

Graph Contrastive Learning with Personalized Augmentation

Xin Zhang*, Qiaoyu Tan†, Xiao Huang*, and Bo Li*

*Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR

†Department of Computer Science and Engineering, Texas A&M University, Texas, USA

xin12.zhang@connect.polyu.hk, qytan@tamu.edu,

xiaohuang@comp.polyu.edu.hk, comp-bo.li@polyu.edu.hk

Abstract—Graph contrastive learning (GCL) has emerged as an effective tool to learn representations for whole graphs in the absence of labels. The key idea is to maximize the agreement between two augmented views of each graph via data augmentation. Existing GCL models mainly focus on applying *identical augmentation strategies* for all graphs within a given scenario. However, real-world graphs are often not monomorphic but abstractions of diverse natures. Even within the same scenario (e.g., macromolecules and online communities), different graphs might need diverse augmentations to perform effective GCL. Thus, blindly augmenting all graphs without considering their individual characteristics may undermine the performance of GCL arts. However, it is non-trivial to achieve personalized allocation since the search space for all graphs is exponential to the number of graphs. To bridge the gap, we propose the first principled framework, termed as *Graph contrastive learning with Personalized Augmentation (GPA)*. It advances conventional GCL by allowing each graph to choose its own suitable augmentation operations. To cope with the huge search space, we design a tailored augmentation selector by converting the discrete space into continuous, which is a plug-and-play module and can be effectively trained with downstream GCL models end-to-end. Extensive experiments across 10 benchmark datasets from different types and domains demonstrate the superiority of GPA against state-of-the-art competitors. Moreover, by visualizing the learned augmentation distributions across different types of datasets, we show that GPA can effectively identify the most suitable augmentations for each graph based on its characteristics. The code is available at <https://github.com/qiaoyu-tan/GPA>.

Index Terms—Graph Contrastive Learning, Personalized Augmentation, Graph Neural Networks.

1 INTRODUCTION

GRAPH contrastive learning (GCL) could learn effective representations of graphs [1] in a self-supervised manner. It attracts considerable attention [2], [3], [4], [5], given that labels are not available in many real-world networked systems. The core idea of GCL is to generate two augmented views for each graph by perturbing it, and then learn representations via maximizing the mutual information between the two views [6], [7]. Existing efforts can be roughly divided into two categories, i.e., node-level [8], [9] and graph-level [10]. In this paper, we focus on the latter.

The performance of GCL is known to be heavily affected by the chosen augmentation types [1], [11], [12], since different augmentations may impose different inductive biases about the data. Intensive recent works have been devoted to exploring effective augmentations for different graph scenarios [8], [10], [13], [14]. Typical augmentation strategies include node dropping, edge perturbation, subgraph sampling, and attribute masking. The best augmentation option is often data-driven and varies in graph scales or types [11], [12]. For example, [11] reveals that edge perturbation may benefit social networks but hurt biochemical molecules. Therefore, manually searching augmentation strategies for a given scenario would involve extensive trials and efforts, hindering the practical usages of GCL. Several studies have been proposed to address this issue by automating GCL.

Despite the recent advances, existing GCL might be sub-optimal for augmentation configuration, since they apply a well-chosen but identical augmentation option to all graphs




Fig. 1. The effect of different augmentation strategies toward four randomly sampled graphs from MUTAG. X-axis denotes the id of augmentation pair. Y-axis is the graph id. The color represents the performance. The darker the color is, the better performance GCL achieves under the corresponding augmentation strategies.

in a dataset. The rationale is: graphs in a scenario usually have different properties because the characteristics of real-world graphs are complex and diverse [15], [16], [17]. For example, by slightly changing the structure of a molecular graph, its target function could be completely different [18]. Similar observations have also been found in many other graph domains, such as social communities and protein-protein interaction networks [19]. Motivated by these observations, a natural question arises: *What are the impacts of different augmentation strategies on a given graph (an instance in a graph dataset)? Can we build a stronger automated GCL by allowing each graph instance to choose its favorable augmentation types?*

To understand the first question, we conduct a pilot study on the MUTAG dataset to test how different augmen-

tation types impact the GCL results. Results in Fig. 1 show that different graphs favor distinct perturbation operations to achieve their best performance. For example, graph 1 performs better when using *edge perturbation* vs. *attribute masking*, while graph 2 prefers *original* vs. *node dropping*. Such personalized phenomenon of graph instances, in terms of their desirable augmentation strategies, has never been explored in GCL. To bridge the gap, in this paper, we propose to develop an effective augmentation selector to identify the most informative perturbation operators for different graphs when performing GCL on a specific dataset.

However, it is a non-trivial and challenging task to perform personalized augmentations in GCL because of two major reasons. First, unlike the traditional GCL setting, the search space in our personalized scenario grows exponentially with the number of graphs N in a dataset. This search space is intolerable in practice since N could be thousands or even tens of thousands [20]. As a result, the conventional trial-and-error approach cannot be directly applied because each trail itself (i.e., testing an augmentation option for a dataset) is time-consuming. Second, the augmentation choices and GCL model naturally depend on and reinforce each other since the contrastive loss consists of augmented views and GCL encoder. That is, learning a better GCL requires well-chosen augmentation strategies [11], while selecting suitable augmentation operators needs the signals of GCL as feedback [10]. Thus, how to perform the effective personalized augmentation selector on a premise of such mutual effect is another challenge.

To tackle the aforementioned challenges, we propose a novel contrastive learning framework, dubbed GPA. Specifically, we aim to investigate two research questions. 1) What are the impacts of different augmentation strategies on a given graph (an instance in a graph dataset)? 2) Can we build a stronger automated GCL by allowing each graph instance to choose its favorable augmentation types? GPA works by iteratively updating a personalized augmentation selector and the GCL method, where the former aims to identify optimal augmentation types for each graph instance, and the latter is trained according to an instance-level contrastive loss defined on those assigned augmentation options. Our **main contributions** are highlighted as follows.

- We focus on augmentation selection for graph contrastive learning, and propose an effective personalized augmentation framework (GPA). To the best of our knowledge, GPA is the first to assign personalized augmentation types to each graph based on its own characteristic when performing GCL. Moreover, we automate the personalized augmentation process, equipping GPA with broader applicability and practicability.
- To cope with the huge search space, we develop a personalized augmentation selector to effectively infer optimal augmentation strategies by relaxing the discrete space to be continuous. To exploit the mutual reinforced effect between augmentation strategies and GCL, GPA is formulated as a bi-level optimization, which learns the augmentation selector and GCL method systematically.

- We conduct extensive experiments to evaluate GPA on multiple graph benchmarks of diverse scales and types. Empirical results demonstrate the superiority of GPA against state-of-the-art GCL competitors.

2 RELATED WORK

There are three lines of research that are closely related to our work, namely, graph representation learning, conventional graph contrastive learning and automated graph contrastive learning. In this section, we briefly review some related works in these three fields. For comprehensive review, please refer to [21], [22] and [23].

2.1 Graph Representation Learning

With the rapid development of graph neural networks (GNNs), a large number of GNN-based graph representation learning frameworks have been proposed [23], [24], [25], [26], [27], [28], [29], which exhibit promising performance. Typically, these methods can be divided into supervised and unsupervised categories. While supervised methods [25], [30], [31], [32], [33], [34] achieve empirical success with the help of labels, reliable labels are often scarce in real-world scenarios. Thus, unsupervised graph learning approaches [35], [36], [37], [38] have broader application potential. For example, one of the well-known methods is GAE [35] which learns graph representations by reconstructing the network structure under the autoencoder approach. Another popular approach GraphSAGE [37] aims to train GNNs by a random-walk based objective.

2.2 Conventional Graph Contrastive Learning

Graph contrastive learning (GCL) has attracted significant attention in the past two years for self-supervised graph learning [39], [40]. It learns GNN encoder by maximizing the agreement between representations of a graph in its different augmented views, so that similar graphs are close to each other, while dissimilar ones are spaced apart. Many GCL efforts have been devoted to node level [7], [9], [41], subgraph level [42], [43], and graph level [11], [44] scenarios, where the key challenge lies in designing effective graph augmentations. Recently, GraphCL [11] introduced four types of graph augmentations, including node dropping, edge perturbation, sub-graph sampling, and node attribute masking, and showed that different graph applications may favor different augmentation combinations. However, the optimal augmentation configuration for a given dataset is mainly determined by either domain experts or extensive trial-and-errors, thus it limits the broader applications of GCL in practice. In parallel, some studies focus on other types of graphs, such as Knowledge Graph [45], [46].

2.3 Automated Graph Contrastive Learning

Recently, some automated GCL methods have been proposed to seek the optimal augmentation without repetitive trials. These methods can be roughly categorized into two types.

One is **GCL with view generation**, which targets to train a view generator to predict the edge or node probabilities

of the original graph and then generate augmented views by sampling edges or nodes based on the possibilities [47]. Two SOTA models in this type are AD-GCL [48] and AutoGCL [49]. Despite its effectiveness, this type of work is limited to simple augmentation strategies, i.e., edge dropping or node masking, and further loses its applicability to complex graphs and scenarios. In our experiments, we include AD-GCL and AutoGCL as representatives of this type since their empirical results are generally better than the work in [47].

The other type, **GCL with augmentation selection**, avoids the limitation of applicability by adopting predefined augmentation strategies. Specifically, the predefined augmentation strategies include simple ones such as node dropping and edge masking and complicated ones like sampling subgraph and motif-based augmentation. Thus, instead of training view generators, this category of work focuses on selecting optimal strategies. A typical model in this category is JOAO [10]. It automatically selects the ideal augmentation strategies for each dataset via assigning each augmentation strategy a projection head and learning the importance of each head.

2.4 Comparison of GPA and JOAO

While JOAO and GPA all focus on training GraphCL in an automated fashion, they differ in two crucial ways: 1) **The research problem is different.** In general, JOAO's primary goal is to find the optimal augmentation strategies for a dataset, focusing on a dataset-level perspective. In contrast, GPA addresses the challenge at the graph level, aiming to identify the most suitable augmentation strategies for individual graphs within a dataset. This distinction allows GPA to achieve greater granularity and adaptability, as it tailors augmentation strategies according to the specific characteristic of each graph instance. By focusing on individual graphs, GPA achieves a level of customization that is pivotal in enhancing the performance of GraphCL models in diverse graph datasets. Specifically, JOAO learns a shared sampling distribution for a given dataset, but GPA learns a sampling vector for each instance (i.e., graph) in the dataset. Notably, due to randomness, JOAO can sample different augmentations to different graphs to some extent. However, since the sampling distribution is fixed, it is "fake" personalization because it does not consider the characteristics of different graphs. In contrast, GPA could explicitly generate the personalized sampling vector to each graph, which is tailored and more challenging as the augmentation space grows exponentially to the dataset size. 2) **The sampler optimization process is different.** JOAO adopts a sampling-based approach to directly train a K -dimensional sampling distribution parameters. However, such a solution cannot be applied to our personalized scenarios because i) the total trainable parameters will be linear to the size of the dataset, i.e., $K \times |\mathcal{V}|$, where \mathcal{V} is the dataset; and ii) it is restricted to transductive settings and cannot generate distribution vectors for new graphs, i.e., unseen graphs during training. To this end, we adopt a new approach to achieve personalized augmentation allocation by parameterizing the sampler as a neural network taking the topology structure and node attributes of a graph as input. It reduces the

trainable parameters of the sampler to be independent of the size of the dataset.

3 PROBLEM STATEMENT

Notation. Let $\mathcal{G} = \{\mathcal{G}_n : 1 \leq n \leq N\}$ denote a graph set with N sample graphs, where $\mathcal{G}_n = (\mathcal{V}, \mathcal{E}) \in \mathcal{G}$ stands for an undirected graph with nodes \mathcal{V} and edges \mathcal{E} . Each node $v \in \mathcal{V}$ in \mathcal{G}_n is described by an F -dimensional feature vector $\mathbf{X}_v \in \mathbb{R}^F$. We use $\mathcal{A} = \{A_k : 1 \leq k \leq K\}$ to denote a set of data augmentation operators, where K is the maximum number of augmentation types of interest. Each augmentation operator $A_k : \mathcal{G}_n \rightarrow \tilde{\mathcal{G}}_n$ transforms a graph into its conceptually similar form with certain prior. In previous GCL studies, they focus on identifying two optimal augmentation types for the whole dataset \mathcal{G} , such as the augmentation pair $A_{i,j} = (A_i, A_j)$, where $A_i, A_j \in \mathcal{A}$. The optimal augmentation pair here is often manually picked via rules of thumb or trial-and-error. However, as shown in Fig. 1, different graphs within the dataset may favor different augmentation combinations. Therefore, we study personalized augmentation selection, i.e., to identify the optimal augmentation strategies for each graph, and formally define the research problem as below.

Definition 1. Personalized augmentation selection. Given a set of graphs $\mathcal{G} = (\mathcal{G}_n : 1 \leq n \leq N)$, and the augmentation space $\mathcal{A} = \{A_k : 1 \leq k \leq K\}$ consisting of K different augmentation types, to perform GCL, personalized augmentation selection aims to find the optimal augmentation pair $A_{i,j}^n = (A_i^n, A_j^n)$ for each graph $\mathcal{G}_n \in \mathcal{G}$. The values of i and j are determined by the characteristic of the n -th sample graph \mathcal{G}_n .

4 METHODOLOGY

In this section, we present the details of the proposed GPA shown in Fig. 2. In a nutshell, it contains two critical components: the *personalized augmentation selector* and the *GCL model*. The former module aims to infer augmentation choices for the downstream GCL methods when training them on the training set, while the latter provides reward to update the augmentation selector based on the validation set. In the following, we first illustrate the exponential selection space of our personalized augmentation setting. Then, we elaborate the details of the augmentation selector and the GCL method. Finally, we show how to jointly optimize the two components in a unified perspective.

4.1 Personalized Augmentation Selection Space

Given the graph dataset $\mathcal{G} = \{\mathcal{G}_n : 1 \leq n \leq N\}$ and a pool of augmentation operators $\mathcal{A} = \{A_1, A_2, \dots, A_K\}$, existing GCL efforts aim to select two informative operators (e.g., $(A_i, A_j \mid 1 \leq i, j \leq K)$) for N graphs to create their augmented views. Since the augmentation operators are shared for the whole dataset, the total selection space is $\binom{K+1}{2}$, i.e., sampling two operators with replacement. This collective selection strategy is widely adopted in existing GCL works. However, as discussed before, various sample graphs may favor different augmentation operators owing to the diversity of graph-structured data. Therefore,

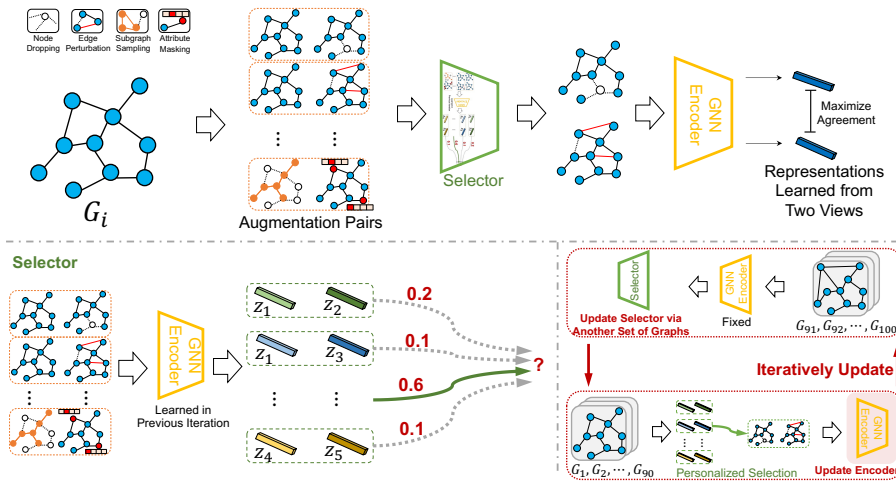


Fig. 2. Illustration of our GPA framework. The *personalized augmentation selector* infers the two most informative augmentation operators, and the *GCL model* trains GNN encoder based on the sampled augmented views. Specifically, the *personalized augmentation selector* is learned to adjust its selection strategy to infer optimal augmentations on each graph, according to the characteristic of each graph and the *GCL model*'s performance, i.e., loss.

we propose to adaptively choose two augmentation strategies for different graphs. Following Definition 1, we define the selected augmentations for each graph \mathcal{G}_n as $A_{i,j}^n = (A_i^n, A_j^n)$. Then, the potential augmentation selection size for each graph is $\binom{K+1}{2}$, and the total selection space for the whole dataset equals to $\binom{K+1}{2}^N$. Although K is empirically small (e.g., $K = 5$) in GCL domain, the total selection space in our personalized setting is still huge and intractable, since the complexity grows exponentially to the number of graphs. For instance, when $K = 5$ and $N = 100$, we already have 15^{100} selection configurations roughly. The situation is more serious in real-world scenarios where N is thousands or even tens of thousands. In this paper, we adopt five essential augmentation operators (i.e. $K = 5$), denoted by $\mathcal{A} = \{\text{Identical}, \text{NodeDrop}, \text{EdgePert}, \text{Subgraph}, \text{AttMask}\}$. These augmentations are initially proposed by the pioneering work [11], and have been demonstrated to be effective for contrastive learning [10]. The augmentation pairs are shown in Table 1, and the details of these augmentations are listed below.

- **NodeDrop.** Given the graph \mathcal{G}_n , NodeDrop randomly discards a fraction of the vertices and their connections. The dropping probability of each node follows the i.i.d. uniform distribution. The underlying assumption is that missing part of vertices does not damage the semantic information of \mathcal{G}_n .
- **EdgePert.** The connectivity in \mathcal{G}_n is perturbed through randomly adding or dropping a certain portion of edges. We also follow the i.i.d. uniform distribution to add/drop each edge. The underlying prior is that the semantic meaning of \mathcal{G}_n is robust to the variance of edges.
- **AttMask.** AttMask masks the attributes of a certain proportion of vertices. Similarly, each node's masking possibility follows the i.i.d. uniform distribution. Attribute masking implies that the absence of some vertex attributes does not affect the semantics of \mathcal{G}_n .

TABLE 1
Distinct graph augmentation pairs.

$A_{i,j}$	A_i, A_j	Augmentation	Augmentation
$A_{1,1}$	A_1, A_1	Identical	Identical
$A_{1,2}$	A_1, A_2	Identical	Node Dropping
$A_{1,3}$	A_1, A_3	Identical	Edge Perturbation
$A_{1,4}$	A_1, A_4	Identical	Subgraph
$A_{1,5}$	A_1, A_5	Identical	Attribute Masking
$A_{2,2}$	A_2, A_2	Node Dropping	Node Dropping
$A_{2,3}$	A_2, A_3	Node Dropping	Edge Perturbation
$A_{2,4}$	A_2, A_4	Node Dropping	Subgraph
$A_{2,5}$	A_2, A_5	Node Dropping	Attribute Masking
$A_{3,3}$	A_3, A_3	Edge Perturbation	Edge Perturbation
$A_{3,4}$	A_3, A_4	Edge Perturbation	Subgraph
$A_{3,5}$	A_3, A_5	Edge Perturbation	Attribute Masking
$A_{4,4}$	A_4, A_4	Subgraph	Subgraph
$A_{4,5}$	A_4, A_5	Subgraph	Attribute Masking
$A_{5,5}$	A_5, A_5	Attribute Masking	Attribute Masking

- **Subgraph.** This augmentation method samples a subgraph from the given graph \mathcal{G}_n based on random walk. It believes that most of the semantic meaning of \mathcal{G}_n can be preserved in its local structure.

In summary, by considering personalized augmentation, the search space per dataset increases from $\binom{K+1}{2}$ to $\binom{K+1}{2}^N$. Therefore, common selection techniques such as rules of thumb or trial-and-errors adopted by prior GCL approaches [10], [11] are no longer appropriate. Thus, a tailored augmentation selector is needed to effectively tackle the challenging personalized augmentation problem.

4.2 Personalized Augmentation Selector

In order to assign different augmentation operators to various sample graphs when performing GCL on a specific dataset, random selection is the intuitive solution. Its key idea is to randomly sample two augmentation types for each graph from the candidate set. Despite the simplicity, the random selection approach fails to control the quality of sam-

pled augmentation operators. Therefore, directly coupling the existing GCL framework with random augmentation selection would lead to performance degradation (shown as the variant: GPA-random in Sec. 5.4).

To address this issue, we focus on data-driven search by making the augmentation selection process learnable. The principle idea is to parameterize our personalized augmentation selector with a deep neural network, which takes a query graph as input and outputs its optimal augmentation choices. There are two main hurdles to achieve this goal: (i) given the exponential augmentation space, how can we make our personalized augmentation selector scale to the real-world dataset with thousands of graphs; (ii) since the topology structure and node attributes are crucial to graph-structured data, how can our personalized augmentation selector exploit this information to produce a more precise augmentation choice? We illustrate our dedicated solutions below.

Given the augmentation pool $\mathcal{A} = \{A_i : 1 \leq i \leq K\}$ and a query graph \mathcal{G}_n , our augmentation selector is required to select the most two informative augmentations, e.g., (A_1^n, A_3^n) , from the candidate set. This selection problem is well-known to be discrete and non-differentiable. Although enormous efforts based on evolution or reinforcement learning have been proposed to address the discrete selection problem, they are still not suitable for such a large selection space (illustrated in Sec. 4.1) and are far from utilizing the properties of graphs. Therefore, we propose to make the search space learnable by relaxing the discrete selection space to be continuous inspired by [50] and further make this relaxation consider the characteristic of graphs. Specifically, given the sampled augmentation pair $A_{i,j}^n = (A_i^n, A_j^n)$ of \mathcal{G}_n , its importance score $\hat{\alpha}_{i,j}^n$ is computed as

$$\hat{\alpha}_{i,j}^n = \frac{\exp(\alpha_{i,j}^n)}{\sum_{i',j'} \exp(\alpha_{i',j'}^n)} \quad (1)$$

$$\alpha_{i,j}^n = g_\theta(f_w(A_i^n(\mathcal{G}_n)) \parallel f_w(A_j^n(\mathcal{G}_n))),$$

where g_θ denotes the score function that takes the representations of augmented views $A_i^n(\mathcal{G}_n)$ and $A_j^n(\mathcal{G}_n)$ as input. This design enforces the score estimation to take into account the topology and node attributes of \mathcal{G}_n . In practice, g_θ is parameterized as a two-layer MLP with a ReLU activation function. f_w is the GNN encoder for graph representation learning. \parallel indicates the concatenation operation. Through Eq. (1), the discrete augmentation selection process reduces to learning a score function g_θ under the consideration of graph characteristic.

After the personalized selector is well-trained, let $\alpha^n = [\hat{\alpha}_{1,1}^n, \dots, \hat{\alpha}_{i,j}^n, \dots, \hat{\alpha}_{K,K}^n]$ denote the vector of importance scores associated with all different augmentation pairs of the graph \mathcal{G}_n . The optimal augmentation choice of \mathcal{G}_n can be obtained by selecting the augmentation pair with the maximum score in α^n . For example, $A^n = (A_i^n, A_j^n)$ if $\hat{\alpha}_{i,j}^n = \arg \max_{i',j'} \hat{\alpha}_{i',j'}^n$.

To summarize, the above equation provides a principled solution to our personalized augmentation setting. On one hand, it allows augmentation selection in such a large search space via the simple forward propagation of a shallow neural network, i.e., g_θ . On the other hand, it can also infer the most informative augmentations for each graph

based on its own characteristic for downstream GCL model training (discussed in Sec. 4.3).

4.3 GCL Model Learning

After the personalized augmentation selector is plugged in, we can adopt it to train the GCL model. Notice that our model is applicable to arbitrary GCL methods that rely on two augmented views as input. In this section, we mainly focus on the most popular and generic GCL architecture - GraphCL [11] as the backbone and leave other specific architectures for future work.

Assume (A_i^n, A_j^n) is the optimal augmentation pair of graph \mathcal{G}_n identified by our personalized augmentation selector, and $f_w(\cdot)$ is the GNN encoder. GraphCL proposes to learn $f_w(\cdot)$ by maximizing the agreement between the two augmented views, i.e., $A_i^n(\mathcal{G}_n), A_j^n(\mathcal{G}_n)$. The GNN encoder can encode the whole graph into a hidden space \mathbb{R}^D . In practice, a shared projection head function $p : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is often applied upon the output of GNN encoder to improve the model capacity. In the following sections, we abuse the notation f_w to denote both the GNN encoding function and the projection function. Based on this notation, we can formally calculate the instance-level contrastive loss as follows:

$$\mathcal{L}(\mathcal{G}_n) = -\log \left(\exp(\text{sim}(f_w(A_i^n(\mathcal{G}_n)), f_w(A_j^n(\mathcal{G}_n)))) / \tau \right) / \sum_{n'=1, n' \neq n}^N \exp(\text{sim}(f_w(A_i^n(\mathcal{G}_n)), f_w(A_j^{n'}(\mathcal{G}_{n'})))) / \tau \quad (2)$$

where $\text{sim}(\cdot, \cdot)$ denotes the cosine similarity function and τ is the temperature parameter. By minimizing Eq. (2), it encourages the two augmented views of the same sample graph to have similar representations, while enforces the augmented representations of disparate graphs to be highly distinct. As the sum operation over all graphs \mathcal{G} in the denominator of Eq. (2) is computationally prohibitive, GCL is often trained under minibatch sampling [11], where the negative views are generated from the augmented graphs within the same minibatch.

Although Eq. (2) looks similar to the traditional contrastive loss, the key difference between them is that the augmentation operators A_i^n, A_j^n are closely related to the sample graph \mathcal{G}_n in our formulation. That is, the augmentation operators learned by our model vary from one graph to another according to their own characteristics, which are previously enforced to be the same regardless of the graph's diverse nature. Benefiting from considering the graph's personality, the GCL method can learn the basic but essential features of graphs and thus achieve more expressive representations for downstream tasks, as empirically verified in Sec. 5.

4.4 Model Optimization

Until now, we have illustrated the detailed personalized augmentation selector as well as the downstream GCL framework, the remaining question is how to effectively train the two modules. The naive solution is to first train the personalized augmentation selector separately, and then

optimize the GCL method using the identified personalized augmentations as input. However, such a task-agnostic approach is sub-optimal, since it does not take the mutual reinforced effect between augmentation and the GCL method into consideration. This is because learning a better GCL method requires optimal augmentation strategies since they can augment more personalized features to distinguish it from other objects. Meanwhile, obtaining suitable augmentation strategies also needs the signals of a better GCL method as guidance for optimization. As a result, without linking the two modules in a principled way, it is almost infeasible to enforce the *personalized augmentation selector* to accurately infer optimal augmentation strategy for improving the performance of the *GCL method*. To this end, we propose to tackle this problem by jointly training the two modules under a bi-level optimization, expressed as:

$$\begin{aligned} & \arg \min_{\theta} \mathcal{L}_{valid}(w^*(\theta), \theta) \\ \text{s.t. } & w^*(\theta) = \arg \min_w \mathcal{L}_{train}(w, \theta), \end{aligned} \quad (3)$$

where θ denotes the trainable parameters of the personalized augmentation selector, and w is the parameters for the GCL method. The upper-level objective $\mathcal{L}_{valid}(w^*(\theta), \theta)$ aims to find θ that minimizes the validation rewards on the validation set given the optimal w^* , and the lower-level objective $\mathcal{L}_{train}(w, \theta^*)$ targets to optimize w by minimizing the contrastive loss based on the training set with θ fixed. We want to remark that GPA only exploits the signals from the self-supervisory task itself without accessing labels. Thus, compared with conventional supervised methods, the validation set here only contains a set of graphs without label information, which is much easier to construct, e.g., randomly sampling 10% of the training set.

By optimizing Eq. (3), the personalized augmentation selector and the target GCL model will be jointly trained to reinforce their reciprocal effects. Since deriving exact solutions for this bi-level problem is indeed analytically intractable, we adopt the alternating gradient descent algorithm to solve it as follows.

4.4.1 Lower-level optimization.

With θ fixed, we can update w with the standard gradient descent procedure as below.

$$w' = w - \xi \nabla_w \mathcal{L}_{train}(w, \theta), \quad (4)$$

where $\mathcal{L}_{train} = \mathbb{E}_{\mathcal{G}_{train}} \mathcal{L}(\mathcal{G}_n)$ with $\mathcal{L}(\mathcal{G}_n)$ is instance-level contrastive loss defined in Eq. (2), \mathcal{G}_{train} denotes the training set, and ξ is the learning rate.

4.4.2 Upper-level optimization.

Since it is not intuitive to directly calculate the gradient *w.r.t.* θ over all augmentation options, we first define the upper-level objective based on Eq. (2) as below:

$$\mathcal{L}_{valid}(w^*(\theta), \theta) = \sum_{A_i, A_j \in \mathcal{A}} \sum_{\mathcal{G}_n \in \mathcal{G}_{valid}} \hat{\alpha}_{i,j}^n \mathcal{L}(\mathcal{G}_n), \quad (5)$$

where $\hat{\alpha}_{i,j}^n$ is the selecting score computed by the score function g_{θ} with parameter θ defined in Eq. (1). \mathcal{G}_{valid}

denotes the validation set. Based on the above loss function, we can update θ by fixing w , expressed as:

$$\theta' = \theta - \xi \nabla_{\theta} \mathcal{L}_{valid}(w^*(\theta), \theta). \quad (6)$$

However, evaluating the gradient *w.r.t.* θ exactly is intractable and computationally expensive, since it requires solving for the optimal $w^*(\theta)$ whenever θ gets updated. In other words, whenever θ gets updated, the low-level objective $\mathcal{L}_{train}(w, \theta)$ needs to be trained until convergence in order to obtain the optimal $w^*(\theta)$ and further calculate the gradient $\nabla_{\theta} \mathcal{L}_{valid}(w^*(\theta), \theta)$, which is computationally intensive. Therefore, to approximate the optimal solution $w^*(\theta)$, we propose to take one step of gradient descent update for w , without solving the lower-level optimization completely by training until convergence. To further compute the gradient of θ , we apply chain rule to differentiate $\mathcal{L}_{train}(w'(\theta), \theta)$ with respect to θ via w' , where w' is defined in Eq. (4). Therefore, the gradient of θ can be approximated as:

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{valid}(w^*(\theta), \theta) & \approx \nabla_{\theta} \mathcal{L}_{valid}(w', \theta) \\ & \approx \nabla_{\theta} \mathcal{L}_{valid}(w - \xi \nabla_w \mathcal{L}_{train}(w, \theta), \theta). \end{aligned} \quad (7)$$

Applying chain rule to Eq. (7), we further obtain the gradient of θ as:

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{valid}(w^*(\theta), \theta) & \approx \nabla_{\theta} \mathcal{L}_{valid}(w', \theta) \\ & \approx \frac{\partial \mathcal{L}_{valid}(w', \theta)}{\partial \theta} \frac{\partial \theta}{\partial \theta} + \frac{\partial \mathcal{L}_{valid}(w', \theta)}{\partial w'} \frac{\partial w'}{\partial \theta} \\ & \approx \nabla_{\theta} \mathcal{L}_{valid}(w', \theta) \\ & \quad - \nabla_{w'} \mathcal{L}_{valid}(w', \theta) \xi \nabla_{\theta, w}^2 \mathcal{L}_{train}(w, \theta). \end{aligned} \quad (8)$$

Since the computation cost of the second term in Eq. (8) is still high, it can be further approximated by the finite difference method:

$$\begin{aligned} & \xi \nabla_{\theta, w}^2 \mathcal{L}_{train}(w, \theta) \nabla_{w'} \mathcal{L}_{valid}(w', \theta) \\ & \approx \frac{\nabla_{\theta} \mathcal{L}_{train}(w^+, \theta) - \nabla_{\theta} \mathcal{L}_{train}(w^-, \theta)}{2\epsilon} \end{aligned} \quad (9)$$

where $w^{\pm} = w \pm \epsilon \nabla_{w'} \mathcal{L}_{valid}(w', \theta)$ and ϵ denotes a small scalar.

By alternating the update rules in Eq. (4) and Eq. (6), we are able to progressively learn the two modules. Although an optimizer with the theoretical guarantee of convergence for the bi-level problem in Eq. (3) remains an open challenge, alternating gradient descent algorithm has been widely adopted to solve similar objectives in Bayesian optimization [51], automatic differentiation [52], and adversarial training [53]. The complete optimization procedure of GPA is shown in the Algorithm 1. We also show some level of empirical convergence in Fig. 3.

4.5 More Discussions on GPA

In addition to augmentation selection, another line of research focuses on generating views via a trainable generator, such as AutoGCL [49] and AD-GCL [48]. Specifically, the augmented view is obtained by sampling the original one based on the edge or node probabilities predicted by the generator. As such, these methods could limit their applications to simple augmentation strategies (e.g., edge dropping and node masking) and cannot be adopted for

Algorithm 1: The framework of GPA

Input: A graph dataset \mathcal{G} , the personalized augmentation selector $g_\theta(\cdot)$, and a GCL model $f_w(\cdot)$;
Output: The well-trained GCL model;

- 1 Split the input graph dataset \mathcal{G} into train \mathcal{G}_{train} and validation \mathcal{G}_{valid} set;
- 2 Initialize the selector parameter θ and the GCL model parameter w ;
- 3 **while not converge do**
- 4 Randomly sample a minibatch of graphs from the training set;
- 5 Infer the optimal augmentation pairs for sampled graphs using the *personalized augmentation selector* $g_\theta(\cdot)$;
- 6 Update parameters w of the GCL model based on the sampled graphs and the identified augmentation types according to Eq. (4);
- 7 Randomly sample a batch of graphs from validation set;
- 8 Compute the rewards based on the sampled validation graphs using the updated w' according to Eq. (5);
- 9 Update parameters θ of the *personalized augmentation selector* according to Eq. (6) and Eq. (9);
- 10 **Return** The well-trained GCL model.

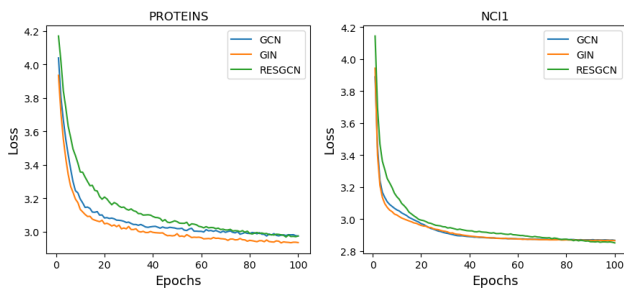


Fig. 3. Empirical training curves of approximate gradient scheme in GPA on datasets PROTEINS and NCI1 with different GNN encoders.

applications that require complex augmentation strategies, such as subgraph sampling and motif-based augmentations (e.g., graphon [54]). In contrast, as an augmentation selection method, GPA has broad applicability, i.e., it is flexible to incorporate new augmentation strategies created by domain experts that might be useful by modifying the augmentation pool without extra effort. We summarize the difference in Table 2.

4.5.1 Complexity analysis of GPA

We analyze the complexity of GPA and illustrate its impact on the GCL backbone –GraphCL [11]. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and the GNN encoder f_w . The time complexity for the GNN backbones used in common graph learning tasks is $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}|)$. Since the proposed GPA method is built upon GraphCL [11] and the personalized augmentation selector g_θ , the additional cost is mainly caused by the selector.

For GraphCL, it performs two encoder computations per iteration plus a prediction head. If we assume the backward cost is similar to that of the forward pass, the complexity of GraphCL is $4C_{\text{encoder}}(|\mathcal{V}| + |\mathcal{E}|) + 2C_{\text{head}}(N_{\text{batch}}) + C_{\text{loss}} + 2C_{\text{aug}}$, where C are constants which rely on the neural architectures, and N_{batch} is the batch size. For the selector g_θ , it takes K encoder computations per iteration to generate graph representations and simple multilayer perceptrons (MLPs) to estimate the sampling probability. Thus its complexity is $KC_{\text{encoder}}(|\mathcal{V}| + |\mathcal{E}|) + C_{\text{selector}}(N_{\text{batch}}) + KC_{\text{aug}}$. Therefore, the total complexity for GPA and GraphCL are $(2K+4)C_{\text{encoder}}(|\mathcal{V}| + |\mathcal{E}|) + (2K+2)C_{\text{aug}}$ and $4C_{\text{encoder}}(|\mathcal{V}| + |\mathcal{E}|) + 2C_{\text{aug}}$, respectively, since the costs of GNN encoder and data augmentation are significantly higher than others. Assume the time cost for data augmentation is 5 times of GNN encoder, i.e., $C_{\text{aug}} = 5C_{\text{encoder}}$, then the time complexity of GPA is $\frac{2K+4+(2K+2)*5}{4+2*5} = \frac{12K+14}{14}$ times that of GraphCL. Given that $K = 5$ in our experiments, the complexity of GPA is around 6 times of GraphCL in theory. However, due to the inefficient implementation of our current code (e.g., for loop in the loss function and perform K augmentations in series), the empirical complexity of GPA is less than 7 times of GraphCL. It is noteworthy that parallel techniques such as multiprocessing could be used to accelerate the training.

5 EXPERIMENTS

We evaluate the performance of GPA on multiple graph datasets with various scales and types. We focus on exploring the following research questions.

- **Q1:** How effective is GPA in performing graph representation learning against state-of-the-art GCL methods in unsupervised and semi-supervised evaluation tasks?
- **Q2:** How effective is the proposed personalized augmentation selector in identifying augmentations across various datasets?
- **Q3:** Compared with random selection, how effective is our proposed personalized augmentation selector?
- **Q4:** What are the impacts of hyperparameters on GPA, such as the embedding dimension d of the score function?

5.1 Datasets and Experiment Settings

Datasets. For a comprehensive comparison, we evaluate the performance of GPA on ten widely used benchmark datasets. Specifically, we include two small molecules networks (NCI1 and MUTAG), two bioinformatics networks (DD and PROTEINS), and five social networks (COLLAB, REDDIT-BINARY, REDDIT-MULTI-5K, IMDB-BINARY, and GITHUB) from TUDatasets [18]. To evaluate the scalability of our model, we also use one large-scale OGB [20] dataset **ogbg-molhiv**. The data statistics are summarized in Table 5.

Learning protocols. Following common protocols [2], [11], we aim to evaluate the performance of GPA in unsupervised and semi-supervised settings. In our model training phase, we randomly split 10% of graphs in each dataset into

TABLE 2

An overview of graph contrastive learning methods. Our GPA model is the first automated GCL effort that supports personalized augmentation in terms of various augmentation types.

	Auto	Simple Augmentation	Complex Augmentation	Applicability	Personalized
GraphCL	-	✓	✓	✓	-
JOAO	✓	✓	✓	✓	-
AD-GCL	✓	✓	-	-	✓
AutoGCL	✓	✓	-	-	✓
Ours	✓	✓	✓	✓	✓

TABLE 3

Unsupervised learning performance for graph classification in TUDatasets (Averaged accuracy \pm std. over 10 runs). The **bold** numbers denote the best performance and the numbers with underline represent the second best performance

Dataset	NCI1	PROTEINS	DD	MUTAG	COLLAB	RDT-B	RDT-M5K	IMDB-B	Avg.Rank
InfoGraph	76.20 \pm 1.06	74.44 \pm 0.31	72.85 \pm 1.78	<u>89.01 \pm 1.13</u>	70.65 \pm 1.13	82.50 \pm 1.42	53.46 \pm 1.03	73.03 \pm 0.87	4.75
GraphCL	77.87 \pm 0.41	74.39 \pm 0.45	<u>78.62 \pm 0.40</u>	86.80 \pm 1.34	71.36 \pm 1.15	89.53 \pm 0.84	55.99 \pm 0.28	71.14 \pm 0.44	3.88
JOAO	78.07 \pm 0.47	74.55 \pm 0.41	<u>77.32 \pm 0.54</u>	87.35 \pm 1.02	69.50 \pm 0.36	85.29 \pm 1.35	55.74 \pm 0.63	70.21 \pm 3.08	5.00
JOAOv2	78.36 \pm 0.53	74.07 \pm 1.10	77.40 \pm 1.15	87.67 \pm 0.79	69.33 \pm 0.34	86.42 \pm 1.45	<u>56.03 \pm 0.27</u>	70.83 \pm 0.25	4.50
AD-GCL	69.67 \pm 0.51	73.59 \pm 0.65	74.49 \pm 0.52	88.62 \pm 1.27	<u>73.32 \pm 0.61</u>	85.52 \pm 0.79	53.00 \pm 0.82	71.57 \pm 1.01	5.13
AutoGCL	82.00 \pm 0.29	<u>75.80 \pm 0.36</u>	77.57 \pm 0.60	88.64 \pm 1.08	70.12 \pm 0.68	88.58 \pm 1.49	56.75 \pm 0.18	<u>73.30 \pm 0.40</u>	2.38
GPA	<u>80.42 \pm 0.41</u>	75.94 \pm 0.25	79.90 \pm 0.35	89.68 \pm 0.80	76.17 \pm 0.10	<u>89.32 \pm 0.38</u>	53.70 \pm 0.19	74.64 \pm 0.35	2.38

TABLE 4

Semi-supervised learning performance for graph classification in TUDatasets (Averaged accuracy \pm std. over 10 runs). The **bold** numbers denote the best performance and the numbers with underline represent the second best performance

Dataset	NCI1	PROTEINS	DD	RDT-B	RDT-M5K	GITHUB	ogbg-molhiv	Avg.Rank
GAE	74.36 \pm 0.24	70.51 \pm 0.17	74.54 \pm 0.68	87.69 \pm 0.40	53.58 \pm 0.13	63.89 \pm 0.52	-	6.67
InfoGraph	74.86 \pm 0.26	72.27 \pm 0.40	75.78 \pm 0.34	88.66 \pm 0.95	53.61 \pm 0.31	65.21 \pm 0.88	-	4.50
GraphCL	74.63 \pm 0.25	74.17 \pm 0.34	76.17 \pm 1.37	<u>89.11 \pm 0.19</u>	52.55 \pm 0.45	65.81 \pm 0.79	55.48 \pm 1.32	4.29
JOAO	74.48 \pm 0.25	72.13 \pm 0.92	75.69 \pm 0.67	88.14 \pm 0.25	52.83 \pm 0.54	65.00 \pm 0.30	56.83 \pm 1.39	5.71
JOAOv2	74.86 \pm 0.39	73.31 \pm 0.48	75.81 \pm 0.73	88.79 \pm 0.65	52.71 \pm 0.28	66.60 \pm 0.60	<u>57.39 \pm 1.39</u>	4.00
AD-GCL	75.18 \pm 0.31	73.96 \pm 0.47	77.91 \pm 0.73	90.10 \pm 0.15	53.49 \pm 0.28	67.13 \pm 0.52	-	2.33
AutoGCL	73.75 \pm 2.25	75.65 \pm 2.40	<u>77.50 \pm 4.41</u>	79.80 \pm 3.47	49.91 \pm 2.70	62.46 \pm 1.51	-	5.83
GPA	75.50 \pm 0.14	<u>74.27 \pm 1.11</u>	76.68 \pm 0.81	89.99 \pm 0.32	54.92 \pm 0.35	68.31 \pm 0.13	60.76 \pm 1.01	1.57

TABLE 5
Statistics of the datasets.

	$ G $	Avg.Nodes	Avg.Edges	#Label
NCI1	4,110	29.87	32.30	2
PROTEINS	1,113	39.06	72.82	2
DD	1,178	284.32	715.66	2
MUTAG	188	17.93	19.79	2
COLLAB	5,000	74.49	2,457.78	3
IMDB-BINARY	1,000	19.77	96.53	2
REDDIT-BINARY	2,000	429.63	497.75	2
REDDIT-MULTI-5K	4,999	508.52	594.87	5
GITHUB	12,725	113.79	234.64	2
ogbg-molhiv	41,127	25.5	27.5	2

the validation set and use the remaining for training. After the model is trained, in unsupervised setting, we train an SVM classifier on the graph representations generated by the trained GCL model, and apply 10-fold cross-validation to evaluate the performance. For semi-supervised setting, we finetune the GCL model (its GNN encoder) with a logistic regression layer for semi-supervised learning, where

the labeled sample ratio is 0.1. To avoid randomness, we repeat the process for ten times and report the averaged results.

Implementation details. Our model is built upon Pytorch and PyG (PyTorch Geometric) library [55]. We train our model with Adam optimizer using a fixed batch size of 128. Similar to GraphCL [11], the default augmentation ratio is set to 0.2 for all augmentation types. In unsupervised setting, we adopt a three-layer GIN [56] encoder with hidden dimension 128 for all datasets. In semi-supervised scenarios, we employ a five-layer ResGCN [57] encoder with dimension 128 for TUDatasets [18], while a five-layer GIN [56] encoder with dimension 300 for OGB datasets as suggested in [20]. There is one hyper-parameter in our model, i.e., the hidden dimension d of score function g_θ . We search d within the set $\{128, 256, 512\}$. The impact of the hidden dimension is analyzed in Sec. 5.5.

Baseline methods. To validate the effectiveness of GPA, we compare against three categories of state-of-the-art competitors. First, to evaluate the effectiveness of contrastive learning, we include one traditional network embedding

method GAE [35]. Second, to study why we need personalized augmentation, we include classic GCL methods that assign identical augmentation strategies for all graphs InfoGraph [2], GraphCL [11], JOAO [10] and its variant JOAOv2 [10]. Note that JOAO and JOAOv2 are two sample-based automated GCL methods that focus on selecting the most suitable predefined augmentation strategies for each dataset. They are the most relevant baselines to our proposed GPA. Third, we also include two view-generated automated GCL methods: AD-GCL [48] and AutoGCL [49]. Ordinarily, they should not be considered as our baselines since GPA belongs to GCL with augmentation selection, not view generation. We add them to comprehensively show the power of GPA.

5.2 Comparison with Baselines

We start by comparing the performance of GPA with the state-of-the-art baseline methods under two settings (Q1). Table 3 & Table 4 report the results of all methods on diverse datasets under unsupervised and semi-supervised settings, respectively. We have the following Observations.

Obs.1. With personalized augmentations for each graph, GPA outperforms vanilla GCL methods with fixed augmentation per dataset. By identifying different augmentations for different sample graphs, GPA performs generally better than vanilla GCL methods on two evaluation scenarios (Table 3 and Table 4). Specifically, in the semi-supervised evaluation task, GPA consistently outperforms GAE, InfoGraph, GraphCL, JOAO, and JOAOv2 on all datasets. In the unsupervised setting, GPA outperforms the vanilla GCL methods on 6 out of 8 datasets. In particular, GPA improves 7.8%, 9.9%, 9.6%, and 6.7% over InfoGraph, JOAOv2, JOAO, and GraphCL on COLLAB in Table 3, respectively. This observation validates the effectiveness of performing personalized augmentation in GCL training.

Obs.2. Across diverse datasets, GPA performs better than (or on par with) the view-generated GCL methods on two evaluation settings. On all datasets originating from diverse domains, GPA generally performs better or sometimes on par with the state-of-the-art AD-GCL and AutoGCL methods, as shown in the average rank. In unsupervised setting, GPA outperforms AD-GCL on all datasets while GPA beats AutoGCL on 6 out of 8 datasets. In semi-supervised setting, GPA outperforms both AD-GCL and AutoGCL on NCI1, RDT-B, RDT-M5K, and GITHUB datasets and achieves the second best on PROTEINS.

Obs.3. GPA scales well on large datasets. To study the scalability of our model, we further conduct experiment on the large-scale OGB dataset: ogbg-molhiv. View-generated GCL methods are excluded for this dataset since they are not officially tested on OGB datasets. From the Table 4, GPA consistently performs better than GraphCL and JOAO. To be specific, GPA improves 9.5% and 5.9% over GraphCL and JOAOv2 on ogbg-molhiv, respectively.

5.3 Personalized Augmentation Analysis

To study the effectiveness of our model in identifying informative augmentation types for various graphs (Q2), we visualize the learned augmentation distribution on Fig. 4.

TABLE 6
Ablation study of GPA under unsupervised setting in terms of mean classification accuracy

	GPA-random	GPA
NCI1	77.71 ± 0.60	80.42 ± 0.41
PROTEINS	74.21 ± 0.39	75.94 ± 0.25
DD	77.25 ± 0.69	79.90 ± 0.35
MUTAG	86.08 ± 2.93	89.68 ± 0.80
IMDB-B	71.80 ± 1.13	74.64 ± 0.35

By comparing across different types of datasets, we observe the following.

Obs.4. By learning from the data, GPA can effectively assign different augmentations for various datasets. Our model GPA can identify different augmentations for different sample graphs, and allow different datasets to have their own augmentation distributions (see Fig. 4). Specifically, on MUTAG, 19 graphs prefer (*Identical, NodeDrop*) augmentations, while 30 graphs favor (*Subgraph, Subgraph*) augmentation combinations. Notice that (*Subgraph, Subgraph*) will generate two different subgraph-perturbation-induced augmented views, owing to sample randomness. Besides, COLLAB more likes the (*AttMask, AttMask*) augmentation pair, while DD prefers (*EdgePert, EdgePert*) operations. These observations empirically echo the necessity of performing personalized augmentation for GCL methods.

Another promising observation is that our model can assign (*Identical, Identical*) choice (i.e., two identical views) to some portion of graphs over all datasets. Given that the mutual information between two identical views (i.e., representations) is always maximized, such pure identical augmentations can be regarded as a skip operation. That is, these graphs abandon themselves during the GCL model training. The possible reason is that the existing augmentation strategies are not suitable to capture their characteristics or damage their semantic meanings. In this case, blindly selecting any combination of other augmentation types may incur huge performance degradation or noise. This observation sheds light on designing more advanced augmentation strategies beyond the current basic augmentations. On the other hand, it verifies the effectiveness of the proposed augmentation selector in skipping noisy graphs during model training by providing identical augmentations.

5.4 Ablation Study

To further investigate the effectiveness of the proposed personalized augmentation selector (Q3), we compare it with a random-search based variant, i.e., GPA-random. GPA-random replaces the personalized augmentation selector with a random mechanism. Specifically, it assigns one random augmentation pair to each graph. Noticed that GPA-random still assigns different augmentations to each graph while GraphCL assigns one pre-defined pair of augmentation strategies to the whole dataset. Table 6 shows the results in terms of unsupervised setting. From the table, we can observe that GPA consistently performs better than GPA-random in all cases. In particular, GPA improves 3.5%, 2.3%, 3.4%, 4.2%, and 4.0% over GPA-random on NCI1,

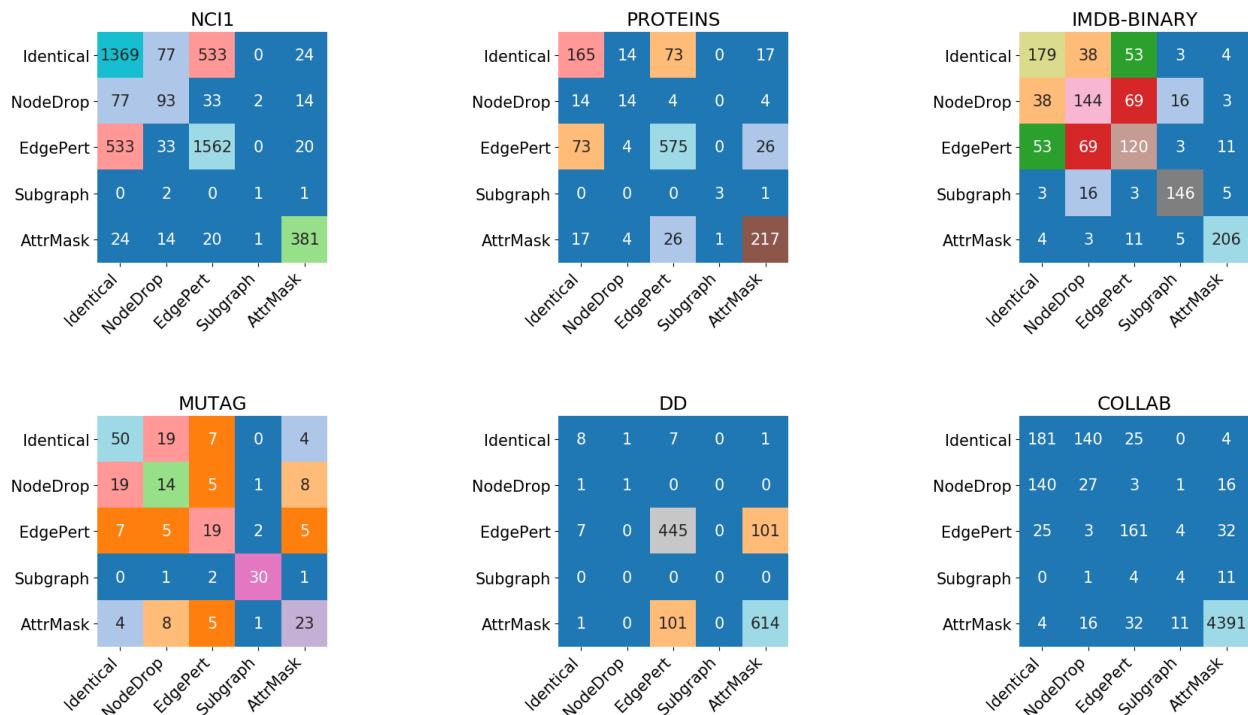


Fig. 4. Augmentation distribution learned by GPA over molecules, bioinformatics, and social networks, in terms of the unsupervised setting.




Fig. 5. Personalized augmentation selector dimension analysis of GPA

PROTEINS, DD, MUTAG, and IMDB-B, respectively. This comparison validates our motivation to develop a tailored and learnable personalized augmentation selector for GCL methods.

5.5 Parameter Sensitivity Analysis

We now study the impact of the parameter, i.e., the hidden dimension d of the score function to answer Q4. Specifically, we search d from the set $\{64, 128, 256, 512\}$ and plot the results of GPA on three representative datasets in Fig. 5. Similar observations are obtained by other datasets. From the figure, we can see that GPA generally performs stably across various dimension choices.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we study the augmentation selection problem for graph contrastive learning. Existing GCL efforts mainly

focus on employing two shared augmentation strategies for all graphs in the dataset based on the assumption that they contain similar nature. Here, we argue that such collective augmentation selection is suboptimal in practice due to the heterogeneity of graph structure data. Different graphs should have different augmentation preferences. To bridge the gap, we propose a novel graph contrastive learning framework with personalized augmentation termed as GPA. GPA not only allows each graph to select its optimal augmentation types, but also automates the selection via a personalized augmentation selector, which can be jointly trained with downstream GCL models under a bi-level optimization. Empirical results on the graph classification task demonstrate the superiority of GPA against state-of-the-art GCL methods in terms of unsupervised and semi-supervised settings, across multiple benchmark datasets with various types such as molecules, bioinformatics, and social networks.

In the future, we will explore how to extend our personalized augmentation idea to the node-level contrastive learning task. Moreover, it would be interesting to investigate whether our model works for other data types, such as image or text data.

REFERENCES

- [1] K. Hassani and A. H. Khasahmadi, "Contrastive multi-view representation learning on graphs," in *International Conference on Machine Learning*. PMLR, 2020, pp. 4116–4126.
- [2] F.-Y. Sun, J. Hoffmann, V. Verma, and J. Tang, "Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization," *International Conference on Learning Representations*, 2020.
- [3] Q. Tan, N. Liu, X. Huang, R. Chen, S.-H. Choi, and X. Hu, "Mgae: Masked autoencoders for self-supervised learning on graphs," *arXiv preprint arXiv:2201.02534*, 2022.

- [4] B. Chen, J. Zhang, X. Zhang, Y. Dong, J. Song, P. Zhang, K. Xu, E. Kharlamov, and J. Tang, "Gccad: Graph contrastive learning for anomaly detection," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [5] X. Xu, F. Zhou, K. Zhang, and S. Liu, "Ccgl: Contrastive cascade graph learning," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [6] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax." *International Conference on Learning Representations (Poster)*, vol. 2, no. 3, p. 4, 2019.
- [7] Z. Peng, W. Huang, M. Luo, Q. Zheng, Y. Rong, T. Xu, and J. Huang, "Graph representation learning via graphical mutual information maximization," in *International World Wide Web Conference*, 2020, pp. 259–270.
- [8] Y. Zhu, Y. Xu, F. Yu, Q. Liu, S. Wu, and L. Wang, "Graph contrastive learning with adaptive augmentation," in *International World Wide Web Conference*, 2021, pp. 2069–2080.
- [9] S. Wan, S. Pan, J. Yang, and C. Gong, "Contrastive and generative graph convolutional networks for graph-based semi-supervised learning," *AAAI Conference on Artificial Intelligence*, 2020.
- [10] Y. You, T. Chen, Y. Shen, and Z. Wang, "Graph contrastive learning automated," *International Conference on Machine Learning*, 2021.
- [11] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5812–5823, 2020.
- [12] W. Jin, T. Derr, H. Liu, Y. Wang, S. Wang, Z. Liu, and J. Tang, "Self-supervised learning on graphs: Deep insights and new direction," *arXiv preprint arXiv:2006.10141*, 2020.
- [13] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax," *International Conference on Learning Representations*, 2019.
- [14] D. Chen, Y. Lin, L. Li, X. R. Li, J. Zhou, X. Sun *et al.*, "Distance-wise graph contrastive learning," *International Joint Conference on Artificial Intelligence*, 2022.
- [15] D. B. West *et al.*, *Introduction to graph theory*. Prentice hall Upper Saddle River, 2001, vol. 2.
- [16] N. Liu, Q. Tan, Y. Li, H. Yang, J. Zhou, and X. Hu, "Is a single vector enough? exploring node polysemy for network embedding," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 932–940.
- [17] J. Huang, X. Huang, and J. Xu, "Truss-based structural diversity search in large graphs," *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [18] C. Morris, N. M. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann, "Tudataset: A collection of benchmark datasets for learning with graphs," in *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*, 2020. [Online]. Available: www.graphlearning.io
- [19] S. Thakoor, C. Tallec, M. G. Azar, R. Munos, P. Veličković, and M. Valko, "Bootstrapped representation learning on graphs," in *International Conference on Learning Representations 2021 Workshop on Geometrical and Topological Representation Learning*, 2021.
- [20] 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," *Advances in Neural Information Processing Systems*, vol. 33, pp. 22 118–22 133, 2020.
- [21] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, J. Zhang, and J. Tang, "Self-supervised learning: Generative or contrastive," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [22] Y. Liu, M. Jin, S. Pan, C. Zhou, Y. Zheng, F. Xia, and P. Yu, "Graph self-supervised learning: A survey," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [23] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [24] L. Wang, W. Yu, W. Wang, W. Cheng, W. Zhang, H. Zha, X. He, and H. Chen, "Learning robust representations with graph denoising policy network," in *IEEE International Conference on Data Mining*. IEEE, 2019, pp. 1378–1383.
- [25] R. G. Iyer, W. Wang, and Y. Sun, "Bi-level attention graph neural networks," in *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2021, pp. 1126–1131.
- [26] H. Wang, J. Zhang, Q. Zhu, and W. Huang, "Augmentation-free graph contrastive learning," *arXiv preprint arXiv:2204.04874*, 2022.
- [27] Q. Tan, N. Liu, and X. Hu, "Deep representation learning for social network analysis," *Frontiers in big Data*, vol. 2, p. 2, 2019.
- [28] H. Li, X. Wang, Z. Zhang, and W. Zhu, "Ood-gnn: Out-of-distribution generalized graph neural network," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [29] X. Fan, M. Gong, Y. Wu, A. Qin, and Y. Xie, "Propagation enhanced neural message passing for graph representation learning," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [30] J. Chen, T. Ma, and C. Xiao, "Fastgcn: Fast learning with graph convolutional networks via importance sampling," in *International Conference on Learning Representations*, 2018.
- [31] M. Ding, J. Tang, and J. Zhang, "Semi-supervised learning on graphs with generative adversarial nets," in *The Conference on Information and Knowledge Management*, 2018, pp. 913–922.
- [32] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations*, 2017.
- [33] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *International Conference on Learning Representations*, 2018.
- [34] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu, "Heterogeneous graph attention network," in *The world wide web conference*, 2019, pp. 2022–2032.
- [35] T. N. Kipf and M. Welling, "Variational graph auto-encoders," *Advances in Neural Information Processing Systems*, 2016.
- [36] A. Garcia Duran and M. Niepert, "Learning graph representations with embedding propagation," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [37] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [38] Z. Hou, X. Liu, Y. Dong, C. Wang, J. Tang *et al.*, "Graph-mae: Self-supervised masked graph autoencoders," *arXiv preprint arXiv:2205.10803*, 2022.
- [39] L. Wu, H. Lin, C. Tan, Z. Gao, and S. Z. Li, "Self-supervised learning on graphs: Contrastive, generative, or predictive," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [40] H. Li, Z. Zhang, X. Wang, and W. Zhu, "Disentangled graph contrastive learning with independence promotion," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [41] Z. Hu, Y. Dong, K. Wang, K.-W. Chang, and Y. Sun, "Gpt-gnn: Generative pre-training of graph neural networks," in *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1857–1867.
- [42] J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang, "Gcc: Graph contrastive coding for graph neural network pre-training," in *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1150–1160.
- [43] Y. Jiao, Y. Xiong, J. Zhang, Y. Zhang, T. Zhang, and Y. Zhu, "Sub-graph contrast for scalable self-supervised graph representation learning," in *IEEE International Conference on Data Mining*. IEEE, 2020, pp. 222–231.
- [44] J. Zeng and P. Xie, "Contrastive self-supervised learning for graph classification," in *AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 10 824–10 832.
- [45] Z. Pan and P. Wang, "Hyperbolic hierarchy-aware knowledge graph embedding for link prediction," in *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2021, pp. 2941–2948.
- [46] R. G. Iyer, Y. Bai, W. Wang, and Y. Sun, "Dual-geometric space embedding model for two-view knowledge graphs," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 676–686.
- [47] Y. You, T. Chen, Z. Wang, and Y. Shen, "Bringing your own view: Graph contrastive learning without prefabricated data augmentations," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 2022, pp. 1300–1309.
- [48] S. Suresh, P. Li, C. Hao, and J. Neville, "Adversarial graph augmentation to improve graph contrastive learning," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [49] Y. Yin, Q. Wang, S. Huang, H. Xiong, and X. Zhang, "Autogcl: Automated graph contrastive learning via learnable view generators," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 8, 2022, pp. 8892–8900.
- [50] H. Liu, K. Simonyan, and Y. Yang, "Darts: Differentiable architecture search," in *International Conference on Learning Representations*, 2018.
- [51] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in Neural Information Processing Systems*, vol. 25, 2012.

- [52] A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots, "Truncated back-propagation for bilevel optimization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1723–1732.
- [53] J. Wang, T. Zhang, S. Liu, P.-Y. Chen, J. Xu, M. Fardad, and B. Li, "Towards a unified min-max framework for adversarial exploration and robustness," *arXiv preprint arXiv:1906.03563*, 2019.
- [54] L. Ruiz, L. Chamon, and A. Ribeiro, "Graphon neural networks and the transferability of graph neural networks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1702–1712, 2020.
- [55] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in *International Conference on Learning Representations Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [56] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *International Conference on Learning Representations*, 2018.
- [57] Y. S. Ting Chen, Song Bian, "Are powerful graph neural nets necessary? a dissection on graph classification," *CoRR*, vol. abs/1905.04579, 2019.



Bo Li is an assistant professor in the Department of Computing at the Hong Kong Polytechnic University. Formerly, he was a Postdoctoral Fellow in the Department of Computer Science at the University of Oxford and in the Department of Electrical and Computer Engineering at the University of Texas at Austin. He received his Ph.D. in Computer Science from Stony Brook University. He is broadly interested in algorithms, AI, and computational economics, including problems related to resource allocation, game theory, online algorithms, and their applications to Blockchain and machine learning.



Xin Zhang received MSc degree in information technology from The Hong Kong Polytechnic University, and the BS degree in economic statistics from Shanghai University of Finance and Economics. Her research interests include graph machine learning, self-supervised learning, and data mining.



Qiaoyu Tan received the BEng degree from the College of Computer Science and Technology, Southwest University, China, in 2017. He is currently working toward the PhD degree in the Department of Computer Science and Engineering, Texas A&M University, College station, Texas. His research interests include graph machine learning, self-supervised learning, and recommendation systems.



Xiao Huang is an assistant professor in the Department of Computing at The Hong Kong Polytechnic University. He received Ph.D. in Computer Engineering from Texas A&M University in 2020, M.S. in Electrical Engineering from Illinois Institute of Technology in 2015, and B.S. in Engineering from Shanghai Jiao Tong University in 2012. He is a program committee member of NeurIPS 2021-2022, KDD 2019-2022, WWW 2022, ICLR 2022, AAI 2021-2022, ICML 2021-2022, IJCAI 2020-2022, CIKM 2019-2022,

WSDM 2021-2022, SDM 2022, ICKG 2020-2021. His research interests include graph neural networks, knowledge graphs, network anomaly detection, and recommender systems.