

SimGRACE: A Simple Framework for Graph Contrastive Learning without Data Augmentation

Jun Xia^{1,2,3†}, Lirong Wu^{1,2†}, Jintao Chen³, Bozhen Hu^{1,2}, Stan Z.Li^{1,2*}

¹ School of Engineering, Westlake University, Hangzhou 310030, China

² Institute of Advanced Technology, Westlake Institute for Advanced Study, Hangzhou 310030, China

³ Zhejiang University, Hangzhou 310058, China

{xiajun,wulirong,hubozhen,stan.zq.li}@westlake.edu.cn, chenjintao@zju.edu.cn

ABSTRACT

Graph contrastive learning (GCL) has emerged as a dominant technique for graph representation learning which maximizes the mutual information between paired graph augmentations that share the same semantics. Unfortunately, it is difficult to preserve semantics well during augmentations in view of the diverse nature of graph data. Currently, data augmentations in GCL that are designed to preserve semantics broadly fall into three unsatisfactory ways. First, the augmentations can be manually picked per dataset by trial-and-errors. Second, the augmentations can be selected via cumbersome search. Third, the augmentations can be obtained by introducing expensive domain-specific knowledge as guidance. All of these limit the efficiency and more general applicability of existing GCL methods. To circumvent these crucial issues, we propose a Simple framework for GRaph Contrastive LEarning, **SimGRACE** for brevity, which does not require data augmentations. Specifically, we take original graph as input and GNN model with its perturbed version as two encoders to obtain two correlated views for contrast. SimGRACE is inspired by the observation that graph data can preserve their semantics well during encoder perturbations while not requiring manual trial-and-errors, cumbersome search or expensive domain knowledge for augmentations selection. Also, we explain why SimGRACE can succeed. Furthermore, we devise adversarial training scheme, dubbed **AT-SimGRACE**, to enhance the robustness of graph contrastive learning and theoretically explain the reasons. Albeit simple, we show that SimGRACE can yield competitive or better performance compared with state-of-the-art methods in terms of generalizability, transferability and robustness, while enjoying unprecedented degree of flexibility and efficiency.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks; Learning latent representations**; • **Mathematics of computing** → **Graph algorithms**.

[†]Equal Contribution, ^{*}Corresponding Author.

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.

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9096-5/22/04...\$15.00

<https://doi.org/10.1145/3485447.3512156>

KEYWORDS

Graph neural networks, graph self-supervised learning, contrastive learning, graph representation learning, robustness

ACM Reference Format:

Jun Xia, Lirong Wu, Jintao Chen, Bozhen Hu, Stan Z.Li. 2022. SimGRACE: A Simple Framework for Graph Contrastive Learning without Data Augmentation. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3485447.3512156>

1 INTRODUCTION

Graph Neural Networks (GNNs), inheriting the power of neural networks and utilizing the structural information of graph data simultaneously, have achieved overwhelming accomplishments in various graph-based tasks, such as node, graph classification or graph generation [10, 13, 22, 50]. However, most existing GNNs are trained in a supervised manner and it is often resource- and time-intensive to collect abundant labeled data [47]. To remedy this issue, tremendous endeavors have been devoted to graph self-supervised learning that learns representations from unlabeled graphs. Among many, graph contrastive learning (GCL) [53–55] follows the general framework of contrastive learning in computer vision domain [42, 46], in which two augmentations are generated for each graph and then maximizes the mutual information between these two augmented views. In this way, the model can learn representations that are invariant to perturbations. For example, GraphCL [54] first designs four types of general augmentations (node dropping, edge perturbation, attribute masking and subgraph) for GCL. However, these augmentations are not suitable for all scenarios because the structural information and semantics of the graphs varies significantly across domains. For example, GraphCL [54] finds that edge perturbation benefits social networks but hurt some biochemical molecules in GCL. Worse still, these augmentations may alter the graph semantics completely even if the perturbation is weak. For example, dropping a carbon atom in the phenyl ring will alter the aromatic system and result in an alkene chain, which will drastically change the molecular properties [40].

To remedy these issues, several strategies have been proposed recently. Typically, GraphCL [54] manually picks data augmentations per dataset by tedious trial-and-errors, which significantly limits the generality and practicality of their proposed framework. To get rid of the tedious dataset-specific manual tuning of GraphCL, JOAO [53] proposes to automate GraphCL in selecting augmentation pairs. However, it suffers more computational overhead to

Table 1: Comparison between state-of-the-art GCL methods (graph-level representation learning) and SimGRACE.

	No manual trial-and-errors	No domain knowledge	Preserving semantics	No cumbersome search	Generality
GraphCL [54]	✗	✓	✗	✓	✗
MoCL [40]	✓	✗	✓	✓	✗
JOAO(v2) [53]	✓	✓	✗	✗	✓
SimGRACE	✓	✓	✓	✓	✓




Figure 1: Comparison of GraphCL [54], MoCL [40] and SimGRACE on MUTAG dataset. The samples of two classes are distinguished by colors (blue & orange). We first train three GNN encoders with these methods respectively and visualise the representations of original graphs with t-SNE in the upper row. Then, we perturb graphs or encoders in their respective ways (edge perturbation for GraphCL, replacing functional group with bioisosteres of similar properties for MoCL, encoder perturbation for SimGRACE) and visualise the representations of perturbed (GraphCL, MoCL) or original (SimGRACE) graphs in the below row. Unlike GraphCL, SimGRACE and MoCL can preserve the class identity semantics well after perturbations. However, MoCL requires expensive domain knowledge as guidance.

search suitable augmentations and still relies on human prior knowledge in constructing and configuring the augmentation pool to select from. To avoid altering the semantics in the general augmentations adopted in GraphCL and JOAO(v2), MoCL [40] proposes to replace valid substructures in molecular graph with bioisosteres that share similar properties. However, it requires expensive domain knowledge as guidance and can not be applied in other domains like social graphs. Hence, a natural question emerges: *Can we emancipate graph contrastive learning from tedious manual trial-and-errors, cumbersome search or expensive domain knowledge?*

To answer this question, instead of devising more advanced data augmentations strategies for GCL, we attempt to break through state-of-the-arts GCL framework which takes semantic-preserved data augmentations as prerequisite. More specifically, we take original graph data as input and GNN model with its perturbed version as two encoders to obtain two correlated views. And then, we maximize the agreement of these two views. With the encoder perturbation as noise, we can obtain two different embeddings for same input as “positive pairs”. Similar to previous works [42, 54],

we take other graph data in the same mini-batch as “negative pairs”. The idea of encoder perturbation is inspired by the observations in Figure 1. The augmentation or perturbation of MoCL and our SimGRACE can preserve the class identity semantics well while GraphCL can not. Also, we explain why SimGRACE can succeed. Besides, GraphCL [54] shows that GNNs can gain robustness using their proposed framework. However, (1) they do not explain why GraphCL can enhance the robustness; (2) GraphCL seems to be immunized to random attacks well while performing unsatisfactory against adversarial attacks. GROC [19] first integrates adversarial transformations into the graph contrastive learning framework and improves the robustness against adversarial attacks. Unfortunately, as the authors pointed out, the robustness of GROC comes at a price of much longer training time because conducting adversarial transformations for each graph is time-consuming. To remedy these deficiencies, we propose a novel algorithm AT-SimGRACE to perturb the encoder in an adversarial way, which introduces less computational overhead while showing better robustness. Theoretically, we explain why AT-SimGRACE can enhance the robustness. We highlight our contributions as follows:

- *Significance:* We emancipate graph contrastive learning from tedious manual trial-and-errors, cumbersome search or expensive domain knowledge which limit the efficiency and more general applicability of existing GCL methods. The comparison between SimGRACE and state-of-the-art GCL methods can be seen in Table 1.
- *Framework:* We develop a novel and effective framework, SimGRACE, for graph contrastive learning which enjoys unprecedented degree of flexibility, high efficiency and ease of use. Moreover, we explain why SimGRACE can succeed.
- *Algorithm:* We propose a novel algorithm AT-SimGRACE to enhance the robustness of graph contrastive learning. AT-SimGRACE can achieve better robustness while introducing minor computational overhead.
- *Experiments:* We experimentally show that the proposed methods can yield competitive or better performance compared with state-of-the-art methods in terms of generalizability, transferability, robustness and efficiency on multiple social and biochemical graph datasets. The code is available at: <https://github.com/junxia97/SimGRACE>.

2 RELATED WORK

2.1 Generative / Predictive self-supervised learning on graphs

Inspired by the success of self-supervised learning in computer vision [20, 42] and natural language processing [8, 23], tremendous

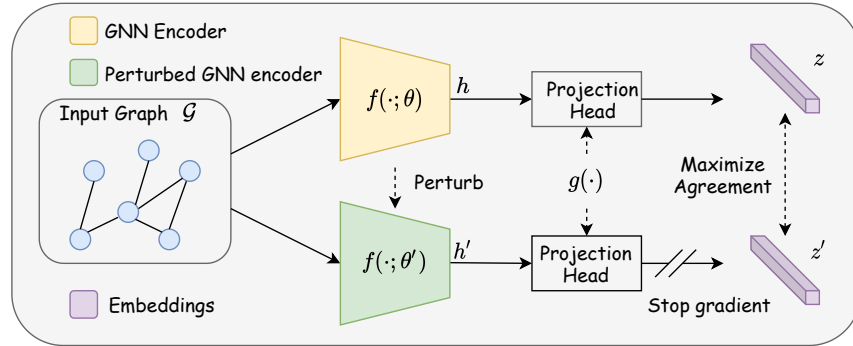


Figure 2: Illustration of SimGRACE, a simple framework of graph contrastive learning. Instead of augmenting the graph data, we feed the original graph \mathcal{G} into a GNN encoder $f(\cdot; \theta)$ and its perturbed version $f(\cdot; \theta')$. After passing a shared projection head $g(\cdot)$, we maximize the agreement between representations z_i and z_j via a contrastive loss.

endeavors have been devoted to graph self-supervised learning that learns representations in an unsupervised manner with designed pretext tasks. Initially, Hu et al. [16] propose two pretext tasks, i.e., predicting neighborhood context and node attributes to conduct node-level pre-training. Besides, they utilize supervised graph-level property prediction and structure similarity prediction as pretext tasks to perform graph-level pre-training. GPT-GNN [17] designs generative task in which node attributes and edges are alternatively generated such that the likelihood of a graph is maximized. Recently, GROVER [34] incorporates GNN into a transformer-style architecture and learns node embedding by predicting contextual property and graph-level motifs. Different from above methods, our SimGRACE follows a contrastive framework that will be introduced in details below.

2.2 Graph Contrastive Learning

Graph contrastive learning can be categorized into two groups. One group can encode useful information by contrasting local and global representations. Initially, DGI [45] and InfoGraph [38] are proposed to obtain expressive representations for graphs or nodes via maximizing the mutual information between graph-level representations and substructure-level representations of different granularity. More recently, MVGRL [14] proposes to learn both node-level and graph-level representation by performing node diffusion and contrasting node representation to augmented graph representations. Another group is designed to learn representations that are tolerant to data transformation. Specifically, they first augment graph data and feed the augmented graphs into a shared encoder and projection head, after which their mutual information is maximized. Typically, for node-level tasks [56, 57], GCA [58] argues that data augmentation schemes should preserve intrinsic structures and attributes of graphs and thus proposes to adopt adaptive augmentations that only perturb unimportant components. DGCL [48] introduces a novel probabilistic method to alleviate the issue of false negatives in GCL. For graph-level tasks, GraphCL [54] proposes four types of augmentations for general graphs and demonstrated that the learned representations can help downstream tasks. However, the success of GraphCL comes at the price of tedious manual trial-and-errors. To tackle this issue, JOAO [53] proposes a unified bi-level

optimization framework to automatically select data augmentations for GraphCL, which is time-consuming and inconvenient. More recently, MoCL [40] proposes to incorporate domain knowledge into molecular graph augmentations in order to preserve the semantics. However, the domain knowledge is extremely expensive. Worse still, MoCL can only work on molecular graph data, which significantly limits their generality. Despite the fruitful progress, they still require tedious manual trial-and-errors, cumbersome search or expensive domain knowledge for augmentation selection. Instead, our SimGRACE breaks through state-of-the-arts GCL framework that takes semantic-preserved data augmentations as prerequisite.

3 METHOD

3.1 SimGRACE

In this section, we will introduce SimGRACE framework in details. As sketched in Figure 2, the framework consists of the following three major components:

(1) Encoder perturbation. A GNN encoder $f(\cdot; \theta)$ and its perturbed version $f(\cdot; \theta')$ first extract two graph-level representations h and h' for the same graph \mathcal{G} , which can be formulated as,

$$h = f(\mathcal{G}; \theta), h' = f(\mathcal{G}; \theta'). \quad (1)$$

The method we proposed to perturb the encoder $f(\cdot; \theta)$ can be mathematically described as,

$$\theta'_l = \theta_l + \eta \cdot \Delta \theta_l; \quad \Delta \theta_l \sim \mathcal{N}(0, \sigma_l^2), \quad (2)$$

where θ_l and θ'_l are the weight tensors of the l -th layer of the GNN encoder and its perturbed version respectively. η is the coefficient that scales the magnitude of the perturbation. $\Delta \theta_l$ is the perturbation term which samples from Gaussian distribution with zero mean and variance σ_l^2 . Also, we show that the performance will deteriorate when we set $\eta = 0$ in section 4.6.1. Note that BGRL [41] and MERIT [18] also update a target network with an online encoder during training. However, SimGRACE differs from them in three aspects: (1) SimGRACE perturbs the encoder with a random Gaussian noise instead of momentum updating; (2) SimGRACE does not require data augmentation while BGRL and MERIT take it as

prerequisite. (3) SimGRACE focuses on graph-level representation learning while BGRL and MERIT only work in node-level tasks.

(2) **Projection head.** As advocated in [42], a non-linear transformation $g(\cdot)$ named projection head maps the representations to another latent space can enhance the performance. In our SimGRACE framework, we also adopt a two-layer perceptron (MLP) to obtain z and z' ,

$$z = g(\mathbf{h}), z' = g(\mathbf{h}'). \quad (3)$$

(3) **Contrastive loss.** In SimGRACE framework, we utilize the normalized temperature-scaled cross entropy loss (NT-Xent) as previous works [30, 37, 46, 54] to enforce the agreement between positive pairs z and z' compared with negative pairs.

During SimGRACE training, a minibatch of N graphs are randomly sampled and then they are fed into a GNN encoder $f(\cdot; \theta)$ and its perturbed version $f(\cdot; \theta')$, resulting in two presentations for each graph and thus $2N$ representations in total. We re-denote z, z' as z_n, z'_n for n -th graph in the minibatch. Negative pairs are generated from the other $N - 1$ perturbed representations within the same mini-batch as in [5, 42, 54]. Denoting the cosine similarity function as $\text{sim}(z, z') = z^T z' / \|z\| \|z'\|$, the contrastive loss for the n -th graph is defined as,

$$\ell_n = -\log \frac{\exp(\text{sim}(z_n, z'_n) / \tau)}{\sum_{n'=1, n' \neq n}^N \exp(\text{sim}(z_n, z_{n'}) / \tau)}, \quad (4)$$

where τ is the temperature parameter. The final loss is computed across all positive pairs in the minibatch.

3.2 Why can SimGRACE work well?

In order to understand why SimGRACE can work well, we first introduce the analysis tools from [43]. Specifically, they identify two key properties related to contrastive learning: *alignment* and *uniformity* and then propose two metrics to measure the quality of representations obtained via contrastive learning. One is the alignment metric which is straightforwardly defined with the expected distance between positive pairs:

$$\ell_{\text{align}}(f; \alpha) \triangleq \mathbb{E}_{(x, y) \sim p_{\text{pos}}} [\|f(x) - f(y)\|_2^\alpha], \quad \alpha > 0 \quad (5)$$

where p_{pos} is the distribution of positive pairs (augmentations of the same sample). This metric is well aligned with the objective of contrastive learning: positive samples should stay close in the embedding space. Analogously, for our SimGRACE framework, we provide a modified metric for *alignment*,

$$\ell_{\text{align}}(f; \alpha) \triangleq \mathbb{E}_{x \sim p_{\text{data}}} [\|f(x; \theta) - f(x; \theta')\|_2^\alpha], \quad \alpha > 0 \quad (6)$$

where p_{data} is the data distribution. We set $\alpha = 2$ in our experiments. The other is the uniformity metric which is defined as the logarithm of the average pairwise Gaussian potential:

$$\ell_{\text{uniform}}(f; \alpha) \triangleq \log \mathbb{E}_{x, y \stackrel{i.i.d.}{\sim} p_{\text{data}}} \left[e^{-t \|f(x; \theta) - f(y; \theta)\|_2^\alpha} \right]. \quad t > 0 \quad (7)$$

In our experiments, we set $t = 2$. The uniformity metric is also aligned with the objective of contrastive learning that the embeddings of random samples should scatter on the hypersphere. We take the checkpoints of SimGRACE, GraphCL and MoCL every 2 epochs during training and visualize the alignment ℓ_{align} and

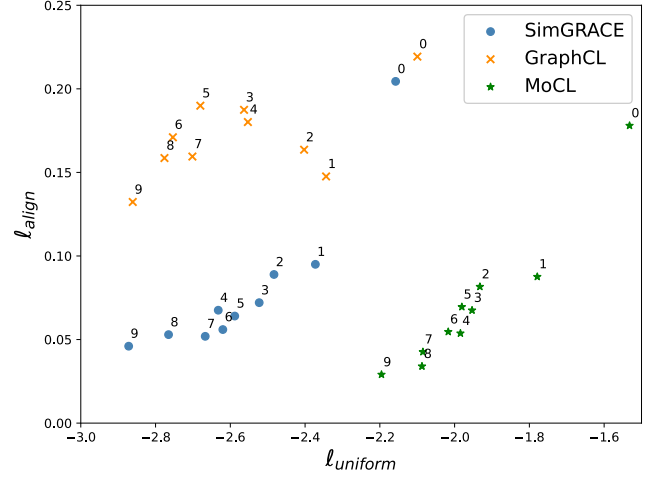


Figure 3: $\ell_{\text{align}} - \ell_{\text{uniform}}$ plot for SimGRACE, GraphCL and MoCL on MUTAG dataset. The numbers around the points are the indexes of epochs. For both ℓ_{align} and ℓ_{uniform} , lower is better.

uniformity ℓ_{uniform} metrics in Figure 3. As can be observed, all the three methods can improve the alignment and uniformity. However, GraphCL achieves a smaller gain on the alignment than SimGRACE and MoCL. In other words, the positive pairs can not stay close in GraphCL because general graph data augmentations (drop edges, drop nodes and etc.) destroy the semantics of original graph data, which degrades the quality of the representations learned by GraphCL. Instead, MoCL augments graph data with domain knowledge as guidance and thus can preserve semantics during augmentation. Eventually, MoCL dramatically improves the alignment. Compared with GraphCL, SimGRACE can achieve better alignment while improving uniformity because encoder perturbation can preserve data semantics well. On the other hand, although MoCL achieves better alignment than SimGRACE via introducing domain knowledge as guidance, it only achieves a small gain on the uniformity, and eventually underperforms SimGRACE.

3.3 AT-SimGRACE

Recently, GraphCL [54] shows that GNNs can gain robustness using their proposed framework. However, they did not explain why GraphCL can enhance the robustness. Additionally, GraphCL seems to be immunized to random attacks well while being unsatisfactory against adversarial attacks. In this section, we aim to utilize Adversarial Training (AT) [11, 24] to improve the adversarial robustness of SimGRACE in a principled way. Generally, AT directly incorporates adversarial examples into the training process to solve the following optimization problem:

$$\min_{\theta} \mathcal{L}'(\theta), \quad \text{where} \quad \mathcal{L}'(\theta) = \frac{1}{n} \sum_{i=1}^n \max_{\|x'_i - x_i\|_p \leq \epsilon} \ell'_i(f(x'_i; \theta), y_i), \quad (8)$$

where n is the number of training examples, x'_i is the adversarial example within the ϵ -ball (bounded by an L_p -norm) centered at

natural example \mathbf{x}_i , f is the DNN with weight θ , $\ell'(\cdot)$ is the standard supervised classification loss (e.g., the cross-entropy loss), and $\mathcal{L}'(\theta)$ is called the "adversarial loss". However, above general framework of AT can not directly be applied in graph contrastive learning because (1) AT requires labels as supervision while labels are not available in graph contrastive learning; (2) Perturbing each graph for the dataset in an adversarial way will introduce heavy computational overhead, which has been pointed out in GROC [19]. To remedy the first issue, we substitute supervised classification loss in Eq. (8) with contrastive loss in Eq. (4). To tackle the second issue, instead of conducting adversarial transformation of graph data, we perturb the encoder in an adversarial way, which is more computationally efficient.

Assuming that Θ is the weight space of GNNs, for any \mathbf{w} and any positive ϵ , we can define the norm ball in θ with radius ϵ centered at \mathbf{w} as,

$$\mathbf{R}(\mathbf{w}; \epsilon) := \{\theta \in \Theta : \|\theta - \mathbf{w}\| \leq \epsilon\}, \quad (9)$$

we choose L_2 norm to define the norm ball in our experiments. With this definition, we can now formulate our AT-SimGRACE as an optimization problem,

$$\min_{\theta} \mathcal{L}(\theta + \Delta),$$

where $\mathcal{L}(\theta + \Delta) = \frac{1}{M} \sum_{i=1}^M \max_{\Delta \in \mathbf{R}(0; \epsilon)} \ell_i(f(\mathcal{G}_i; \theta + \Delta), f(\mathcal{G}_i; \theta))$,

(10)

where M is the number of graphs in the dataset. We propose Algorithm 1 to solve this optimization problem. Specifically, for inner maximization, we forward I steps to update Δ in the direction of increasing the contrastive loss using gradient ascent algorithm. With the output perturbation Δ of inner maximization, the outer loops update the weights θ of GNNs with mini-batched SGD.

Algorithm 1: Encoder perturbation of AT-SimGRACE

Data: Graph dataset $\mathcal{D} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_M\}$, contrastive loss ℓ , batch size N , initial encoder weights θ , inner iterations I , inner learning rate ζ , outer learning rate γ , norm ball radius ϵ .

```

1 for each mini-batch do
2   Sample  $\mathcal{D}_B = \{\mathcal{G}_i\}_{i=1}^N$  from  $\mathcal{D}$ ;
3   Initialize perturbation:  $\Delta \leftarrow 0$ ;
4   for  $t = 0, 1, 2, \dots, I - 1$  do
5     Update perturbation:
6      $\Delta \leftarrow \Delta + \zeta \sum_{i=1}^N \nabla_{\theta} \ell_i(f(\mathcal{G}_i; \theta + \Delta), f(\mathcal{G}_i; \theta)) / N$ ;
7     if  $\|\Delta\|_2 > \epsilon$  then
8       | Normalize perturbation:  $\Delta \leftarrow \epsilon \Delta / \|\Delta\|_2$ ;
9     end
10  end
11 Update weights:
12   $\theta' \leftarrow \theta - \gamma \sum_{i=1}^N \nabla_{\theta} \ell_i(f(\mathcal{G}_i; \theta + \Delta), f(\mathcal{G}_i; \theta)) / N$ .
13 end

```

3.4 Theoretical Justification

In this section, we aim to explain the reasons why AT-SimGRACE can enhance the robustness of graph contrastive learning. To start, it is widely accepted that flatter loss landscape can bring robustness [3, 31, 44]. For example, as formulated in Eq. 8, adversarial training (AT) enhances robustness via restricting the change of loss when the input of models is perturbed indeed. Thus, we want to theoretically justify why AT-SimGRACE works via validating that AT-SimGRACE can flatten the loss landscape. Inspired by previous work [29] that connects sharpness of loss landscape and PAC-Bayes theory [25, 26], we utilize PAC-Bayes framework to derive guarantees on the expected error. Assuming that the prior distribution P over the weights is a zero mean, σ^2 variance Gaussian distribution, with probability at least $1 - \delta$ over the draw of M graphs, the expected error of the encoder can be bounded as:

$$\mathbb{E}_{\{\mathcal{G}_i\}_{i=1}^M, \Delta} [\mathcal{L}(\theta + \Delta)] \leq \mathbb{E}_{\Delta} [\mathcal{L}(\theta + \Delta)] + 4\sqrt{\frac{KL(\theta + \Delta \| P) + \ln \frac{2M}{\delta}}{M}}. \quad (11)$$

We choose Δ as a zero mean spherical Gaussian perturbation with variance σ^2 in every direction, and set the variance of the perturbation to the weight with respect to its magnitude $\sigma = \alpha \|\theta\|$. Besides, we substitute $\mathbb{E}_{\Delta} [\mathcal{L}(\theta + \Delta)]$ with $\mathcal{L}(\theta) + \mathbb{E}_{\Delta} [\mathcal{L}(\theta + \Delta)] - \mathcal{L}(\theta)$. Then, we can rewrite Eq. 11 as:

$$\mathbb{E}_{\{\mathcal{G}_i\}_{i=1}^M} [\mathcal{L}(\theta + \Delta)] \leq \mathcal{L}(\theta) + \underbrace{\{\mathbb{E}_{\Delta} [\mathcal{L}(\theta + \Delta)] - \mathcal{L}(\theta)\}}_{\text{Expected sharpness}} + 4\sqrt{\frac{1}{M} \left(\frac{1}{2\alpha} + \ln \frac{2M}{\delta} \right)}. \quad (12)$$

It is obvious that $\mathbb{E}_{\Delta} [\mathcal{L}(\theta + \Delta)] \leq \max_{\Delta} [\mathcal{L}(\theta + \Delta)]$ and the third term $4\sqrt{\frac{1}{M} \left(\frac{1}{2\alpha} + \ln \frac{2M}{\delta} \right)}$ is a constant. Thus, AT-SimGRACE optimizes the worst-case of sharpness of loss landscape $\max_{\Delta} [\mathcal{L}(\theta + \Delta)] - \mathcal{L}(\theta)$ to the bound of the expected error, which explains why AT-SimGRACE can enhance the robustness.

4 EXPERIMENTS

In this section, we conduct experiments to evaluate SimGRACE and AT-SimGRACE through answering the following research questions.

- **RQ1. (Generalizability)** Does SimGRACE outperform competitors in unsupervised and semi-supervised settings?
- **RQ2. (Transferability)** Can GNNs pre-trained with SimGRACE show better transferability than competitors?
- **RQ3. (Robustness)** Can AT-SimGRACE perform better than existing competitors against various adversarial attacks?
- **RQ4. (Efficiency)** How about the efficiency (time and memory) of SimGRACE? Does it more efficient than competitors?
- **RQ5. (Hyperparameters Sensitivity)** Is the proposed SimGRACE sensitive to hyperparameters like the magnitude of the perturbation η , training epochs and batch size?

Table 2: Comparing classification accuracy with baselines under the same experiment setting. The top three accuracy or rank for each dataset are emphasized in bold. A.R. denotes average rank. - indicates that results are not available in published papers.

Methods	NCI1	PROTEINS	DD	MUTAG	COLLAB	RDT-B	RDT-M5K	IMDB-B	A.R. ↓
GL	–	–	–	81.66 ± 2.11	–	77.34 ± 0.18	41.01 ± 0.17	65.87 ± 0.98	8.3
WL	80.01 ± 0.50	72.92 ± 0.56	–	80.72 ± 3.00	–	68.82 ± 0.41	46.06 ± 0.21	72.30 ± 3.44	6.2
DGK	80.31 ± 0.46	73.30 ± 0.82	–	87.44 ± 2.72	–	78.04 ± 0.39	41.27 ± 0.18	66.96 ± 0.56	5.5
node2vec	54.89 ± 1.61	57.49 ± 3.57	–	72.63 ± 10.20	–	–	–	–	9.0
sub2vec	52.84 ± 1.47	53.03 ± 5.55	–	61.05 ± 15.80	–	71.48 ± 0.41	36.68 ± 0.42	55.26 ± 1.54	10.2
graph2vec	73.22 ± 1.81	73.30 ± 2.05	–	83.15 ± 9.25	–	75.78 ± 1.03	47.86 ± 0.26	71.10 ± 0.54	6.7
MVGRL	–	–	–	75.40 ± 7.80	–	82.00 ± 1.10	–	63.60 ± 4.20	8.3
InfoGraph	76.20 ± 1.06	74.44 ± 0.31	72.85 ± 1.78	89.01 ± 1.13	70.65 ± 1.13	82.50 ± 1.42	53.46 ± 1.03	73.03 ± 0.87	3.8
GraphCL	77.87 ± 0.41	74.39 ± 0.45	78.62 ± 0.40	86.80 ± 1.34	71.36 ± 1.15	89.53 ± 0.84	55.99 ± 0.28	71.14 ± 0.44	3.1
JOAO	78.07 ± 0.47	74.55 ± 0.41	77.32 ± 0.54	87.35 ± 1.02	69.50 ± 0.36	85.29 ± 1.35	55.74 ± 0.63	70.21 ± 3.08	4.3
JOAOv2	78.36 ± 0.53	74.07 ± 1.10	77.40 ± 1.15	87.67 ± 0.79	69.33 ± 0.34	86.42 ± 1.45	56.03 ± 0.27	70.83 ± 0.25	3.6
SimGRACE	79.12 ± 0.44	75.35 ± 0.09	77.44 ± 1.11	89.01 ± 1.31	71.72 ± 0.82	89.51 ± 0.89	55.91 ± 0.34	71.30 ± 0.77	2.0

4.1 Experimental Setup

4.1.1 Datasets. For unsupervised and semi-supervised learning, we use datasets from the benchmark TUDataset [27], including graph data for various social networks [2, 52] and biochemical molecules [9, 32]. For transfer learning, we perform pre-training on ZINC-2M and PPI-306K and finetune the model with various datasets including PPI, BBBP, ToxCast and SIDER.

4.1.2 Evaluation Protocols. Following previous works for graph-level self-supervised representation learning [39, 53, 54], we evaluate the generalizability of the learned representations on both unsupervised and semi-supervised settings. In unsupervised setting, we train SimGRACE using the whole dataset to learn graph representations and feed them into a downstream SVM classifier with 10-fold cross-validation. For semi-supervised setting, we pre-train GNNs with SimGRACE on all the data and did finetuning & evaluation with K ($K = \frac{1}{\text{label rate}}$) folds for datasets without the explicit training/validation/test split. For datasets with the train/validation/test split, we pre-train GNNs with the training data, finetuning on the partial training data and evaluation on the validation/test sets. More details can be seen in the appendix.

4.1.3 Compared baselines. We compare SimGRACE with state-of-the-arts graph kernel methods including GL [36], WL [35] and DGK [52]. Also, we compare SimGRACE with other graph self-supervised learning methods: GAE [21], node2vec [12], sub2vec [1], graph2vec [28], EdgePred [16], AttrMasking [16], ContextPred [16], Infomax (DGI) [45], InfoGraph [39] and instance-instance contrastive methods GraphCL [54], JOAO(v2) [53].

4.2 Unsupervised and semi-supervised learning (RQ1)

For unsupervised representation learning, as can be observed in Table 2, SimGRACE outperforms other baselines and always ranks top three on all the datasets. Generally, SimGRACE performs better on biochemical molecules compared with data augmentation based methods. The reason is that the semantics of molecular graphs are

Table 3: Transfer learning comparison with other pre-training schemes. The top-3 accuracy for each dataset are emphasized in bold.

Pre-Train dataset	PPI-306K	ZINC 2M		
Fine-Tune dataset	PPI	BBBP	ToxCast	SIDER
No Pre-Train	64.8 ± 1.0	65.8 ± 4.5	63.4 ± 0.6	57.3 ± 1.6
EdgePred	65.7 ± 1.3	68.8 ± 0.8	62.7 ± 0.4	58.4 ± 0.8
AttrMasking	65.2 ± 1.6	67.3 ± 2.4	64.1 ± 0.6	60.4 ± 0.7
ContextPred	64.4 ± 1.3	64.3 ± 2.8	64.2 ± 0.5	61.0 ± 0.7
GraphCL	67.88 ± 0.85	68.0 ± 2.0	63.9 ± 0.6	60.9 ± 0.6
JOAO	64.43 ± 1.38	69.68 ± 0.67	62.40 ± 0.57	60.53 ± 0.88
JOAOv2	63.94 ± 1.59	70.22 ± 0.98	62.94 ± 0.48	59.97 ± 0.79
SimGRACE	70.25 ± 1.22	71.25 ± 0.86	63.36 ± 0.52	60.59 ± 0.96

more fragile compared with social networks. General augmentations (drop nodes, drop edges and etc.) adopted in other baselines will not alter the semantics of social networks dramatically. For semi-supervised task, as can be observed in Table 4, we report two semi-supervised tasks with 1% and 10% label rate respectively. In 1% setting, SimGRACE outperforms previous baselines by a large margin or matching the performance of SOTA methods. For 10% setting, SimGRACE performs comparably to SOTA methods including GraphCL and JOAO(v2) whose augmentations are derived via expensive trial-and-errors or cumbersome search.

4.3 Transferability (RQ2)

To evaluate the transferability of the pre-training scheme, we conduct experiments on transfer learning on molecular property prediction in chemistry and protein function prediction in biology following previous works [16, 49, 54]. Specifically, we pre-train and fine-tune the models with different datasets. For pre-training, learning rate is tuned in {0.01, 0.1, 1.0} and epoch number in {20, 40, 60, 80, 100} where grid search is performed. As sketched in Table 3, there is no universally beneficial pre-training scheme especially for the out-of-distribution scenario in transfer learning. However, SimGRACE shows competitive or better transferability than other pre-training schemes, especially on PPI dataset.

Table 4: Comparing classification accuracy with baselines under the same semi-supervised setting. The top three accuracy or rank are emphasized in bold. – indicates that label rate is too low for a given dataset size. L.R. and A.R. are short for label rate and average rank, respectively.

L.R.	Methods	NCI1	PROTEINS	DD	COLLAB	RDT-B	RDT-M5K	A.R. ↓
1%	No pre-train.	60.72 ± 0.45	–	–	57.46 ± 0.25	–	–	8.5
	Augmentations	60.49 ± 0.46	–	–	58.40 ± 0.97	–	–	8.0
	GAE	61.63 ± 0.84	–	–	63.20 ± 0.67	–	–	5.5
	Infomax	62.72 ± 0.65	–	–	61.70 ± 0.77	–	–	4.0
	ContextPred	61.21 ± 0.77	–	–	57.60 ± 2.07	–	–	7.5
	GraphCL	62.55 ± 0.86	–	–	64.57 ± 1.15	–	–	2.0
	JOAO	61.97 ± 0.72	–	–	63.71 ± 0.84	–	–	4.5
	JOAOv2	62.52 ± 1.16	–	–	64.51 ± 2.21	–	–	3.0
SimGRACE	64.21 ± 0.65	–	–	64.28 ± 0.98	–	–	2.0	
10%	No pre-train.	73.72 ± 0.24	70.40 ± 1.54	73.56 ± 0.41	73.71 ± 0.27	86.63 ± 0.27	51.33 ± 0.44	7.7
	Augmentations	73.59 ± 0.32	70.29 ± 0.64	74.30 ± 0.81	74.19 ± 0.13	87.74 ± 0.39	52.01 ± 0.20	7.0
	GAE	74.36 ± 0.24	70.51 ± 0.17	74.54 ± 0.68	75.09 ± 0.19	87.69 ± 0.40	33.58 ± 0.13	6.3
	Infomax	74.86 ± 0.26	72.27 ± 0.40	75.78 ± 0.34	73.76 ± 0.29	88.66 ± 0.95	53.61 ± 0.31	3.7
	ContextPred	73.00 ± 0.30	70.23 ± 0.63	74.66 ± 0.51	73.69 ± 0.37	84.76 ± 0.52	51.23 ± 0.84	8.3
	GraphCL	74.63 ± 0.25	74.17 ± 0.34	76.17 ± 1.37	74.23 ± 0.21	89.11 ± 0.19	52.55 ± 0.45	2.8
	JOAO	74.48 ± 0.27	72.13 ± 0.92	75.69 ± 0.67	75.30 ± 0.32	88.14 ± 0.25	52.83 ± 0.54	4.2
	JOAOv2	74.86 ± 0.39	73.31 ± 0.48	75.81 ± 0.73	75.53 ± 0.18	88.79 ± 0.65	52.71 ± 0.28	2.5
SimGRACE	74.60 ± 0.41	74.03 ± 0.51	76.48 ± 0.52	74.74 ± 0.28	88.86 ± 0.62	53.97 ± 0.64	2.3	

Table 5: Performance under three adversarial attacks for GNN with different depth following the protocols in [7].

Methods	Two-Layer			Three-Layer			Four-Layer		
	No Pre-Train	GraphCL	AT-SimGRACE	No Pre-Train	GraphCL	AT-SimGRACE	No Pre-Train	GraphCL	AT-SimGRACE
Unattack	93.20	94.73	94.24	98.20	98.33	99.32	98.87	99.00	99.13
RandSampling	78.73	80.68	81.73	92.27	92.60	94.27	95.13	97.40	97.67
GradArgmax	69.47	69.26	75.13	64.60	89.33	93.00	95.80	97.00	96.60
RL-S2V	42.93	42.20	44.86	41.93	61.66	66.00	70.20	84.86	85.29

Table 6: Comparisons of efficiency on three graph datasets. Note that we do not take the time for manual trial-and-errors of GraphCL into consideration. In fact, picking the suitable augmentations manually for GraphCL is much more time-consuming. All the three methods are evaluated on a 32GB V100 GPU.

Dataset	Algorithm	Training Time	Memory
PROTEINS	GraphCL	111s	1231MB
	JOAOv2	4088s	1403MB
	SimGRACE	46 s	1175 MB
COLLAB	GraphCL	1033s	10199MB
	JOAOv2	10742s	7303MB
	SimGRACE	378 s	6547 MB
RDT-B	GraphCL	917s	4135MB
	JOAOv2	10278s	3935MB
	SimGRACE	280 s	2729 MB

4.4 Adversarial robustness (RQ3)

Following previous works [7, 54], we perform on synthetic data to classify the component number in graphs, facing the RandSampling, GradArgmax and RL-S2V attacks, to evaluate the robustness of AT-SimGRACE. To keep fair, we adopt Structure2vec [6] as the GNN encoder as in [7, 54]. Besides, we pretrain the GNN encoder for 150 epochs because it takes longer time for the convergence of adversarial training. We set the inner learning rate $\zeta = 0.001$ and the radius of perturbation ball $\epsilon = 0.01$. As demonstrated in Table 5, AT-SimGRACE boosts the robustness of GNNs dramatically compared with training from scratch and GraphCL under three typical evasion attacks.

4.5 Efficiency (Training time and memory cost) (RQ4)

In Table 6, we compare the performance of SimGRACE with the state-of-the-arts methods including GraphCL and JOAOv2 in terms of their training time and the memory overhead. Here, the training

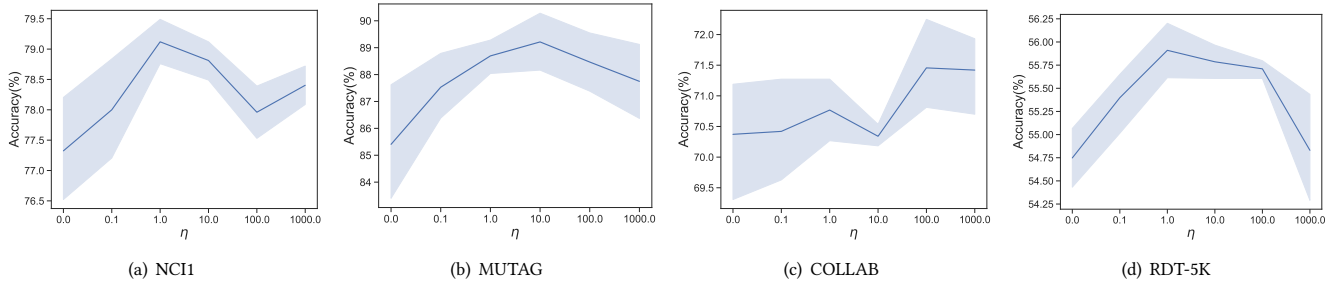


Figure 4: Performance versus magnitude of the perturbation (η) in unsupervised representation learning task.

time refers to the time for pre-training stage of the semi-supervised task and the memory overhead refers to total memory costs of model parameters and all hidden representations of a batch. As can be observed, SimGRACE runs near 40-90 times faster than JOAOv2 and 2.5-4 times faster than GraphCL. If we take the time for manual trial-and-errors in GraphCL into consideration, the superiority of SimGRACE will be more pronounced. Also, SimGRACE requires less computational memory than GraphCL and JOAOv2. In particular, the efficiency of SimGRACE can be more prominent on large-scale social graphs, such as COLLAB and RDT-B.

4.6 Hyper-parameters sensitivity analysis (RQ5)

4.6.1 Magnitude of the perturbation. As can be observed in Figure 4, weight perturbation is crucial in SimGRACE. If we set the magnitude of the perturbation as zero ($\eta = 0$), the performance is usually the lowest compared with other setting of perturbation across these four datasets. This observation aligns with our intuition. Without perturbation, SimGRACE simply compares two original samples as a negative pair while the positive pair loss becomes zero, leading to homogeneously pushes all graph representations away from each other, which is non-intuitive to justify. Instead, appropriate perturbations enforce the model to learn representations invariant to the perturbations through maximizing the agreement between a graph and its perturbation. Besides, well aligned with previous works [15, 33] that claim "hard" positive pairs and negative pairs can boost the performance of contrastive learning, we can observe that larger magnitude (within an appropriate range) of the perturbation can bring consistent improvement of the performance. However, over-large perturbations will lead to performance degradation because the semantics of graph data are not preserved.

4.6.2 Batch-size and training epochs. Figure 5 demonstrates the performance of SimGRACE trained with various batch size and epochs. Generally, larger batch size or training epochs can bring better performance. The reason is that larger batch size will provide more negative samples for contrasting. Similarly, training longer also provides more new negative samples for each sample because the split of total datasets is more various with more training epochs. In our experiments, to keep fair, we follow the same settings of other competitors [53, 54] via training the GNN encoder with batch size as 128 and number of epochs as 20. In fact, we can further




Figure 5: Performance of SimGRACE trained with different batch size and epochs on NCI1 dataset.

improve the performance of SimGRACE with larger batch size and longer training time.

5 CONCLUSIONS

In this paper, we propose a simple framework (SimGRACE) for graph contrastive learning. Although it may appear simple, we demonstrate that SimGRACE can outperform or match the state-of-the-art competitors on multiple graph datasets of various scales and types, while enjoying unprecedented degree of flexibility, high efficiency and ease of use. We emancipate graph contrastive learning from tedious manual tuning, cumbersome search or expensive domain knowledge. Furthermore, we devise adversarial training schemes to enhance the robustness of SimGRACE in a principled way and theoretically explain the reasons. There are two promising avenues for future work: (1) exploring if encoder perturbation can work well in other domains like computer vision and natural language processing. (2) applying the pre-trained GNNs to more real-world tasks including social analysis and biochemistry.

ACKNOWLEDGMENTS

This work is supported in part by the Science and Technology Innovation 2030 - Major Project (No. 2021ZD0150100) and National Natural Science Foundation of China (No. U21A20427).

REFERENCES

- [1] Bijaya Adhikari, Yao Zhang, Naren Ramakrishnan, and Aditya B. Prakash. 2018. Sub2Vec: Feature Learning for Subgraphs. *ADVANCES IN KNOWLEDGE DISCOVERY AND DATA MINING, PAKDD 2018, PT II* (2018), 170–182.
- [2] Rozemberczki Benedek, Kiss Oliver, and Sarkar Rik. 2020. An API Oriented Open-source Python Framework for Unsupervised Learning on Graphs. (2020).
- [3] Liu Chen, Salzmann Mathieu, Lin Tao, Tomioka Ryota, and Süssstrunk Sabine. 2020. On the Loss Landscape of Adversarial Training: Identifying Challenges and How to Overcome Them. *NIPS 2020* (2020).
- [4] Ting Chen, Song Bian, and Yizhou Sun. 2019. Are Powerful Graph Neural Nets Necessary? A Dissection on Graph Classification. *arXiv: Learning* (2019).
- [5] Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. 2017. On Sampling Strategies for Neural Network-based Collaborative Filtering. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2017), 767–776.
- [6] Hanjun Dai, Bo Dai, and Le Song. 2016. Discriminative Embeddings of Latent Variable Models for Structured Data. *ICML* (2016).
- [7] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial Attack on Graph Structured Data. *international conference on machine learning* (2018), 1123–1132.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *north american chapter of the association for computational linguistics* (2019).
- [9] D. Paul Dobson and J. Andrew Doig. 2003. Distinguishing Enzyme Structures from Non-enzymes Without Alignments. *Journal of Molecular Biology* (2003), 771–783.
- [10] Yuanqi Du, Shiyu Wang, Xiaojie Guo, Hengning Cao, Shujie Hu, Junji Jiang, Aishwarya Varala, Abhinav Angirekula, and Liang Zhao. 2021. GraphGT: Machine Learning Datasets for Deep Graph Generation and Transformation. (2021).
- [11] J. Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. *international conference on learning representations* (2015).
- [12] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. *KDD* (2016), 855–864.
- [13] Xiaojie Guo, Yuanqi Du, and Liang Zhao. 2021. Deep Generative Models for Spatial Networks. In *KDD*. 505–515. <https://doi.org/10.1145/3447548.3467394>
- [14] Kaveh Hassani and Amir Hosein Khasahmadi. 2020. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*. PMLR, 4116–4126.
- [15] Chih-Hui Ho and Nuno Vasconcelos. 2020. Contrastive Learning with Adversarial Examples. *NIPS 2020* (2020).
- [16] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2020. Strategies for Pre-training Graph Neural Networks. *ICLR* (2020).
- [17] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. 2020. GPT-GNN: Generative Pre-Training of Graph Neural Networks. *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining Virtual Event CA USA July, 2020* (2020), 1857–1867.
- [18] Ming Jin, Yizhen Zheng, Yuan-Fang Li, Chen Gong, Chuan Zhou, and Shirui Pan. 2021. Multi-Scale Contrastive Siamese Networks for Self-Supervised Graph Representation Learning. *IJCAI* (2021), 1477–1483.
- [19] Nikola Jovanović, Zhao Meng, Lukas Faber, and Roger Wattenhofer. 2021. Towards robust graph contrastive learning. *arXiv preprint arXiv:2102.13085* (2021).
- [20] He Kaiming, Fan Haoqi, Wu Yuxin, Xie Saining, and Girshick Ross. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. *CVPR* (2020), 9726–9735.
- [21] N. Thomas Kipf and Max Welling. 2016. Variational Graph Auto-Encoders. *CoRR* (2016).
- [22] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [23] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *ICLR* (2020).
- [24] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2019. Towards Deep Learning Models Resistant to Adversarial Attacks. *international conference on learning representations* (2019).
- [25] A. David McAllester. 1999. PAC-Bayesian model averaging. *COLT* (1999), 164–170.
- [26] A. David McAllester and Jonathan Baxter. 1999. Some PAC-Bayesian Theorems. *Machine Learning* (1999), 355–363.
- [27] Christopher Morris, M. Nils Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. (2020).
- [28] Annamalai Narayanan, Mahinthan Chandramohan, Rajasekar Venkatesan, Lihui Chen, Yang Liu, and Shantanu Jaiswal. 2017. graph2vec: Learning Distributed Representations of Graphs. *arXiv: Artificial Intelligence* (2017).
- [29] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. 2017. Exploring Generalization in Deep Learning. *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 30 (NIPS 2017)* (2017), 5947–5956.
- [30] van den Aäron Oord, Yazhe Li, and Oriol Vinyals. 2019. Representation Learning with Contrastive Predictive Coding. *arXiv: Learning* (2019).
- [31] Zhao Pu, Chen Pin-Yu, Das Payel, Karthikeyan Ramamurthy Natesan, and Lin Xue. 2020. Bridging Mode Connectivity in Loss Landscapes and Adversarial Robustness. *ICLR* (2020).
- [32] Kaspar Riesen and Horst Bunke. 2008. IAM Graph Database Repository for Graph Based Pattern Recognition and Machine Learning. *SSPR/SPR* (2008), 287–297.
- [33] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. 2021. Contrastive Learning with Hard Negative Samples. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=CR1XOQ0UTH>
- [34] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying WEI, Wenbing Huang, and Junzhou Huang. 2020. Self-Supervised Graph Transformer on Large-Scale Molecular Data. *NIPS 2020* (2020).
- [35] Nino Shervashidze, Pascal Schweitzer, Jan van Erik Leeuwen, Kurt Mehlhorn, and M. Karsten Borgwardt. 2011. Weisfeiler-Lehman Graph Kernels. *Journal of Machine Learning Research* (2011), 2539–2561.
- [36] Nino Shervashidze, V. N. S. Vishwanathan, H. Tobias Petri, Kurt Mehlhorn, and M. Karsten Borgwardt. 2009. Efficient graphlet kernels for large graph comparison. *AISTATS* (2009), 488–495.
- [37] Kihyuk Sohn. 2016. Improved Deep Metric Learning with Multi-class N-pair Loss Objective. *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 29 (NIPS 2016)* (2016), 1849–1857.
- [38] Fan-Yun Sun, Jordan Hoffman, Vikas Verma, and Jian Tang. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. *ICLR* (2020).
- [39] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. 2019. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. *arXiv preprint arXiv:1908.01000* (2019).
- [40] Mengying Sun, Jing Xing, Huijun Wang, Bin Chen, and Jiayu Zhou. 2021. MoCL: Contrastive Learning on Molecular Graphs with Multi-level Domain Knowledge. *KDD 2021* (2021).
- [41] Shantanu Thakoor, Corentin Tallec, Mohammad Gheshlaghi Azar, Remi Munos, Petar Velicković, and Michal Valko. 2021. Bootstrapped Representation Learning on Graphs. In *ICLR 2021 Workshop on Geometrical and Topological Representation Learning*. <https://openreview.net/forum?id=QrzVRAA49Ud>
- [42] Chen Ting, Kornblith Simon, Norouzi Mohammad, and Hinton Geoffrey. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *ICML* (2020), 1597–1607.
- [43] Wang Tongzhou and Isola Phillip. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. *ICML* (2020), 9929–9939.
- [44] Vinay Prabhudev, Dian Yap Ang, Xu Joyce, and Whaley John. 2019. Understanding Adversarial Robustness Through Loss Landscape Geometries. (2019).
- [45] Petar Velickovic, William Fedus, L. William Hamilton, Pietro Liò, Yoshua Bengio, and Devon R. Hjelm. 2019. Deep Graph Infomax. *ICLR* (2019).
- [46] Zhirong Wu, Yuanjun Xiong, X Stella Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), 3733–3742.
- [47] Jun Xia, Haitao Lin, Yongjie Xu, Lirong Wu, Zhangyang Gao, Siyuan Li, and Stan Z. Li. 2021. Towards Robust Graph Neural Networks against Label Noise. https://openreview.net/forum?id=H38f_9b90BO
- [48] Jun Xia, Lirong Wu, Jintao Chen, Ge Wang, and Stan Z. Li. 2021. Debaised Graph Contrastive Learning. *CoRR* abs/2110.02027 (2021). [arXiv:2110.02027](https://arxiv.org/abs/2110.02027)
- [49] JUN XIA, Jiangbin Zheng, Cheng Tan, Ge Wang, and Stan Z Li. 2022. Towards Effective and Generalizable Fine-tuning for Pre-trained Molecular Graph Models. *bioRxiv* (2022). <https://doi.org/10.1101/2022.02.03.479055> [arXiv:https://www.biorxiv.org/content/early/2022/02/06/2022.02.03.479055.full.pdf](https://www.biorxiv.org/content/early/2022/02/06/2022.02.03.479055.full.pdf)
- [50] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=ryGs6iA5Km>
- [51] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks? *international conference on learning representations* (2019).
- [52] Pinar Yanardag and V. N. S. Vishwanathan. 2015. Deep Graph Kernels. *ACM Knowledge Discovery and Data Mining* (2015), 1365–1374.
- [53] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. 2021. Graph Contrastive Learning Automated. *arXiv preprint arXiv:2106.07594* (2021).
- [54] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph Contrastive Learning with Augmentations. 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., 5812–5823. <https://proceedings.neurips.cc/paper/2020/file/3fe230348e9a12c1320749e3f9fa4cd-Paper.pdf>

- [55] Jiaqi Zeng and Pengtao Xie. 2020. Contrastive self-supervised learning for graph classification. *arXiv* (2020).
- [56] Yanqiao Zhu, Yichen Xu, Qiang Liu, and Shu Wu. 2021. An Empirical Study of Graph Contrastive Learning. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, Joaquin Vanschoren and Serena Yeung (Eds.). Curran Associates, Inc.
- [57] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2020. Deep Graph Contrastive Representation Learning. In *ICML Workshop on Graph Representation Learning and Beyond*. <https://arxiv.org/abs/2006.04131>
- [58] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2021. Graph Contrastive Learning with Adaptive Augmentation. *WWW* (2021), 2069–2080.

A APPENDIX: DATASETS IN VARIOUS SETTINGS

A.1 Unsupervised learning & Semi-supervised learning

Table 7: Datasets statistics for unsupervised and semi-supervised experiments.

Datasets	Category	Graph Num.	Avg. Node	Avg. Degree
NCI1	Biochemical Molecules	4110	29.87	1.08
PROTEINS	Biochemical Molecules	1113	39.06	1.86
DD	Biochemical Molecules	1178	284.32	715.66
MUTAG	Biochemical Molecules	188	17.93	19.79
COLLAB	Social Networks	5000	74.49	32.99
RDT-B	Social Networks	2000	429.63	1.15
RDB-M	Social Networks	2000	429.63	497.75
IMDB-B	Social Networks	1000	19.77	96.53

For unsupervised setting, experiments are performed for 5 times each of which corresponds to a 10-fold evaluation, with mean and standard deviation of accuracies (%) reported. For semi-supervised learning, we perform experiments with 1% (if there are over 10 samples for each class) and 10% label rate for 5 times, each of which corresponds to a 10-fold evaluation, with mean and standard deviation of accuracies (%) reported. For pre-training, learning rate is tuned in {0.1, 1.0, 5.0, 10.0} and epoch number in {20, 40, 60, 80, 100} where grid search is performed. All datasets used in both unsupervised and semi-supervised experiments can be seen in Table 7.

A.2 Transfer learning

Table 8: Datasets statistics for transfer learning.

Datasets	Category	Utilization	Graph Num.	Avg. Node	Avg. Degree
ZINC-2M	Biochemical Molecules	Pre-Training	2,000,000	26.62	57.72
PPI-306K	Protein-Protein Intersection Networks	Pre-Training	306,925	39.82	729.62
BBBP	Biochemical Molecules	Finetuning	2,039	24.06	51.90
ToxCast	Biochemical Molecules	Finetuning	8,576	18.78	38.52
SIDER	Biochemical Molecules	Finetuning	1,427	33.64	70.71

The datasets utilized in transfer learning can be seen in Table 8. ZINC-2M and PPI-306K are used for pre-training and the left ones are for fine-tuning.

B GNN ARCHITECTURES IN VARIOUS SETTINGS

To keep fair, we adopt the same GNNs architectures with previous competitors. Specifically, for unsupervised task, GIN [51] with 3 layers and 32 hidden dimensions is adopted as the encoder. For semi-supervised task, we utilize ResGCN [4] with 5 layers and 128 hidden dimensions. For transfer learning, we adopt GIN with the default setting in [16] as the GNN-based encoder. For experiments on adversarial robustness, Structure2vec is adopted as the GNN-based encoder as in [7].