification. *ifMixup* first samples random graph pairs from the training data, and then creates a synthetic graph through mixing each selected sample pair, using a mixing ratio sampled from a Beta distribution. For mixing, *ifMixup* first adds disconnected dummy nodes to make the two selected graphs have the same input size, so that it can align the two graphs with an arbitrary node order. Next, it simultaneously performs linear interpolation between the aligned node feature vectors and the aligned edge representations of the two graph inputs, through leveraging the straight-forward mixing strategy as that of the original Mixup on image pairs. The newly generated mixed graphs, which have a much larger number of graphs with changing local neighborhood properties than the original training dataset, are then used for training to regularize the GNNs learning.

We conduct extensive experiments, using eight benchmarking tasks from various domains, showing that our strategy can effectively regularize the graph classification to improve its predictive accuracy, outperforming popular graph augmen-




Figure 1: The proposed mixing schema. The left figure depicts the source graph pair, and the middle subfigure pictures the adding dump node process, and the right subfigure is the resulting mixed graph with mixing ratio  $\lambda$  of 0.75.

tation approaches and GNN methods. We also theoretically prove that such a simple mixing strategy processes an unique property under a mild assumption: the mixed graph preserves all the information from its original graph pair, namely the mixing process is information lossless.

## Related Work

Most of the graph augmentation strategies are for node classification tasks, and heavily focus on perturbing nodes and edges in one given graph (Hamilton, Ying, and Leskovec 2017; Zhang et al. 2018b; Chen et al. 2020; Zhou, Shen, and Xuan 2020; Qiu et al. 2020; You et al. 2020; Wang et al. 2020b; Fu et al. 2020; Wang et al. 2020a; Song et al. 2021; Zhao et al. 2021; Liu et al. 2022; Zhao et al. 2022; Ding et al. 2022; Guo and Sun 2022). Unlike these approaches, our proposed strategy leverages a pair of graphs, instead of one graph, to augment the learning of graph level classification.

Despite its great success in augmenting data for image and text classification (Zhang et al. 2018a; Guo, Mao, and Zhang 2019a; Verma et al. 2018; Guo, Mao, and Zhang 2019b; Guo 2020; Jindal et al. 2020; Kim, Choo, and Song 2020), Mixup has been less explored for graph learning. Mixup-Graph (Wang et al. 2021) leverages a simple way to avoid dealing with the arbitrary structure in the input space for mixing a graph pair, through mixing the graph representation resulting from passing the graph through the GNNs. Similarly, a concurrent work G-Mixup (Han et al. 2022) first interpolates represented graph generators (i.e., graphons) of different classes, and then leverages the mixed graphons for sampling to generate synthetic graphs.

To our best knowledge, there are only a few attempts on mixing graph inputs (Park, Shim, and Yang 2022; Yoo, Shim, and Kang 2022). These methods, however, rely on cumbersome processes to first select subgraphs from the given graph pair and then design new edges to connect these subgraphs to form the mixed graph. In contrast, our proposed input mixing strategy is much simpler, following the straight-forward mixing idea from the original Mixup (Zhang et al. 2018a).

## Preliminaries

**Graph Classification with GNNs** In the graph classification problem we consider, the input  $G$  is a graph labelled with node features, or a “node-feathered graph”. Specifically, assume that the graph has  $n$  nodes, each identified with an integer ID in  $[n] := \{1, 2, \dots, n\}$ . The set of edges  $E$  of the graph is a collection of unordered pairs  $(i, j)$ ’s of node-IDs, as the current paper considers only undirected graphs (although there is no difficulty to extend the setting to directed

graphs). Associated with node  $i$ , there is a feature vector  $v(i)$  of dimension  $d$ . We will use  $v$  to denote the collection of all feature vectors and one may simply regards  $v$  as a matrix of dimension  $n \times d$ . Thus, each input node-feathered graph  $G$  is essentially specified via the pair  $(v, E)$ . The output  $y$  is a class label in a finite set  $\mathcal{Y} := \{1, 2, \dots, C\}$ , which will be expressed as a (one-hot) probability vector over the label set  $\mathcal{Y}$ . The classification problem is to find a mapping that predicts the label  $y$  for a node-feathered graph  $G$ . The training data  $\mathcal{G}$  of this task is a collection of such  $(G, y)$  pairs.

Modern GNNs use the graph structure and node features to learn a distributed vector to represent a graph. The learning follows the “message passing” mechanism for neighborhood aggregation. It iteratively updates the embedding of a node  $h_v$  by aggregating representations/embeddings of its neighbors. The entire graph representation  $h_G$  is then obtained through a READOUT function, which aggregates embeddings from all nodes of the graph. Formally, representation  $h_i^k$  of node  $i$  at the  $k$ -th layer of a GNN is defined as:

$$h_i^k = \text{AGGREGATE}(h_i^{k-1}, h_j^{k-1} | j \in \mathcal{N}(i), W^k), \quad (1)$$

where  $W^k$  denotes the trainable weights at layer  $k$ ,  $\mathcal{N}(i)$  denotes the set of all nodes adjacent to  $i$ , and AGGREGATE is an aggregation function implemented by specific GNN model (popular ones include Max, Mean, Sum pooling operations), and  $h_i^0$  is typically initialized as the input node feature  $v(i)$ . The graph representation  $h_G$  aggregates node representations  $h_v$  using the READOUT graph pooling function:

$$h_G = \text{READOUT}(h_i^k | i \in [n]). \quad (2)$$

This graph representation is then mapped to label  $y$  using a standard classification network (e.g., a softmax layer).

**Mixup Interpolation for Images** Mixup was introduced by (Zhang et al. 2018a) as an interpolation-based regularizer for image classification. It regularizes the learning of deep classification models by training with synthetic samples, which are created by linearly interpolating a pair of randomly selected training samples as well as their modeling targets. In detail, let  $(x_A, y_A)$  and  $(x_B, y_B)$  be two training instances, in which  $x_A$  and  $x_B$  refer to the input images and  $y_A$  and  $y_B$  refer to their corresponding labels. For a randomly chosen such training pair, Mixup generates a sample as follows.

$$\tilde{x} = \lambda x_A + (1 - \lambda)x_B, \quad (3)$$

$$\tilde{y} = \lambda y_A + (1 - \lambda)y_B, \quad (4)$$

where  $\lambda$  is a scalar mixing ratio, sampled from a Beta( $\alpha, \beta$ ) distribution with hyper-parameters  $\alpha$  and  $\beta$ . Such synthetic instances  $(\tilde{x}, \tilde{y})$ ’s are then used for training.

Motivated by the simplicity and effectiveness of Mixup in regularizing image classification models, we are naturally motivated to design a similar ‘‘Mixup’’ scheme for graph data, in particular, the node-featured graphs, as are the interest of this paper. When this is possible, we may use the synthetic instances  $(\tilde{G}, \tilde{y})$ 's to learn the model parameter  $\theta$  by minimizing the loss  $\mathcal{L}$ :

$$\min_{\theta} \mathbb{E}_{(G_A, y_A) \sim \mathcal{G}, (G_B, y_B) \sim \mathcal{G}, \lambda \sim \text{Beta}(\alpha, \beta)} [\mathcal{L}(\tilde{G}, \tilde{y})]. \quad (5)$$

To mix  $(G_A, y_A)$  and  $(G_B, y_B)$ , it is straight-forward to apply Equation (4) to obtain the mixed label  $\tilde{y}$ . The key question of investigation is how to mix  $G_A$  and  $G_B$  to obtain  $\tilde{G}$ . We will answer this question next.

## The Proposed Method

### Graph Mixing Strategy

Given a node featured graph  $G = (v, E)$ , we represent  $E$  as a binary matrix  $e$  with  $n$  rows and  $n$  columns, in which  $e(i, j) = 1$  if  $(i, j) \in E$ , and  $e(i, j) = 0$  otherwise. Thus instead of expressing  $G$  as  $(v, E)$  we express it as  $(v, e)$ . The mixing of  $G_A = (v_A, e_A)$  with  $G_B = (v_B, e_B)$  to obtain  $\tilde{G} = (\tilde{v}, \tilde{e})$  can simply be done as follows:

$$\tilde{e} = \lambda e_A + (1 - \lambda) e_B. \quad (6)$$

$$\tilde{v} = \lambda v_A + (1 - \lambda) v_B. \quad (7)$$

In order for the above mixing rule to be well defined, we need the two graphs to have the same number of nodes. For this purpose we define  $n = \max(n_A, n_B)$ , where  $n_A$  and  $n_B$  are the number of nodes in instances  $A$  and  $B$  respectively. If  $G_A$  or  $G_B$  has less than  $n$  nodes, we simply introduce dummy node to the graph and make them disconnected from the existing nodes. The feature vectors for the dummy nodes are set to the all-zero vector.

This mixing process is illustrated in Figure 1, where the left is the source graph pair, and the middle depicts the added dummy node (i.e., the node with dotted double circles). The right is the resulting mixed graph, where edge weights sit next to their edges.

It is worth noting that the resulting mixed graphs, through Equations 6 and 7, contain edges with weights between  $[0, 1]$ . As a result, during training, this will require the GNN networks be able to take the edge weights into account for message passing. Fortunately, the two popular GNN networks, namely GCN (Kipf and Welling 2017) and GIN (Xu et al. 2019), can naturally cope with weighted edges in their implementation of Equation 1.

GCN handles edge weights naturally by enabling adjacency matrix to have values between zero and one (Kipf and Welling 2017), instead of either zero or one, representing edge weights:

$$\mathbf{h}_i^k = \sigma \left( W^k \cdot \left( \sum_{j \in \mathcal{N}(i) \cup \{i\}} \frac{e(i, j)}{\sqrt{\hat{d}_j \hat{d}_i}} \mathbf{h}_j^{k-1} \right) \right), \quad (8)$$

where  $\hat{d}_i = 1 + \sum_{j \in \mathcal{N}(i)} e(i, j)$ ;  $W^k$  stands for the trainable weights at layer  $k$ ;  $\sigma$  denotes the non-linearity transformation, i.e. the ReLU function.

---

### Algorithm 1: The mixing schema in ifMixup

---

**Input:** a graph pair  $G_A = (v_A, e_A)$  and  $G_B = (v_B, e_B)$  (all edges in  $e$  have weight 1) **Parameter:** mixing ratio  $\lambda \in (0, 1)$

**Output:** a mixed graph  $\tilde{G} = (\tilde{v}, \tilde{e})$

- 1: Compute max node number,  $n = \max(n_A, n_B)$
  - 2: **if**  $n_A < n$  **then**
  - 3:   Add  $n - n_A$  dummy nodes to  $G_A$
  - 4: **else if**  $n_B < n$  **then**
  - 5:   Add  $n - n_B$  dummy nodes to  $G_B$
  - 6: **end if**
  - 7:  $\tilde{e} = \lambda e_A + (1 - \lambda) e_B$
  - 8:  $\tilde{v} = \lambda v_A + (1 - \lambda) v_B$
  - 9: **return** mixed graph  $\tilde{G} = (\tilde{v}, \tilde{e})$
- 

To enable GIN to handle soft edge weight, we replace the sum operation of the isomorphism operator in GIN with a weighted sum calculation. That is, the GIN updates node representations as:

$$\mathbf{h}_i^k = \text{MLP}^k \left( (1 + \epsilon^k) \cdot \mathbf{h}_i^{k-1} + \sum_{j \in \mathcal{N}(i)} e(i, j) \cdot \mathbf{h}_j^{k-1} \right), \quad (9)$$

where  $\epsilon^k$  is a learnable parameter. The pseudo-code of ifMixup is described in Algorithm 1.

### Information Lossless Property

In this section, we show that a mixed sample generated by ifMixup preserves all the information from its original sample pair used for mixing, namely being information lossless.

Such information lossless is owing to the fact that the mixing strategy in ifMixup guarantees that the original two node-featured graphs  $G_A$  and  $G_B$  recoverable from the mixed graph  $\tilde{G}$  under a mild assumption. In detail, to be able to recover the original graphs  $G_A$  and  $G_B$ , both the *graph topology* and the *node features* of the  $G_A$  and  $G_B$  graphs need be recovered from the mixed instance  $\tilde{G}$ . We refer to the former as *edge invertibility* and the latter as *node feature invertibility*, as will be discussed in detail next.

**Lemma 0.1 (Edge Invertibility).** *Let  $\tilde{e}$  be constructed using Equation 6 with  $\lambda \neq 0.5$ . Consider equation*

$$se + (1 - s)e' = \tilde{e}$$

*with unknowns  $s$ ,  $e$  and  $e'$ , where  $s$  is a scalar and  $e$  and  $e'$  are binary (i.e.,  $\{0, 1\}$ -valued)  $n \times n$  matrices. There are exactly two solutions to this equation:*

$$\begin{cases} s = \lambda, & e = e_A, & e' = e_B, \text{ or} \\ s = 1 - \lambda, & e = e_B, & e' = e_A \end{cases}$$

By this lemma, we see that if the mixing coefficient  $\lambda \neq 0.5$ , from the mixed edge representation  $\tilde{e}$ , we can always recover  $e_A$  and  $e_B$  (and hence  $E_A$  and  $E_B$ ) and their corresponding weights used for mixing. Note that if  $\lambda$  is drawn from a continuous distribution over  $(0, 1)$ , the probability it takes value 0.5 is zero. That is, the connectivity of the original two graphs can be perfectly recovered from the

mixed edge representation  $\tilde{e}$ . In other words, from  $\tilde{G}$ , ifMixup can perfectly recover the topologies (edge connections) of  $G_A$  and  $G_B$ .

**Lemma 0.2 (Node Feature Invertibility).** *Suppose that the node feature vectors for all instances in the task take values from a finite set  $V \subset \mathbb{R}^d$  and that  $V$  is linearly independent. Let  $\tilde{v}$  be constructed using Equation 7. Let  $V^* = V \cup \{\mathbf{0}\}$ , where  $\mathbf{0}$  denotes the zero vector in  $\mathbb{R}^d$ . Consider equation*

$$\tilde{v} = sv + (1 - s)v'$$

*in which  $n \times d$  matrices  $v$  and  $v'$  are unknowns with rows taking value in  $V^*$ . For any fixed  $s \in (0, 1)$ , there is exactly one solution of  $(v, v')$  for this equation.*

By this lemma, we see that if the node *feature set*  $V$  is linearly independent, ifMixup can perfectly recover the node features of the two original graphs from the mixed graph. In other words, from  $\tilde{G}$ , ifMixup can recover all the node features for  $G_A$  and  $G_B$ .

In summary, these two lemmas state that, given a mixed graph  $\tilde{G}$  ifMixup can perfectly recover the original two node-featured graphs  $G_A$  and  $G_B$ , both their graph topologies and all the node features. Such invertibility makes sure that a mixed graph in ifMixup preserves all the information from its original sample pair, being information lossless.

## Discussion

**Manifold Intrusion** The manifold intrusion degrades the performance of Mixup-like methods (Guo, Mao, and Zhang 2019b). It refers to a form of under-fitting, resulting from conflicts between the labels of the synthetic examples and the labels of original training data (Guo, Mao, and Zhang 2019b). Under the context of graph classification, the manifold intrusion represents that the generated graphs have identical topology but different label than some original graphs.

**Theorem 0.3 (Intrusion-Freeness).** *Suppose that  $\lambda \neq 1/2$  and that the condition for Lemma 0.2 is satisfied. Then for any mixed node-featured graph  $\tilde{G} = (\tilde{v}, \tilde{e})$  constructed using Equations 6 and 7, the two original node-feature graph  $G_A$  and  $G_B$  can be uniquely recovered.*

*Proof:* Since  $\lambda \neq 0.5$ , by Lemma 0.1 we can recover  $e_A$ ,  $e_B$  and  $\lambda$  from  $\tilde{e}$ . Given  $\lambda$  and the condition for Lemma 0.2, we can recover  $v_A$  and  $v_B$  from  $\tilde{v}$ .  $\square$

By this theorem, there is no *other pair*  $(G'_A, G'_B)$  (different from  $G_A$  and  $G_B$ ) from the training set  $\mathcal{G}$  that can be mixed into  $\tilde{G}$  using any  $\lambda$  due to the *uniqueness* in the invertibility. That is, it is impossible that the mixed graph  $\tilde{G}$  to coincide with another mixed graph  $\tilde{G}'$  resulting from any other graph pairs or with any other graphs in the training set  $\mathcal{G}$  (which has one-hot label). Thus, manifold intrusion does not occur under the theorem.

**Impact of Node Alignment in Mixing** It is worth noting that, ifMixup mixes different graphs that are not naturally well-aligned, which seems “counter-intuitive”. Nevertheless, such “counter-intuition” technique has been proved to work very well for augmenting images and texts, where Mixup

has been successfully deployed to mix two unrelated images (Zhang et al. 2018a; Thulasidasan et al. 2019) or two semantically different words in sentences (Guo, Mao, and Zhang 2019a; Guo 2020). Also, to this end, ifMixup leverages an arbitrary node order for mixing after aligning the two input graphs by adding dummy nodes. As a result, different node order for the same input graph pair will result in a different mixed graph. We empirically evaluate the impact of such node alignment in the Ablation section in the experiments.

**Impact of Graph Size for Mixing** ifMixup leverages adding dummy nodes to align the input graph pair, which is essentially a padding-based strategy. In this sense, the size of the resulting mixed graph, through mixing two graphs with large difference in node numbers, can be different from that of its source input pair. This is also expected to impact ifMixup’s predictive performance. We also conduct experiments to evaluate such effect in the Ablation section.

**Graph with Weighted Edges and Edge Features** ifMixup assumes that the given graphs for training have binary edges, which is a widely adopted setting in learning from topological graphs. In the applications of graphs with real-valued edge weights, two main strategies can be used by ifMixup to cope with such situation in the literature. One common strategy is to add the edge weights to the features of the nodes (Hu et al. 2020a; Li et al. 2020). Another approach is to treat the edge weights the same way as node features in the Aggregation function of the GNNs (Gilmer et al. 2017; Xu et al. 2019; Kipf and Welling 2017; Hu et al. 2020b). It is also worth noting that the above two approaches can also be used by ifMixup to handle graphs that come with edge attributes.

## Experiments

### Settings

**Datasets** We use eight graph classification tasks from the graph benchmark datasets collection TUDatasets (Morris et al. 2020): PTC\_MR, NCI109, NCI1, and MUTAG for small molecule classification, ENZYMES and PROTEINS for protein categorization, and IMDB-M and IMDB-B for social networks classification. These datasets have been widely used for benchmarking such as in (Xu et al. 2019) and can be downloaded directly using PyTorch Geometric (Fey and Lenssen 2019)’s build-in function online <sup>1</sup>. The social networks datasets IMDB-M and IMDB-B have no node features, and we use the node degrees as feature as in (Xu et al. 2019).

We note that, the node features of all these eight established benchmark datasets are *one-hot coded* as implemented by the built-in `torch_geometric.datasets` functions in the Pytorch Geometric platform. As a result, the node *feature sets* of all these eight popular benchmark datasets are *linearly independent*, satisfying the mild assumption as stated in Lemma 0.2 for ifMixup to be completely information lossless.

**Comparison Baselines** We compare our method with seven baselines, including MixupGraph (Wang et al. 2021), DropEdge (Rong et al. 2020), DropNode (Hamilton, Ying, and Leskovec 2017; Chen, Ma, and Xiao 2018; Huang et al.

<sup>1</sup><https://chrsmrrs.github.io/datasets/docs/datasets>

	GCN Baseline	ifMixup	MixupGraph	DropEdge	DropNode	Attr. Masking	Rel. Impr.
PTC_MR	0.621±0.018	0.654±0.003	0.633±0.012	0.653±0.007	0.648±0.018	<b>0.666±0.009</b>	5.31%
NCI109	0.803±0.001	<b>0.820±0.005</b>	0.801±0.005	0.801±0.001	0.793±0.015	0.772±0.002	2.12%
NCI1	0.804±0.005	<b>0.819±0.004</b>	0.808±0.004	0.811±0.002	0.805±0.019	0.789±0.001	1.87%
MUTAG	0.850±0.011	<b>0.879±0.003</b>	0.860±0.006	0.855±0.008	0.829±0.006	0.806±0.004	3.41%
ENZYMES	0.541±0.001	<b>0.570±0.014</b>	0.551±0.016	0.566±0.006	0.532±0.006	0.506±0.021	5.36%
PROTEINS	0.742±0.003	<b>0.753±0.008</b>	0.742±0.003	0.750±0.003	0.748±0.001	0.748±0.005	1.48%
IMDB-M	0.515±0.002	<b>0.523±0.004</b>	0.513±0.003	0.514±0.001	0.512±0.003	0.514±0.003	1.55%
IMDB-B	0.758±0.004	<b>0.763±0.003</b>	0.759±0.002	0.762±0.004	0.761±0.005	0.756±0.003	0.66%

Table 1: Accuracy with GCN. We report mean accuracy over 3 runs of 10-fold cross validation with standard deviations (denoted  $\pm$ ). The relative improvement of ifMixup over the GCN is provided in the last column of the table. Best results are in Bold.

	GAT	GATv2	ifMixup
PTC_MR	0.635±0.020	0.647±0.002	<b>0.654±0.003</b>
NCI109	0.769±0.005	0.768±0.007	<b>0.820±0.005</b>
NCI1	0.784±0.004	0.791±0.014	<b>0.819±0.004</b>
MUTAG	0.831±0.002	0.835±0.010	<b>0.879±0.003</b>
ENZYMES	0.468±0.011	0.477±0.021	<b>0.570±0.014</b>
PROTEINS	0.743±0.002	0.735±0.010	<b>0.753±0.008</b>
IMDB-M	0.514±0.004	0.511±0.001	<b>0.523±0.004</b>
IMDB-B	0.755±0.017	0.753±0.002	<b>0.763±0.003</b>

Table 2: Comparison of GAT, GATv2, and ifMixup with GCN as baseline. We report mean accuracy over 3 runs of 10-fold cross validation with standard deviations (denoted  $\pm$ ). The relative improvement of ifMixup over the baseline GCN is provided in the last column. Best results are in Bold.

2018), Attribute Masking and Baseline. MixupGraph applies Mixup on graph classification. It leverages a simple way to avoid dealing with the arbitrary structure for mixing a graph pair, through mixing the entire graph representation resulting from the READOUT function of the GNNs. DropEdge and DropNode are two widely used graph perturbation strategies for graph augmentation. DropEdge randomly removes a set of existing edges from a given graph. DropNode randomly deletes a portion of nodes and their connected edges. Attribute Masking represents one of the state-of-the-art methods for graph data augmentation (You et al. 2020). We here assign random attribute values to 20% of the graph nodes. Also, since ifMixup utilizes weighted edges, we also further compare our method with two state-of-the-art attention-based graph networks: GAT (Veličković et al. 2018) and GATv2 (Brody, Alon, and Yahav 2022), which leverage learned edge weights for neighbour aggregation. For these two methods, we also used the implementation from the Pytorch Geometric (Fey and Lenssen 2019) platform.

For the Baseline model, we use two popular GNNs network architectures: GCN (Kipf and Welling 2017) and GIN (Xu et al. 2019). GCN leverages spectral-based convolutional operation to learn spectral features of graph through aggregation, benefiting from a normalized adjacency matrix, while GIN leverages the nodes’ spatial relations to aggregate neighborhood features, representing the state-of-the-art GNN network architecture. We use their implementations in the PyTorch Geometric platform<sup>2</sup>. Note that, for the GCN, we use the GCN with Skip Connection (He et al. 2016) as that in (Li et al. 2019).

<sup>2</sup>[https://github.com/pyg-team/pytorch\\_geometric](https://github.com/pyg-team/pytorch_geometric)

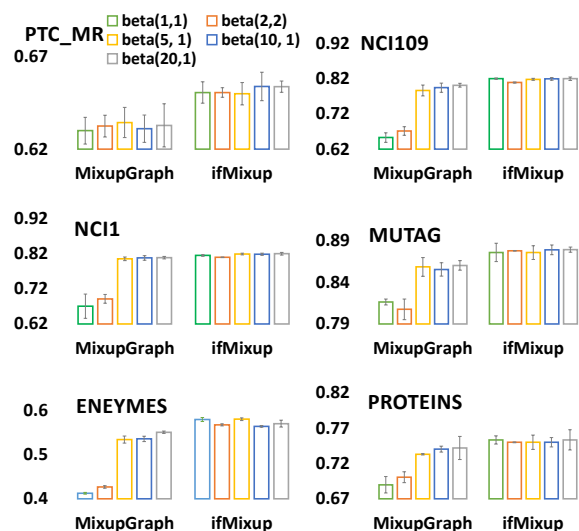


Figure 2: Accuracy of MixupGraph and ifMixup with mixing ratios sampled from Beta(1, 1), Beta(2, 2), Beta(5, 1), Beta(10, 1) and Beta(20, 1) on the six benchmark datasets.

**Detail Settings** We follow the evaluation protocol and hyperparameters search of GIN (Xu et al. 2019) and DropEdge (Rong et al. 2020). We evaluate the models using 10-fold cross validation, and report the mean and standard deviation of three runs on a cluster with GPU nodes of NVidia V100 with 32 GB memory. Each fold is trained with 350 epochs with AdamW optimizer (Kingma and Ba 2015), and the initial learning rate is reduced by half every 50 epochs. The hyper-parameters we search for all models on each dataset are as follows: (1) initial learning rate  $\in \{0.01, 0.0005\}$ ; (2) hidden unit of size  $\in \{64, 128\}$ ; (3) batch size  $\in \{32, 128\}$ ; (4) dropout ratio after the dense layer  $\in \{0, 0.5\}$ ; (5) DropNode and DropEdge drop ratio  $\in \{20\%, 40\%\}$ ; (6) number of layers in GNNs  $\in \{3, 5, 8\}$ ; (7) Beta distribution for ifMixup, MixupGraph and Manifold Mixup  $\in \{\text{Beta}(1, 1), \text{Beta}(2, 2), \text{Beta}(20, 1)\}$ . Following GIN (Xu et al. 2019) and DropEdge (Rong et al. 2020), we report the case giving the best 10-fold average cross-validation accuracy.

## Results of GCN (with Skip Connection) as Baseline

The accuracy obtained by the GCN (with Skip Connection) baseline, ifMixup, MixupGraph, DropEdge, DropNode and Attribute Masking with GCN (with Skip Connection) on the eight datasets are presented in Table 1 (best results in Bold).

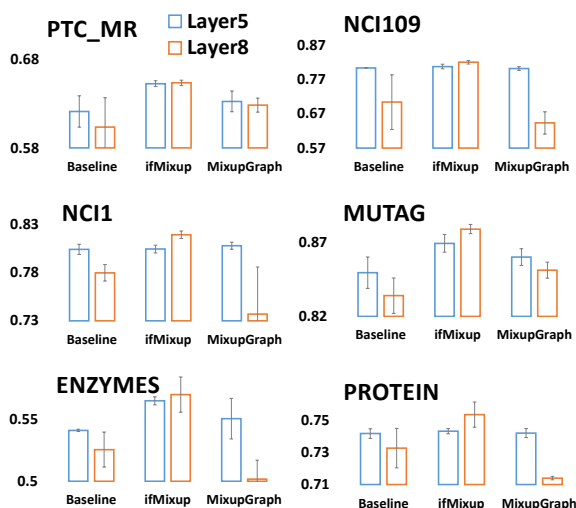


Figure 3: Accuracy on the six benchmark datasets when varying the network depth of GCN, ifMixup, and MixupGraph.

Results in Table 1 show that ifMixup outperformed all the five comparison models against all the eight datasets, except for Attribute Masking on the PTC\_MR dataset. For example, when comparing with the GCN baseline, ifMixup obtained a relative accuracy improvement of 5.36%, 5.31%, and 3.41% on the ENZYMES, PTC\_MR, and MUTAG datasets, respectively. When considering the comparison with the Mixup-like approach MixupGraph, ifMixup also obtained superior accuracy on all the eight datasets. For example, ifMixup was able to improve the accuracy over MixupGraph from 80.1%, 80.8%, 63.3%, and 51.3% to 82.0%, 81.9%, 65.4%, and 52.3% on the NCI109, NCI1, PTC\_MR and IMDB-M datasets, respectively. When comparing with GAT and GATv2, Table 2 shows that ifMixup outperformed the two attention-based strategies, and the improvement is with a large margin in NCI09, NCI1, MUTAG, and ENZYMES.

## Ablation Study

**Sensitivity of Mixing Ratio** We evaluate the sensitivity of the graph mixing ratio to the two Mixup-like approaches: ifMixup and MixupGraph. We present the accuracy obtained by these two methods, with Beta distribution as Beta(1, 1), Beta(2, 2), Beta(5, 1), Beta(10, 1) and Beta(20, 1) on the six molecule datasets. Results are presented in Figure 2. Note that, the mixing ratios sampled from Beta(1, 1) follow a uniform distribution between (0, 1), and those sampled from Beta(2, 2) follow a Bell-Shaped distribution. Those mixing ratios have a wide range. On the other hand, ratios sampled from Beta(5, 1), Beta(10, 1) and Beta(20, 1) mostly fall in the range of (0.8, 1), representing a conservative mixing ratios.

Results in Figure 2 show that both MixupGraph and ifMixup obtained superior results on the six testing datasets with Beta(20, 1). Nevertheless, MixupGraph seemed to very sensitive to the mixing ratio distribution. For example, when Beta distributions were (1, 1) and (2, 2) (first two bars in Figure 2), MixupGraph significantly degraded its accuracy on all the six tasks (except for PTC\_MR). In contrast, ifMixup was robust to the five Beta distributions we tested.

**Impact of GNN Depth** We also evaluate the accuracy obtained by GCN, ifMixup and MixupGraph on the six molecule datasets, when varying the number of layers of the GCN. The results for all the six datasets are depicted in Figure 3. Results in this figure show that when increasing the GCN networks from 5 layers (blue bars) to 8 layers (red bars), both GCN and MixupGraph seemed to degrade its performance on all the six datasets. For example, for the NCI109 and NCI1 datasets, MixupGraph resulted in about 10% of accuracy drop when increasing the number of layers in GCNs (with Skip Connection) from 5 to 8. On the contrary, ifMixup was able to increase the accuracy on all the six datasets tested.

**Impact of Node Order** In this section, we evaluate the impact of node order permutation for ifMixup. We randomly shuffle the node order of one of the graphs in the graph pair before being mixing by ifMixup. We here vary the shuffling frequency: shuffling for every  $k$  epochs. Here we test  $k$  with 60, 30, 10, and 1 (frequency from low to high). We conduct experiments using 8-layer GCN on the NCI109 and NCI1 datasets, and present the results in Table 4. Results in Table 4 show that shuffling the graph for each epoch (*high frequency*) may hurt the performance (e.g., for the NCI1 dataset), but ifMixup seems to be insensitive to *less frequent* shuffling scenarios, such as shuffling every 60, 30, and 10 epochs.

We suspect that the above accuracy degradation of ifMixup when shuffling graph nodes for each epoch was due to the limit of the modeling power of the ifMixup with 8-layer GCN. To verify this hypothesis, we experiment with deeper GCN networks for ifMixup by doubling the network depth, namely using 16 layers. The results are presented in Table 5.

Table 5 indicates that when increasing ifMixup’s modeling capability by using deeper GCN networks, i.e., 16 layers, shuffling for each epoch seemed not to degrade ifMixup’s accuracy, but was able to improve its accuracy. For example, ifMixup that shuffles the graph nodes for each epoch increased the accuracy from 81.4% to 82.2% when its depth was increased from 8 layers to 16 layers on the NCI1 dataset. As shown in Table 5, ifMixup with shuffling achieved the best accuracy for both NCI109 and NCI1 when using 16-layer networks while shuffling the graph nodes for each epoch.

These results suggest that, when the modeling capability is limited, frequent shuffling in ifMixup may hurt the model’s predictive accuracy. This is because shuffling the graph nodes before mixing graph pairs can significantly increase the input variety, compared to that of without shuffling. When we have a model with strong modeling capability, shuffling can help improve the model’s predictive accuracy because the shuffling can significantly increase both the training sample size and the input variety.




Figure 4: The average numbers of graph nodes (left) and edges (right) for training GCN and ifMixup.

	GIN Baseline	ifMixup	MixupGraph	DropEdge	DropNode	Attr. Masking	Rel. Impr.
PTC_MR	0.644±0.007	<b>0.672±0.005</b>	0.631±0.005	0.669±0.003	0.663±0.006	0.657± 0.004	4.35%
NCI109	0.820±0.002	<b>0.837±0.004</b>	0.822±0.008	0.792±0.002	0.796±0.002	0.822± 0.002	2.07%
NCI1	0.818±0.009	<b>0.839±0.004</b>	0.822±0.001	0.791±0.005	0.785±0.003	0.825± 0.002	2.57%
MUTAG	0.886±0.011	<b>0.890±0.006</b>	0.884±0.009	0.854±0.003	0.859±0.003	0.881± 0.004	0.45%
ENZYMES	0.526±0.014	0.543±0.005	0.521±0.007	0.488±0.015	0.528±0.002	<b>0.544± 0.013</b>	3.23%
PROTEINS	0.745±0.003	<b>0.754±0.002</b>	0.744±0.005	0.749±0.002	0.751±0.005	0.748 ±0.010	1.21%
IMDB-M	0.519±0.001	<b>0.532±0.001</b>	0.518±0.004	0.517±0.003	0.516±0.002	0.520± 0.004	2.50%
IMDB-B	0.762±0.004	<b>0.765±0.005</b>	0.761±0.001	0.762±0.005	0.764±0.006	0.761± 0.004	0.39%

Table 3: Accuracy with GIN. We report mean scores over 3 runs of 10-fold cross validation with standard deviations (denoted  $\pm$ ). The relative improvement of ifMixup over the baseline GIN is provided in the last column of the table. Best results are in Bold.

Shuffling Scenario	NCI109	NCI1
no shuffling	0.820±0.005	0.819±0.004
every 60 epochs	0.819±0.003	0.818±0.002
every 30 epochs	0.819±0.003	0.819±0.005
every 10 epochs	0.820±0.001	0.821±0.002
every 1 epoch	0.819±0.004	0.814±0.003

Table 4: Accuracy of ifMixup with various node shuffling scenarios (with increasing shuffling frequency) for mixing.

Layer	Shuf. Scenario	NCI109	NCI1
8	no shuffling	<b>0.820±0.005</b>	<b>0.819±0.004</b>
	shuf. per epoch	0.819±0.004	0.814±0.003
16	no shuffling	0.820±0.003	0.820±0.001
	shuf. per epoch	<b>0.821±0.003</b>	<b>0.822±0.004</b>

Table 5: Accuracy of ifMixup with and without node shuffling vs. network depth. Best results are in Bold.

**Impact of Graph Size** In this section, we evaluate the graph size change in training due to mixing up graph pairs by ifMixup and its impact. In Figure 4, we show the histograms of the number of graph nodes and edges for both the original graphs and mixed graphs on NCI109. From Figure 4, one can see that ifMixup *slightly* shifted the node and edge distributions to the right, but with very similar distributions<sup>3</sup>.

To further evaluate the impact of such change, we implement a variant of ifMixup, where the pair of graphs being mixed have the same number of nodes. We conduct experiments using 8-layer GCN on the NCI109 and NCI1 datasets, and present the results in Table 6. From Table 6, we can see that mixing random pair outperformed the three other pairing scenarios: mixing graphs with the same size with and without shuffling, and mixing with a shuffled version of the same graph. Interestingly, as shown in the second last two rows of the table, node shuffling seems to improve the performance when mixing random graphs with the same number of nodes.

### Results of Using GIN as Baseline

We also evaluate our method using the GIN (Xu et al. 2019) network architecture. The accuracy obtained by the GIN baseline, ifMixup, MixupGraph, DropEdge, DropNode, and Attribute Masking using GIN as baseline on the eight test datasets are presented in Table 3 (best results are in **Bold**).

Table 3 shows that, similar to the GCN case, the ifMixup with GIN as baseline outperformed all the five comparison

<sup>3</sup>This also suggests that ifMixup adds little additional computation cost to the original GNN models.

Mixing Pair	NCI109	NCI1
random pair	0.820±0.005	0.819±0.004
same size/no shuffling	0.787±0.003	0.795±0.002
same size/shuffling	0.816±0.002	0.812±0.004
self-paired/shuffling	0.806 ±0.003	0.807±0.003

Table 6: Accuracy of ifMixup with different pairing scenarios in terms of the number of nodes in a graph pair. Shuffling here refers to graph node order permutation.

models against all the eight datasets, except for Attribute Masking on the ENZYMES dataset. For example, when comparing with GIN, ifMixup obtained a relative accuracy improvement of 4.35%, 3.23%, and 2.57% on the PTC\_MR, ENZYMES, and NCI1 datasets, respectively. When comparing with the Mixup-like approach MixupGraph, ifMixup also obtained higher accuracy on all the eight datasets. For example, ifMixup was able to improve the accuracy over MixupGraph from 82.2%, 82.2%, 63.1%, and 51.8% to 83.7%, 83.9%, 67.2%, and 53.2% on the NCI109, NCI1, PTC\_MR, and IMDB-M datasets, respectively.

## Conclusion and Future Work

We proposed a very simple, straight-forward, input mixing strategy for Mixup on graph classification. Our method directly mixes up a pair of graph inputs, which is easy to be implemented. We also proved that such a simple mixing strategy processes an unique property under a mild assumption: the mixed graph can preserve all the information from its original graph pair used for mixing. We showed, using eight benchmark graph classification tasks from different domains, that our method obtained superior accuracy to popular graph augmentation approaches and graph classification methods.

It is worth noting that, ifMixup focuses on graph level classification. It becomes more challenging to guarantee information losses for mixed graphs on the node level. For example, it has been well observed that GNNs suffer from over-smoothing (Li, Han, and Wu 2018), where the representations of all nodes in a graph may converge to a subspace that makes their representations unrelated to the *input* (Li, Han, and Wu 2018; Wu et al. 2019). That is, for node classification, the learning dynamics of the GNNs play a vital role on the forming of the node embeddings. It is our interest to extend ifMixup for graph node classification in the future.

## Acknowledgements

This project was supported in part by a National Research Council of Canada (NRC) Collaborative R&D grant (AI4D-CORE-07).

## References

- Alon, U.; and Yahav, E. 2021. On the Bottleneck of Graph Neural Networks and its Practical Implications. In *International Conference on Learning Representations*.
- Brody, S.; Alon, U.; and Yahav, E. 2022. How Attentive are Graph Attention Networks? In *International Conference on Learning Representations*.
- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and Relieving the Over-Smoothing Problem for Graph Neural Networks from the Topological View. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 3438–3445. AAAI Press.
- Chen, J.; Ma, T.; and Xiao, C. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. *CoRR*.
- Ding, K.; Xu, Z.; Tong, H.; and Liu, H. 2022. Data Augmentation for Deep Graph Learning: A Survey. arXiv:2202.08235.
- Fey, M.; and Lenssen, J. E. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Fu, K.; Mao, T.; Wang, Y.; Lin, D.; Zhang, Y.; Zhan, J.; an Sun, X.; and Li, F. 2020. TS-Extractor: large graph exploration via subgraph extraction based on topological and semantic information. *Journal of Visualization*, 1 – 18.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural Message Passing for Quantum Chemistry. *CoRR*, abs/1704.01212.
- Guo, H. 2020. Nonlinear Mixup: Out-Of-Manifold Data Augmentation for Text Classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 4044–4051.
- Guo, H.; Mao, Y.; and Zhang, R. 2019a. Augmenting Data with Mixup for Sentence Classification: An Empirical Study. *CoRR*, abs/1905.08941.
- Guo, H.; Mao, Y.; and Zhang, R. 2019b. MixUp as Locally Linear Out-of-Manifold Regularization. In *AAAI2019*.
- Guo, H.; and Sun, S. 2022. SoftEdge: Regularizing Graph Classification with Random Soft Edges. arXiv:2204.10390.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. NIPS’17.
- Han, X.; Jiang, Z.; Liu, N.; and Hu, X. 2022. G-Mixup: Graph Data Augmentation for Graph Classification. arXiv:2202.07179.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020a. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *CoRR*.
- Hu, W.; Liu, B.; Gomes, J.; Zitnik, M.; Liang, P.; Pande, V.; and Leskovec, J. 2020b. Strategies for Pre-training Graph Neural Networks. arXiv:1905.12265.
- Huang, W.; Zhang, T.; Rong, Y.; and Huang, J. 2018. Adaptive Sampling towards Fast Graph Representation Learning. In *NeurIPS*.
- Jiang, B.; Kloster, K.; Gleich, D. F.; and Gribskov, M. 2017. AptRank: an adaptive PageRank model for protein function prediction on bi-relational graphs. *Bioinformatics*, 33(12): 1829–1836.
- Jindal, A.; Ghosh Chowdhury, A.; Didolkar, A.; Jin, D.; Sawhney, R.; and Shah, R. R. 2020. Augmenting NLP models using Latent Feature Interpolations. In *Proceedings of the 28th International Conference on Computational Linguistics*.
- Kim, J.-H.; Choo, W.; and Song, H. O. 2020. Puzzle Mix: Exploiting Saliency and Local Statistics for Optimal Mixup. In *International Conference on Machine Learning (ICML)*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.
- Li, G.; Müller, M.; Thabet, A.; and Ghanem, B. 2019. DeepGCNs: Can GCNs Go as Deep as CNNs? In *The IEEE International Conference on Computer Vision (ICCV)*.
- Li, G.; Xiong, C.; Thabet, A. K.; and Ghanem, B. 2020. DeeperGCN: All You Need to Train Deeper GCNs. *CoRR*, abs/2006.07739.
- Li, Q.; Han, Z.; and Wu, X. 2018. Deeper Insights Into Graph Convolutional Networks for Semi-Supervised Learning. In McIlraith, S. A.; and Weinberger, K. Q., eds., *AAAI*.
- Liu, G.; Zhao, T.; Xu, J.; Luo, T.; and Jiang, M. 2022. Graph Rationalization with Environment-based Augmentations. In *KDD*.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*.
- Park, J.; Shim, H.; and Yang, E. 2022. Graph Transplant: Node Saliency-Guided Graph Mixup with Local Structure Preservation. In *AAAI*, 7966–7974.
- Qiu, J.; Chen, Q.; Dong, Y.; Zhang, J.; Yang, H.; Ding, M.; Wang, K.; and Tang, J. 2020. GCC: Graph Contrastive Coding for Graph Neural Network Pre-Training, 1150–1160.
- Rong, Y.; Huang, W.; Xu, T.; and Huang, J. 2020. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *International Conference on Learning Representations*.
- Song, R.; Giunchiglia, F.; Zhao, K.; and Xu, H. 2021. Topological Regularization for Graph Neural Networks Augmentation. *CoRR*, abs/2104.02478.
- Thulasidasan, S.; Chennupati, G.; Bilmes, J. A.; Bhattacharya, T.; and Michalak, S. 2019. On Mixup Training: Improved Calibration and Predictive Uncertainty for Deep



- Neural Networks. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *NeurIPS*.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. *International Conference on Learning Representations*.
- Verma, V.; Lamb, A.; Beckham, C.; Courville, A.; Mitliagkas, I.; and Bengio, Y. 2018. Manifold Mixup: Encouraging Meaningful On-Manifold Interpolation as a Regularizer. *CoRR*.
- Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; and Hooi, B. 2020a. GraphCrop: Subgraph Cropping for Graph Classification. *CoRR*, abs/2009.10564.
- Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; and Hooi, B. 2021. Mixup for Node and Graph Classification. In *Proceedings of the Web Conference 2021*, 3663–3674.
- Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; Liu, J.; and Hooi, B. 2020b. NodeAug: Semi-Supervised Node Classification with Data Augmentation. In *KDD '20*, 207–217.
- Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Yu, P. S. 2019. A Comprehensive Survey on Graph Neural Networks. *CoRR*.
- Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; and Pande, V. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science*, 9(2): 513–530.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- Yoo, J.; Shim, S.; and Kang, U. 2022. Model-Agnostic Augmentation for Accurate Graph Classification. *arXiv*.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 5812–5823. Curran Associates, Inc.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding Deep Learning (Still) Requires Rethinking Generalization. *arXiv:1611.03530*.
- Zhang, G.; He, H.; and Katabi, D. 2019. Circuit-GNN: Graph Neural Networks for Distributed Circuit Design. In *International Conference on Machine Learning*.
- Zhang, H.; Cissé, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018a. mixup: Beyond Empirical Risk Minimization. In *ICLR2018*.
- Zhang, Y.; Pal, S.; Coates, M.; and Üstebay, D. 2018b. Bayesian graph convolutional neural networks for semi-supervised classification. *arXiv:1811.11103*.
- Zhao, T.; Liu, G.; Günneman, S.; and Jiang, M. 2022. Graph Data Augmentation for Graph Machine Learning: A Survey. *arXiv preprint arXiv:2202.08871*.
- Zhao, T.; Liu, Y.; Neves, L.; Woodford, O.; Jiang, M.; and Shah, N. 2021. Data Augmentation for Graph Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11015–11023.
- Zhou, J.; Shen, J.; and Xuan, Q. 2020. Data Augmentation for Graph Classification. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, 2341–2344.