

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

3-2024

### Iterative graph self-distillation

Hanlin ZHANG

Shuai LIN

Weiyang LIU

Pan ZHOU

Singapore Management University, panzhou@smu.edu.sg

Jian TANG

*See next page for additional authors*

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Graphics and Human Computer Interfaces Commons](#)

---

#### Citation

ZHANG, Hanlin; LIN, Shuai; LIU, Weiyang; ZHOU, Pan; TANG, Jian; LIANG, Xiaodan; and XING, Eric. Iterative graph self-distillation. (2024). *IEEE Transactions on Knowledge and Data Engineering*. 36, (3), 1161-1169. Available at: [https://ink.library.smu.edu.sg/sis\\_research/8992](https://ink.library.smu.edu.sg/sis_research/8992)

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

---

**Author**

Hanlin ZHANG, Shuai LIN, Weiyang LIU, Pan ZHOU, Jian TANG, Xiaodan LIANG, and Eric XING

# Iterative Graph Self-distillation

Hanlin Zhang<sup>1</sup>, Shuai Lin<sup>2</sup>, Weiyang Liu<sup>3,4</sup>, Pan Zhou<sup>5</sup>, Jian Tang<sup>6,7</sup>, Xiaodan Liang<sup>2</sup>, Eric P. Xing<sup>1</sup>  
<sup>1</sup>Carnegie Mellon University, <sup>2</sup>Sun Yat-sen University, <sup>3</sup>University of Cambridge, <sup>4</sup>Max Planck Institute for Intelligent Systems, <sup>5</sup>Sea AI Lab <sup>6</sup>Mila-Québec AI Institute, <sup>7</sup>HEC Montréal

## ABSTRACT

Recently, there has been increasing interest in the challenge of how to discriminatively vectorize graphs. To address this, we propose a method called Iterative Graph Self-Distillation (IGSD) which learns graph-level representation in an unsupervised manner through instance discrimination using a self-supervised contrastive learning approach. IGSD involves a teacher-student distillation process that uses graph diffusion augmentations and constructs the teacher model using an exponential moving average of the student model. The intuition behind IGSD is to predict the teacher network representation of the graph pairs under different augmented views. As a natural extension, we also apply IGSD to semi-supervised scenarios by jointly regularizing the network with both supervised and self-supervised contrastive loss. Finally, we show that finetuning the IGSD-trained models with self-training can further improve the graph representation power. Empirically, we achieve significant and consistent performance gain on various graph datasets in both unsupervised and semi-supervised settings, which well validates the superiority of IGSD.

## KEYWORDS

graph representation learning, self-supervised learning, contrastive learning

### ACM Reference Format:

Hanlin Zhang<sup>1</sup>, Shuai Lin<sup>2</sup>, Weiyang Liu<sup>3,4</sup>, Pan Zhou<sup>5</sup>, Jian Tang<sup>6,7</sup>, Xiaodan Liang<sup>2</sup>, Eric P. Xing<sup>1</sup>. 2021. Iterative Graph Self-distillation. In . ACM, New York, NY, USA, 9 pages.

## 1 INTRODUCTION

Graphs are ubiquitous representations encoding relational structures across various domains. Learning low-dimensional vector representations of graphs is critical in various domains ranging from social science [35] to bioinformatics [11, 58]. Many graph neural networks (GNNs) [14, 27, 52] have been proposed to learn node and graph representations by aggregating information from every node’s neighbors via non-linear transformation and aggregation functions. However, the key limitation of existing GNN architectures is that they often require a huge amount of labeled data to be competitive but annotating graphs like drug-target interaction networks is challenging since it needs domain-specific expertise.

Therefore, unsupervised learning on graphs has been long studied, such as graph kernels [43] and matrix-factorization approaches [2].

Inspired by the recent success of self-supervised representation learning in various domains like images [8, 20] and texts [39], most related works in the graph domain either follow the paradigm of self-supervised pretraining on the pretext tasks or contrastive learning via *InfoMax principle* [22]. The former often needs meticulous designs of hand-crafted tasks [23, 55] while the latter is dominant in self-supervised graph representation learning, which trains encoders to maximize the mutual information (MI) between the representations of the global graph and local patches (such as subgraphs) [18, 45, 48]. However, context-instance contrastive approaches usually need to sample subgraphs as local views to contrast with global graphs. And they usually require an additional discriminator for scoring local-global pairs and negative samples, which can be computationally prohibitive [47]. Besides, the performance is also very sensitive to the choice of encoders and MI estimators [47]. Moreover, context-instance contrastive approaches cannot be handily extended to the semi-supervised setting since local subgraphs lack labels that can be utilized for training. Therefore, we are seeking an approach that learns the entire graph representation by contrasting the whole graph directly to address the above challenges.

Motivated by recent progress on instance discrimination contrastive learning [16, 51], we propose the Iterative Graph Self-Distillation (IGSD), a teacher-student framework to learn graph representations by contrasting graph instances directly. The high-level idea of IGSD is based on graph contrastive learning where we pull similar graphs together and push dissimilar graph away. However, the performance of conventional contrastive learning largely depends on how negative samples are selected. To learn discriminative representations and avoid collapsing to trivial solutions, a large set of negative samples [8, 20] or a special mining strategy [20, 42] is necessary. In order to alleviate the dependency on negative sample mining and still be able to learn discriminative graph representations, we propose to use self-distillation as a strong regularization to guide the graph representation learning.

In the IGSD framework, graph instances are augmented as several views to be encoded and projected into a latent space where we define a similarity metric for consistency-based training. The parameters of the teacher network are iteratively updated as an exponential moving average of the student network parameters [16, 20, 46], allowing the knowledge transfer between them. As a merely small amount of labeled data is often available in many real-world applications, we further extend IGSD to the semi-supervised setting such that it can effectively utilize graph-level labels while considering arbitrary amounts of positive pairs belonging to the same class. Moreover, to leverage the information from pseudo-labels with high confidence, we develop a self-training algorithm based on the supervised contrastive loss for fine-tuning.

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

Ljubljana '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 Association for Computing Machinery.

We experiment with real-world datasets in various scales and compare the performance of IGSD with state-of-the-art graph representation learning methods. Experimental results show that IGSD achieves competitive performance in both self-supervised and semi-supervised settings with different encoders and data augmentation choices. With the help of self-training, our performance can exceed state-of-the-art baselines by a large margin.

In summary, we make the following contributions in this paper:

- We propose a self-distillation framework called IGSD for self-supervised graph-level representation learning where the teacher-student distillation is performed for contrasting graph pairs under different augmented views using graph diffusion.
- We further extend IGSD to the semi-supervised scenarios, where the labeled data are utilized effectively with the supervised contrastive loss and self-training.
- We empirically show that IGSD surpasses state-of-the-art methods in semi-supervised graph classification and molecular property prediction tasks and achieves performance competitive with state-of-the-art approaches in self-supervised graph classification tasks.

## 2 RELATED WORK

**Graph Representation Learning** Traditionally, graph kernels are widely used for learning node and graph representations. This common process includes meticulous designs like decomposing graphs into substructures and using kernel functions like Weisfeiler-Leman graph kernel [43] to measure graph similarity between them. However, they usually require non-trivial hand-crafted substructures and domain-specific kernel functions to measure the similarity while yields inferior performance on downstream tasks like node classification and graph classification. Moreover, they often suffer from poor scalability [5] and great memory consumption [29] due to some procedures like path extraction and recursive subgraph construction. Recently, there has been increasing interest in Graph Neural Network (GNN) approaches for graph representation learning and many GNN variants have been proposed [27, 40, 52]. However, they mainly focus on supervised settings.

**Contrastive Learning** Modern unsupervised learning in the form of contrastive learning can be categorized into two types: context-instance contrast and context-context contrast [31]. The context-instance contrast, or so-called global-local contrast focuses on modeling the belonging relationship between the local feature of a sample and its global context representation. Most unsupervised learning models on graphs like DGI [48], InfoGraph [45], CMC-Graph [19] fall into this category, following the *InfoMax principle* to maximize the the mutual information (MI) between the input and its representation. However, estimating MI is notoriously hard in context-instance contrastive learning and in practice tractable lower bound on this quantity is maximized instead. And maximizing tighter bounds on MI can result in worse representations without stronger inductive biases in sampling strategies, encoder architecture and parametrization of MI estimators [47]. Besides, the intricacies of negative sampling in context-instance contrastive approaches impose key research challenges like improper amount of negative samples or biased negative sampling [10, 47]. Another

line of contrastive learning approaches called instance discrimination (or context-context contrast) directly study the relationships between the global representations of different samples as what metric learning does. For instance, our method is inspired by BYOL [15], a method bootstraps the representations of the whole images directly. And Graph Barlow Twins [4] utilizes a cross-correlation-based loss instead of negative samples for learning representations. Focusing on global representations between samples and corresponding augmented views also allows instance-level supervision to be incorporated naturally like introducing supervised contrastive loss [26] into the framework for learning powerful representations. Graph Contrastive Coding (GCC) [38] is a pioneer to leverage instance discrimination as the pretext task for structural information pre-training. GraphCL [54] aims at learning generalizable, transferable, and robust representations of graph data in an unsupervised manner. However, our work is fundamentally different from theirs. GCC aims to find common and transferable structural patterns across different graph datasets and the contrastive scheme is done through subgraph instance discrimination. GraphCL focuses on the impact of various combinations of graph augmentations on multiple datasets and studies unsupervised setting only. On the contrary, our model aims at learning graph-level representation in both unsupervised and semi-supervised settings by directly contrasting graph instances such that data augmentation strategies and graph labels can be utilized naturally and effectively.

**Knowledge Distillation** Knowledge distillation [21] is a method for transferring knowledge from one architecture to another, allowing model compression and inductive biases transfer. Self-distillation [12] is a special case when two architectures are identical, which can iteratively modify regularization and reduce over-fitting if perform suitable rounds [33]. However, they often focus on closing the gap between the predictive results of student and teacher rather than defining similarity loss in latent space for contrastive learning.

**Semi-supervised Learning** Modern semi-supervised learning can be categorized into two kinds: multi-task learning and consistency training between two separate networks. Most widely used semi-supervised learning methods take the form of multi-task learning:  $\arg \min_{\theta} \mathcal{L}_l(D_l, \theta) + w \mathcal{L}_u(D_u, \theta)$  on labeled data  $D_l$  and unlabeled data  $D_u$ . By regularizing the learning process with unlabeled data, the decision boundary becomes more plausible. Another mainstream of semi-supervised learning lies in introducing student network and teacher network and enforcing consistency between them [30, 32, 46]. [1] studies the empirical effects of consistency regularization, which validates the motivation of the teacher-student model: introducing a slow-moving average teacher network to measures consistency against a student one, thus providing a consistency-based training paradigm where two networks can be mutually improved. It has been shown that semi-supervised learning performance can be greatly improved via unsupervised pre-training of a (big) model, supervised fine-tuning on a few labeled examples, and distillation with unlabeled examples for refining and transferring the task-specific knowledge [9]. However, whether task-agnostic self-distillation would benefit semi-supervised learning is still underexplored.

**Data augmentation** Data augmentation strategies on graphs are limited since defining views of graphs is a non-trivial task. There are two common choices of augmentations on graphs (1)

feature-space augmentation and (2) structure-space augmentation. A straightforward way is to corrupt the adjacency matrix which preserves the features but adds or removes edges from the adjacency matrix with some probability distribution [48]. [57] improves performance in GNN-based semi-supervised node classification via edge prediction. Empirical results show that the diffusion matrix can serve as a denoising filter to augment graph data for improving graph representation learning significantly both in supervised [28] and unsupervised settings [18]. [18] shows the benefits of treating diffusion matrix as an augmented view of multi-view contrastive graph representation learning.

### 3 PRELIMINARIES

#### 3.1 Formulation

**Self-supervised/Unsupervised Graph Representation Learning** Given a set of unlabeled graphs  $\mathcal{G} = \{G_i\}_{i=1}^N$ , we aim at learning the low-dimensional representation of every graph  $G_i \in \mathcal{G}$  favorable for downstream tasks like graph classification.

**Semi-supervised Graph Representation Learning** Consider a whole dataset  $\mathcal{G} = \mathcal{G}_L \cup \mathcal{G}_U$  composed by labeled data  $\mathcal{G}_L = \{(G_i, y_i)\}_{i=1}^l$  and unlabeled data  $\mathcal{G}_U = \{G_i\}_{i=l+1}^{l+u}$  (usually  $u \gg l$ ), our goal is to learn a model that can make predictions on graph labels for unseen graphs. And with  $K$  augmentations, we get  $\mathcal{G}'_L = \{(G'_k, y'_k)\}_{k=1}^{Kl}$  and  $\mathcal{G}'_U = \{G'_k\}_{k=l+1}^{K(l+u)}$  as our training data.

#### 3.2 Graph Representation Learning

We represent a graph instance as  $G(\mathcal{V}, \mathcal{E})$  with the node set  $\mathcal{V}$  and the edge set  $\mathcal{E}$ . The dominant ways of graph representation learning are graph neural networks with neural message passing mechanisms [17]: for every node  $v \in \mathcal{V}$ , node representation  $\mathbf{h}_v^k$  is iteratively computed from the features of their neighbor nodes  $\mathcal{N}(v)$  using a differentiable aggregation function. Specifically, at the iteration  $k$  we get the node embedding as Eq. (1):

$$\mathbf{h}_v^k = \sigma(\mathbf{W}^k \cdot \text{CONCAT}(\mathbf{h}_v^{k-1}, \text{AGGREGATE}_k(\{\mathbf{h}_u^{k-1}, \forall u \in \mathcal{N}(v)\}))) \quad (1)$$

Then the graph-level representations can be attained by aggregating all node representations using a readout function like summation or set2set pooling [50].

#### 3.3 Graph Data Augmentation

It has been shown that the learning performance of GNNs can be improved via graph diffusion, which serves as a homophily-based denoising filter on both features and edges in real graphs [28]. The transformed graphs can also serve as effective augmented views in contrastive learning [18]. Inspired by that, we transform a graph  $G$  with transition matrix  $T$  via graph diffusion and sparsification  $\mathbf{S} = \sum_{k=0}^{\infty} \theta_k T^k$  into a new graph with adjacency matrix  $\mathbf{S}$  as an augmented view in our framework. While there are many design choices in coefficients  $\theta_k$  like heat kernel, we employ Personalized PageRank (PPR) with  $\theta_k^{\text{PPR}} = \alpha(1-\alpha)^k$  due to its superior empirical performance [18]. As another augmentation choice, we randomly remove edges of graphs to attain corrupted graphs as augmented views to validate the robustness of models to different augmentation choices.

### 4 ITERATIVE GRAPH SELF-DISTILLATION

Intuitively, the goal of contrastive learning on graphs is to learn graph representations that are close in the metric space for positive

pairs (graphs with the same labels) and far between negative pairs (graphs with different labels). To achieve this goal, IGSD employs the teacher-student distillation to iteratively refine representations by contrasting latent representations embedded by two networks and using additional predictor and EMA update to avoid collapsing to trivial solutions. Overall, IGSD encourages the closeness of augmented views from the same graph instances while pushing apart the representations from different ones.

#### 4.1 Iterative Graph Self-Distillation Framework

In IGSD, we introduce a teacher-student architecture comprises two networks in similar structure composed by encoder  $f_\theta$ , projector  $g_\theta$  and predictor  $h_\theta$ . We denote the components of the teacher network and the student network as  $f_{\theta'}, g_{\theta'}$  and  $f_\theta, g_\theta, h_\theta$  respectively.

The overview of IGSD is illustrated in Figure 1. In IGSD, the procedure of contrastive learning on negative pairs is described as follows: we first augment the original input graphs  $G_j$  to get augmented view(s)  $G'_j$ . Then  $G'_j$  and different graph instance  $G_i$  are fed respectively into two encoders  $f_\theta, f_{\theta'}$  for extracting graph representations  $\mathbf{h}, \mathbf{h}' = f_\theta(G_i), f_{\theta'}(G'_j)$  with iterative message passing in Eq. (1) and readout functions. The following projectors  $g_\theta, g_{\theta'}$  transform graph representations to projections  $\mathbf{z}, \mathbf{z}'$  via  $\mathbf{z} = g_\theta(\mathbf{h}) = W^{(2)}\sigma(W^{(1)}\mathbf{h})$  and  $\mathbf{z}' = g_{\theta'}(\mathbf{h}') = W'^{(2)}\sigma(W'^{(1)}\mathbf{h}')$ , where  $\sigma$  denotes a ReLU nonlinearity. To prevent collapsing into a trivial solution [15], a specialized predictor is used in the student network for attaining the prediction  $h_\theta(\mathbf{z}) = W_h^{(2)}\sigma(W_h^{(1)}\mathbf{z})$  of the projection  $\mathbf{z}$ . For positive pairs, we follow the same procedure except feeding the original and augmented view of the same graph into two networks respectively.

To contrast latents  $h_\theta(\mathbf{z})$  and  $\mathbf{z}'$ , we use  $L_2$  norm in the latent space to approximate the semantic distance in the input space and the consistency loss can be defined as the mean square error between the normalized prediction  $h_\theta(\mathbf{z})$  and projection  $\mathbf{z}'$ . By passing two graph instances  $G_i$  and  $G_j$  symmetrically, we can obtain the overall consistency loss:

$$\mathcal{L}^{\text{con}}(G_i, G_j) = \left\| h_\theta(\mathbf{z}_i) - \mathbf{z}'_j \right\|_2^2 + \left\| h_\theta(\mathbf{z}'_i) - \mathbf{z}_j \right\|_2^2 \quad (2)$$

With the consistency loss, the teacher network provides a regression target to train the student network, and its parameters  $\theta'$  are updated as an exponential moving average (EMA) of the student parameters  $\theta$  after weights of the student network have been updated using gradient descent:

$$\theta'_t \leftarrow \tau\theta'_{t-1} + (1-\tau)\theta_t \quad (3)$$

With the above iterative self-distillation procedure, we can aggregate information for averaging model weights over each training step instead of using the final weights directly [1]. It should be noted that maintaining a slow-moving average network is also employed in some models like MoCo [20] with different motivations: MoCo uses an EMA of encoder and momentum encoder to update the encoder, ensuring the consistency of dictionary keys in the memory bank. On the other hand, IGSD uses a moving average network to produce prediction targets, enforcing the consistency of teacher and student for training the student network.




Figure 1: Overview of IGSD. Illustration of our framework in the case where we augment input graphs  $G$  once to get  $G'$  for only one forward pass. Blue and red arrows denote contrast on positive and negative pairs respectively.

## 4.2 Self-supervised Learning with IGSD

In IGSD, to contrast the anchor  $G_i$  with other graph instances  $G_j$  (i.e. negative samples), we employ the following self-supervised InfoNCE objective [37]:

$$\mathcal{L}^{\text{self-sup}} = -\mathbb{E}_{G_i \sim \mathcal{G}} \left[ \log \frac{\exp(-\mathcal{L}_{i,i}^{\text{con}})}{\exp(-\mathcal{L}_{i,i}^{\text{con}}) + \sum_{j=1}^{N-1} \mathbb{1}_{i \neq j} \cdot \exp(-\mathcal{L}_{i,j}^{\text{con}})} \right]$$

where  $\mathcal{L}_{i,j}^{\text{con}} = \mathcal{L}^{\text{con}}(G_i, G_j)$ . At the inference time, as semantic interpolations on samples, labels and latents result in better representations and can improve learning performance greatly [3, 49, 56], we obtain the graph representation  $\tilde{h}$  by interpolating the latent representations  $\mathbf{h} = f_\theta(G)$  and  $\mathbf{h}' = f_{\theta'}(G)$  with Mixup function  $\text{Mix}_\lambda(a, b) = \lambda \cdot a + (1 - \lambda) \cdot b$ :

$$\tilde{h} = \text{Mix}_\lambda(\mathbf{h}, \mathbf{h}') \quad (4)$$

## 4.3 Semi-supervised Learning with IGSD

To bridge the gap between self-supervised pretraining and downstream tasks, we extend our model to the semi-supervised setting. In this scenario, it is straightforward to plug in the self-supervised loss as a regularizer for representation learning. However, the instance-wise supervision limited to standard supervised learning may lead to biased negative sampling problems [10]. To tackle this challenge, we can use a small amount of labeled data further to generalize the similarity loss to handle arbitrary number of positive samples belonging to the same class:

$$\mathcal{L}^{\text{supcon}} = \sum_{i=1}^{Kl} \frac{1}{KN_{y'_i}} \sum_{j=1}^{Kl} \mathbb{1}_{i \neq j} \cdot \mathbb{1}_{y'_i = y'_j} \cdot \mathcal{L}^{\text{con}}(G_i, G_j) \quad (5)$$

where  $N_{y'_i}$  denotes the total number of samples in the training set that have the same label  $y'_i$  as anchor  $i$ . Thanks to the graph-level contrastive nature of IGSD, we are able to alleviate the biased negative sampling problems [26] with supervised contrastive loss,

which is crucial [10] but unachievable in most context-instance contrastive learning models since subgraphs are generally hard to assign labels to. Besides, with this loss we are able to fine-tune our model effectively using self-training where pseudo-labels are assigned iteratively to unlabeled data.

With the standard supervised loss like cross entropy or mean square error  $\mathcal{L}(\mathcal{G}_L, \theta)$ , the overall objective can be summarized as:

$$\mathcal{L}^{\text{semi}} = \mathcal{L}(\mathcal{G}_L, \theta) + w\mathcal{L}^{\text{self-sup}}(\mathcal{G}_L \cup \mathcal{G}_U, \theta) + w'\mathcal{L}^{\text{supcon}}(\mathcal{G}_L, \theta) \quad (6)$$

Common semi-supervised learning methods use consistency regularization to measure discrepancy between predictions made on perturbed unlabeled data points for better prediction stability and generalization [36]. By contrast, our methods enforce consistency constraints between latents from different views, which acts as a regularizer for learning directly from labels.

Labeled data provides extra supervision about graph classes and alleviates biased negative sampling. However, they are costly to attain in many areas. Therefore, we develop a contrastive self-training algorithm to leverage label information more effectively than cross entropy in the semi-supervised scenario. In the algorithm, we train the model using a small amount of labeled data and then fine-tune it by iterating between assigning pseudo-labels to unlabeled examples and training models using the augmented dataset. In this way, we harvest massive pseudo-labels for unlabeled examples.

With increasing size of the augmented labeled dataset, the discriminative power of IGSD can be improved iteratively by contrasting more positive pairs belonging to the same class. In this way, we accumulate high-quality pseudo-labels after each iteration to compute the supervised contrastive loss in Eq.(5) and make distinction from conventional self-training algorithms [41]. On the contrary, traditional self-training can use pseudo-labels for computing cross entropy only.

|                 | Datasets        | MUTAG             | IMDB-B            | IMDB-M            | NCI1              | COLLAB            | PTC               |
|-----------------|-----------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| Datasets        | # graphs        | 188               | 1000              | 1500              | 4110              | 5000              | 344               |
|                 | # classes       | 2                 | 2                 | 3                 | 2                 | 3                 | 2                 |
|                 | Avg # nodes     | 17.9              | 19.8              | 13.0              | 29.8              | 74.5              | 25.5              |
| Graph Kernels   | Random Walk     | 83.7 ± 1.5        | 50.7 ± 0.3        | 34.7 ± 0.2        | OMR               | OMR               | 57.9 ± 1.3        |
|                 | SHORTEST PATH   | 85.2 ± 2.4        | 55.6 ± 0.2        | 38.0 ± 0.3        | 51.3 ± 0.6        | 49.8 ± 1.2        | 58.2 ± 2.4        |
|                 | GRAPHLET KERNEL | 81.7 ± 2.1        | 65.9 ± 1.0        | 43.9 ± 0.4        | 53.9 ± 0.4        | 56.3 ± 0.6        | 57.3 ± 1.4        |
|                 | WL SUBTREE      | 80.7 ± 3.0        | 72.3 ± 3.4        | 47.0 ± 0.5        | 55.1 ± 1.6        | 50.2 ± 0.9        | 58.0 ± 0.5        |
|                 | DEEP GRAPH      | 87.4 ± 2.7        | 67.0 ± 0.6        | 44.6 ± 0.5        | 54.5 ± 1.2        | 52.1 ± 1.0        | 60.1 ± 2.6        |
|                 | MLG             | 87.9 ± 1.6        | 66.6 ± 0.3        | 41.2 ± 0.0        | >1 Day            | >1 Day            | <b>63.3 ± 1.5</b> |
|                 | GCKN            | 87.2 ± 6.8        | 70.5 ± 3.1        | 50.8 ± 0.8        | 70.6 ± 2.0        | 54.3 ± 1.0        | 58.4 ± 7.6        |
| self-supervised | GRAPH2VEC       | 83.2 ± 9.6        | 71.1 ± 0.5        | 50.4 ± 0.9        | 73.2 ± 1.8        | 47.9 ± 0.3        | 60.2 ± 6.9        |
|                 | INFOGRAPH       | 89.0 ± 1.1        | 74.2 ± 0.7        | 49.7 ± 0.5        | 73.8 ± 0.7        | 67.6 ± 1.2        | 61.7 ± 1.7        |
|                 | CMC-GRAPH       | 89.7 ± 1.1        | 74.2 ± 0.7        | 51.2 ± 0.5        | 75.0 ± 0.7        | 68.9 ± 1.9        | 62.5 ± 1.7        |
|                 | GCC             | 86.4 ± 0.5        | 71.9 ± 0.5        | 48.9 ± 0.8        | 66.9 ± 0.2        | <b>75.2 ± 0.3</b> | 58.4 ± 1.2        |
|                 | GRAPHCL         | 86.8 ± 1.3        | 71.1 ± 0.4        | 48.4 ± 0.8        | <b>77.9 ± 0.4</b> | 71.4 ± 1.2        | 58.4 ± 1.7        |
|                 | OURS (RANDOM)   | 85.7 ± 2.1        | 71.6 ± 1.2        | 49.2 ± 0.6        | 75.1 ± 0.4        | 65.8 ± 1.0        | 57.6 ± 1.5        |
|                 | OURS            | <b>90.2 ± 0.7</b> | <b>74.7 ± 0.6</b> | <b>51.5 ± 0.3</b> | 75.4 ± 0.3        | 70.4 ± 1.1        | 61.4 ± 1.7        |

**Table 1: Graph classification accuracies (%) for kernels and self-supervised methods on 6 datasets. We report the mean and standard deviation of final results with five runs. ‘>1 Day’ represents that the computation exceeds 24 hours. ‘OMR’ means out of memory error.**

## 5 EXPERIMENTS

### 5.1 Experimental Setup

**Evaluation Tasks.** We conduct experiments by comparing with state-of-the-art models on three tasks. In **graph classification** tasks, we experiment in both the **self-supervised setting** where we only have access to all unlabeled samples in the dataset and the **semi-supervised setting** where we use a small fraction of labeled examples and treat the rest as unlabeled ones by ignoring their labels. In **molecular property prediction** tasks where labels are expensive to obtain, we only consider the **semi-supervised setting**.

**Datasets.** For graph classification tasks, we employ several widely-used graph kernel datasets [25] for learning and evaluation: 3 bioinformatics datasets (MUTAG, PTC, NCI1) and 3 social network datasets (COLLAB, IMDB-B, IMDB-M) with statistics summarized in Table 1. In the semi-supervised graph regression tasks, we use the QM9 dataset containing 134,000 drug-like organic molecules [40] with 9 heavy atoms and select the first ten physicochemical properties as regression targets for training and evaluation.

**Baselines.** In the self-supervised graph classification, we compare with the following representative baselines: InfoGraph [45], CMC-Graph [19], GCC [38], GraphCL[54], Graph2Vec [34] and Graph Kernels including Random Walk Kernel [13], Shortest Path Kernel [24], Graphlet Kernel [44], Weisfeiler-Lehman Sub-tree Kernel (WL SubTree) [43], Deep Graph Kernels [53], Multi-Scale Laplacian Kernel (MLG) [29] and Graph Convolutional Kernel Network (GCKN) [7]<sup>1</sup>.

For the semi-supervised graph classification, we compare our method with competitive baselines like InfoGraph, InfoGraph\* and Mean Teachers. And the GIN baseline doesn’t have access to the unlabeled data. In the semi-supervised molecular property prediction tasks, baselines include InfoGraph, InfoGraph\* and Mean Teachers [46].

<sup>1</sup>We report our reproduced results of GCKN for fair comparisons since the original GCKN paper adopts different train-test splits for nested 10-fold cross-validation, which is different from the Stratified10fold ones of other contrastive learning works [18, 45] and ours.

**Model Configuration.** In our framework, We use GCNs [27] and GINs [52] as encoders to attain node representations for self-supervised and semi-supervised graph classification respectively. For semi-supervised molecular property prediction, we employ message passing neural networks (MPNNs) [14] as backbone encoders to encode molecular graphs with rich edge attributes. All projectors and predictors are implemented as two-layer MLPs.

In semi-supervised molecular property prediction tasks, we generate multiple views based on edge attributes (bond types) of rich-annotated molecular graphs for improving performance. Specifically, we perform label-preserving augmentation to attain multiple diffusion matrixes of every graph on different edge attributes while ignoring others respectively. The diffusion matrix gives a denser graph based on each type of edges to leverage edge features better. We train our models using different numbers of augmented training data and select the amount using cross validation.


For self-supervised graph classification, we adopt LIB-SVM [6] with  $C$  parameter selected in  $\{1e-3, 1e-2, \dots, 1e2, 1e3\}$  as our downstream classifier. Then we use 10-fold cross validation accuracy as the classification performance and repeat the experiments 5 times to report the mean and standard deviation. For semi-supervised graph classification, we randomly select 5% of training data as labeled data while treat the rest as unlabeled one and report the best test set accuracy in 300 epochs. Following the experimental setup in [45], we randomly choose 5000, 10000, 10000 samples for training, validation and testing respectively and the rest are treated as unlabeled training data for the molecular property prediction tasks.

**Hyper-Parameters** For hyper-parameter tuning, we select number of GCN layers over  $\{2, 8, 12\}$ , batch size over  $\{16, 32, 64, 128, 256, 512\}$ , number of epochs over  $\{20, 40, 100\}$  and learning rate over  $\{1e-4, 1e-3\}$  in self-supervised graph classification. The hyper-parameters we tune for semi-supervised graph classification and molecular property prediction are the same in [52] and [45], respectively.

In all experiments, we fix the fixed  $\alpha = 0.2$  for PPR graph diffusion and set the weighting coefficient of Mixup function to be 0.5 and tune our projection hidden size over  $\{1024, 2048\}$  and projection

| Datasets                            | IMDB-B      | IMDB-M      | COLLAB      | NCI1        |
|-------------------------------------|-------------|-------------|-------------|-------------|
| MEAN TEACHERS                       | 69.0        | 49.3        | 72.5        | 71.1        |
| INFOGRAPH*                          | 71.0        | 49.3        | 67.6        | 71.1        |
| GIN (SUPERVISED ONLY)               | 67.0        | 50.0        | 71.4        | 67.9        |
| OURS (SELF-SUPP)                    | 72.0        | 50.0        | 72.6        | 70.6        |
| OURS (SUPCON)                       | 75.0        | 52.0        | 73.4        | 67.9        |
| GIN (SUPERVISED ONLY)+SELF-TRAINING | 72.0        | 51.3        | 70.4        | 74.0        |
| OURS (SELF-SUPP)+SELF-TRAINING      | 73.0        | 54.0        | 71          | 72.5        |
| OURS (SUPCON)+SELF-TRAINING         | <b>77.0</b> | <b>55.3</b> | <b>73.6</b> | <b>77.1</b> |

**Table 2: Graph classification accuracies (%) of semi-supervised experiments on 4 datasets. We report the best results on test set in 300 epochs.**



**Figure 2: Semi-supervised molecular property prediction results in terms of mean absolute error (MAE) ratio. The histogram shows error ratio with respect to supervised results (1.0) of every semi-supervised models. Lower scores are better and a model outperforms the supervised baseline when the score is less than 1.0.**

size over {256, 512}. We start self-training after 30 epochs and tune the number of iterations over {20, 50}, pseudo-labeling threshold over {0.9, 0.95}.

## 5.2 Numerical Results

**Results on self-supervised graph classification.** We first present the results of the self-supervised setting in Table 1. All graph kernels give inferior performance except in the PTC dataset. The Random Walk kernel runs out of memory and the Multi-Scale Laplacian Kernel suffers from a long running time (exceeds 24 hours) in two larger datasets. IGSD outperforms state-of-the-art baselines like InfoGraph, CMC-Graph and GCC in most datasets, showing that IGSD can learn expressive graph-level representations for downstream classifiers. Besides, our model still achieve competitive results in datasets like IMDB-M and NCI1 with random dropping augmentation, which demonstrates the robustness of IGSD with different data augmentation strategies.

**Results on semi-supervised graph classification.** We further apply our model to semi-supervised graph classification tasks with results demonstrated in Table 4, where we set  $w$  and  $w'$  in Eq. (6) to be 1 and 0 as Ours (Self-supp) while 0 and 1 as Ours (SupCon). In this setting, our model performs better than Mean Teachers and InfoGraph\*. Both the self-supervised loss and supervised contrastive loss provide extra performance gain compared with GIN using supervised data only. Besides, both of their performance can be improved significantly combined using self-training especially with supervised contrastive loss. It makes empirical sense

since self-training iteratively assigns pseudo-labels with high confidence to unlabeled data, which provides extra supervision on their categories under contrastive learning framework.

### Results on semi-supervised molecular property prediction.

We present the regression performance of our model measured in the QM9 dataset in Figure 2. We display the performance of our model and baselines as mean square error ratio with respect to supervised results and our model outperforms all baselines in 9 out of 10 tasks compared with strong baselines InfoGraph, InfoGraph\* and Mean Teachers. And in some tasks like R2 (5), U0 (7) and U (8), IGSD achieves significant performance gains against its counterparts, which demonstrates the ability to transfer knowledge learned from self-supervised data for supervised tasks.

## 5.3 Ablation Studies and Analysis

**Effects of self-training.** We first investigate the effects of self-training for our model performance in Table 4. Results show that self-training can improve the GIN baseline and our models with self-supervised loss (self-supp) or supervised contrastive loss (SupCon). The improvement is even more significant combined with supervised contrastive loss since high-quality pseudo-labels provide additional information of graph categories. Moreover, our self-training algorithm consistently outperforms the traditional self-training baseline, which further validates the superiority of our model.






Figure 3: Semi-supervised graph classification accuracy with different proportion of labeled data.




Figure 4: Self-supervised performance with different batch size (number of negative pairs)

**Effects of different amount of negative pairs.** We then conduct ablation experiments on the amount of negative pairs by varying batch size over {16, 32, 64, 128} with results on IMDB-B dataset shown in Figure 4. Both methods contrast negative pairs batch-wise and increasing batch size improves the performance of IGSD while degrades CMC-Graph. When batch size is greater than 32, IGSD outperforms CMC-Graph and the performance gap becomes larger as the batch size increases, which means IGSD is better at leveraging negative pairs for learning effective representations than CMC-Graph.

**Effects of different proportion of labeled data.** We also investigate the performance of different models with different proportion of labeled data with IMDB-B dataset. As illustrated in Figure 3, IGSD outperforms the strong InfoGraph\* baseline given different amount of labeled data consistently. And the performance gain is most significant when the fraction of labeled data is 10% since our models can leverage labels more effectively by regularizing original self-supervised learning objective when labels are scarce.

**Effects of the Projector.** We further investigate the performance with and without the projector on 4 datasets. Results in the Table 3 show that dropping the projector degrades the performance, which indicates the necessity of a projector. It also aligns with the




Figure 5: Effects of projector hidden size on Graph classification accuracies (%).

conclusions of previous contrastive learning works in the image domain like [8], which have empirically shown that using additional projector can improve downstream performance.

Moreover, to investigate the effects of projectors on model performance, we fix the output size of layers in encoders so that their output size is always 512. Then we conducted ablation experiments on different size of the projection head on IMDB-B with the results in Fig. 5. The performance is insensitive to the projection size while a larger projection size could slightly improve the self-supervised learning performance.

| Datasets           | MUTAG    | IMDB-B   | IMDB-M   | PTC      |
|--------------------|----------|----------|----------|----------|
| IGSD               | 90.2±0.7 | 74.7±0.6 | 51.5±0.3 | 61.4±1.7 |
| IGSD w/o PROJECTOR | 87.7±0.9 | 74.2±0.6 | 51.1±0.6 | 56.7±0.9 |

Table 3: Effects of projectors on Graph classification accuracies (%).

## 6 CONCLUSIONS

In this paper, we propose IGSD, a novel graph-level representation learning framework via self-distillation. Our framework iteratively performs teacher-student distillation by instance discrimination on augmented views of graph instances using graph diffusion. Experimental results in both self-supervised and semi-supervised settings show that IGSD is not only able to learn expressive graph representations competitive with state-of-the-art models but also effective with different choices of encoders and augmentation strategies. In the future, we plan to apply our framework to other graph learning tasks and investigate the design of view generators to generate effective views automatically.

## REFERENCES

- [1] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. 2019. There Are Many Consistent Explanations of Unlabeled Data: Why You Should Average. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rgkKBhA5Y7>
- [2] Mikhail Belkin and Partha Niyogi. 2002. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in neural information processing systems*. 585–591.
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. 2019. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*. 5049–5059.
- [4] Piotr Bielak, Tomasz Kajdanowicz, and Nitesh V. Chawla. 2021. Graph Barlow Twins: A self-supervised representation learning framework for graphs. [arXiv:2106.02466](https://arxiv.org/abs/2106.02466) [cs.LG]
- [5] Karsten M Borgwardt and Hans-Peter Kriegel. 2005. Shortest-path kernels on graphs. In *Fifth IEEE international conference on data mining (ICDM'05)*. IEEE, 8–pp.
- [6] Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)* 2, 3 (2011), 1–27.
- [7] Dexiong Chen, Laurent Jacob, and Julien Mairal. 2020. Convolutional kernel networks for graph-structured data. In *International Conference on Machine Learning*. PMLR, 1576–1586.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [9] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. 2020. Big Self-Supervised Models are Strong Semi-Supervised Learners. *Advances in Neural Information Processing Systems* 33 (2020), 22243–22255.
- [10] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. 2020. Debiased Contrastive Learning. In *NeurIPS*.
- [11] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*. 2224–2232.
- [12] Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. Born again neural networks. In *International Conference on Machine Learning*. PMLR, 1607–1616.
- [13] Thomas Gärtner, Peter Flach, and Stefan Wrobel. 2003. On graph kernels: Hardness results and efficient alternatives. In *Learning theory and kernel machines*. Springer, 129–143.
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*. PMLR, 1263–1272.
- [15] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 21271–21284. <https://proceedings.neurips.cc/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf>
- [16] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. [arXiv:2006.07733](https://arxiv.org/abs/2006.07733) [cs.LG]
- [17] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*. 1024–1034.
- [18] Kaveh Hassani and Amir Hosein Khasahmadi. 2020. Contrastive Multi-View Representation Learning on Graphs. In *Proceedings of International Conference on Machine Learning*. 3451–3461.
- [19] Kaveh Hassani and Amir Hosein Khasahmadi. 2020. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*. PMLR, 4116–4126.
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Jun 2020). <https://doi.org/10.1109/cvpr42600.2020.00975>
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. [arXiv preprint arXiv:1503.02531](https://arxiv.org/abs/1503.02531) (2015).
- [22] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2018. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*.
- [23] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2019. Strategies for Pre-training Graph Neural Networks. In *International Conference on Learning Representations*.
- [24] Hisashi Kashima, Koji Tsuda, and Akihiro Inokuchi. 2003. Marginalized kernels between labeled graphs. In *Proceedings of the 20th international conference on machine learning (ICML-03)*. 321–328.
- [25] Kristian Kersting, Nils M. Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. 2016. Benchmark Data Sets for Graph Kernels. <http://graphkernels.cs.tu-dortmund.de>
- [26] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised Contrastive Learning. *Advances in Neural Information Processing Systems* 33 (2020).
- [27] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.
- [28] Johannes Klicpera, Stefan Weissenberger, and Stephan Günnemann. 2019. Diffusion improves graph learning. *Advances in Neural Information Processing Systems* 32 (2019), 13354–13366.
- [29] Risi Kondor and Horace Pan. 2016. The multiscale laplacian graph kernel. In *Advances in Neural Information Processing Systems*. 2990–2998.
- [30] Dong-Hyun Lee. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, Vol. 3.
- [31] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2021. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [32] Takeru Miyato, Shin-Ichi Maeda, Masanori Koyama, and Shin Ishii. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (Aug 2019), 1979–1993. <https://doi.org/10.1109/tpami.2018.2858821>
- [33] Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. 2020. Self-distillation amplifies regularization in hilbert space. [arXiv preprint arXiv:2002.05715](https://arxiv.org/abs/2002.05715) (2020).
- [34] Annamalai Narayanan, Mahinthan Chandramohan, Rajasekar Venkatesan, Lihui Chen, Yang Liu, and Shantanu Jaiswal. 2017. graph2vec: Learning distributed representations of graphs. [arXiv preprint arXiv:1707.05005](https://arxiv.org/abs/1707.05005) (2017).
- [35] Mark EJ Newman and Michelle Girvan. 2004. Finding and evaluating community structure in networks. *Physical review E* 69, 2 (2004), 026113.
- [36] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic Evaluation of Deep Semi-Supervised Learning Algorithms. *Advances in Neural Information Processing Systems* 31 (2018), 3235–3246.
- [37] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. [arXiv preprint arXiv:1807.03748](https://arxiv.org/abs/1807.03748) (2018).
- [38] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. 2020. Gcc: Graph contrastive coding for graph neural network pre-training. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1150–1160.
- [39] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- [40] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data* 1, 1 (2014), 1–7.
- [41] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. 2005. Semi-Supervised Self-Training of Object Detection Models. In *Proceedings of the Seventh IEEE Workshops on Application of Computer Vision (WACV/MOTION'05)-Volume 1-Volume 01*. 29–36.
- [42] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.
- [43] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. 2011. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research* 12, 9 (2011).
- [44] Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. 2009. Efficient graphlet kernels for large graph comparison. In *Artificial Intelligence and Statistics*. 488–495.
- [45] Fan-Yun Sun, Jordan Hoffman, Vikas Verma, and Jian Tang. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=r1lfF2NYvH>
- [46] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems* 30 (2017).
- [47] Michael Tschannen, Josip Djolonga, Paul K Rubenstein, Sylvain Gelly, and Mario Lucic. 2019. On mutual information maximization for representation learning. [arXiv preprint arXiv:1907.13625](https://arxiv.org/abs/1907.13625) (2019).
- [48] Petar Velicković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2018. Deep graph infomax. [arXiv preprint arXiv:1809.10341](https://arxiv.org/abs/1809.10341) (2018).

- [49] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. 2019. Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*. PMLR, 6438–6447.
- [50] Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. 2015. Order Matters: Sequence to sequence for sets. arXiv:1511.06391 [stat.ML]
- [51] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3733–3742.
- [52] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*.
- [53] Pinar Yanardag and SVN Vishwanathan. 2015. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1365–1374.
- [54] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems* 33 (2020), 5812–5823.
- [55] Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. 2020. When does self-supervision help graph convolutional networks?. In *International Conference on Machine Learning*. PMLR, 10871–10880.
- [56] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*.
- [57] Tong Zhao, Yozen Liu, Leonardo Neves, Oliver Woodford, Meng Jiang, and Neil Shah. 2020. Data Augmentation for Graph Neural Networks. arXiv:2006.06830 [cs.LG]
- [58] Yadi Zhou, Fei Wang, Jian Tang, Ruth Nussinov, and Feixiong Cheng. 2020. Artificial intelligence in COVID-19 drug repurposing. *The Lancet Digital Health* (2020).

## A APPENDIX

### A.1 The comparison with GCKN in different evaluation settings

The evaluation setting in original GCKN paper is different from ours and other contrastive learning works like InfoGraph [45] and CMC-Graph [19]: GCKN uses different train-test splits for nested 10-fold cross validation with parameter of SVC is in  $\frac{1}{n} \times np.logspace(-3, 4, 60)$ , requiring 10 times more computation. While we conducted experiments using Stratified10fold cross validation with parameters in SVC in  $\{1e-3, 1e-2, \dots, 1e2, 1e3\}$  to compare with GCKN with results as Table 1 in Section 5.2.

We also evaluated the representations using the train-test splitting indexes provided in GCKN for nested 10-fold cross validation with parameter of SVC is in  $\frac{1}{n} \times np.logspace(-3, 4, 60)$  and the results are as follows:

| Datasets | MUTAG      | IMDB-B     | IMDB-M     | NCI1       | PTC        |
|----------|------------|------------|------------|------------|------------|
| GCKN     | 87.6 ± 2.8 | 76.6 ± 3.8 | 53.1 ± 3.3 | 83.4 ± 2.0 | 60.9 ± 6.9 |
| Ours     | 91.6 ± 7.6 | 75.4 ± 3.9 | 52.3 ± 3.3 | 79.4 ± 2.3 | 63.2 ± 7.8 |

**Table 4: Graph classification accuracies (%) of self-supervised experiments (in the GCKN setting) on 6 datasets.**