

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

1-2024

Affinity uncertainty-based hard negative mining in graph contrastive learning

Chaoxi NIU

Guansong PANG

Singapore Management University, gspang@smu.edu.sg

Ling CHEN

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), and the [OS and Networks Commons](#)

Citation

NIU, Chaoxi; PANG, Guansong; and CHEN, Ling. Affinity uncertainty-based hard negative mining in graph contrastive learning. (2024). *IEEE Transactions on Neural Networks and Learning Systems*. 1-11.
Available at: https://ink.library.smu.edu.sg/sis_research/8611

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 cherylids@smu.edu.sg.

Affinity Uncertainty-based Hard Negative Mining in Graph Contrastive Learning

Chaoxi Niu, Guansong Pang, *Member, IEEE*, Ling Chen, *Senior Member, IEEE*

Abstract—Hard negative mining has shown effective in enhancing self-supervised contrastive learning (CL) on diverse data types, including graph contrastive learning (GCL). Existing hardness-aware CL methods typically treat negative instances that are most similar to the anchor instance as hard negatives, which helps improve the CL performance, especially on image data. However, this approach often fails to identify the hard negatives but leads to many false negatives on graph data. This is mainly due to that the learned graph representations are not sufficiently discriminative due to over-smooth representations and/or non-i.i.d. issues in graph data. To tackle this problem, this paper proposes a novel approach that builds a discriminative model on *collective affinity* information (i.e., two sets of pairwise affinities between the negative instances and the anchor instance) to mine hard negatives in GCL. In particular, the proposed approach evaluates how confident/uncertain the discriminative model is about the affinity of each negative instance to an anchor instance to determine its hardness weight relative to the anchor instance. This uncertainty information is then incorporated into existing GCL loss functions via a weighting term to enhance their performance. The enhanced GCL is theoretically grounded that the resulting GCL loss is equivalent to a triplet loss with an *adaptive* margin being exponentially proportional to the learned uncertainty of each negative instance. Extensive experiments on 10 graph datasets show that our approach i) consistently enhances different state-of-the-art GCL methods in both graph and node classification tasks, and ii) significantly improves their robustness against adversarial attacks. Code is available at <https://github.com/mala-lab/AUGCL>.

Index Terms—Graph contrastive learning, Hard negative mining, Uncertainty estimation, Affinity learning.

I. INTRODUCTION

GRAPH is ubiquitous and plays an important role in various fields [1]–[3], such as social networks, bioinformatics, chemistry, etc. Due to its non-Euclidean nature, learning expressive graph representations is one crucial foundation of different graph mining tasks, such as graph classification and node classification. In recent years, graph neural networks (GNNs) have become predominant in achieving this goal. Most existing GNNs focus on supervised or semi-supervised learning settings [4]–[6], where class label information is required for training the GNNs. However, obtaining such information is hard or costly, especially for graph data which is at large scale and/or demands strong domain knowledge

to accurately perform the data annotation. Recently, self-supervised learning of GNNs [7], [8] which can learn graph representations without accessing ground truth labels was introduced to tackle this issue and has attracted significant research interests.

Graph contrastive learning (GCL) has become one of the most popular self-supervised methods for graph representation learning [9]–[19]. It focuses on learning representations by maximizing the mutual information between augmentations of the same instance, in which the augmentations of the same graph/node are often treated as positive instances, with the other graphs/nodes as negative instances [7], [8].

Despite the impressive successes achieved by current GCL methods, their learning capability can be largely limited by the way they choose negative samples [20]–[23]. One commonly used negative selection approach is to randomly select negative instances from a sufficiently large batch or a memory bank, and then treat all negative instances equally in contrastive learning. However, this approach typically overlooks the negative instances that can provide more information for contrastive learning than the others. These informative negative instances are commonly referred to as *hard negative* instances. Many prior studies [20], [21], [23] have demonstrated the critical importance of hard negative instances over counterparts (e.g., easy negatives that are distant from the positive in both semantics and representations) in learning discriminative features and enabling fast convergence.

Many recent contrastive learning (CL) methods [20]–[22], [24], [25] thus incorporate hard negative mining methods into their training process to leverage these hard negative instances. These methods typically treat negative instances that are most similar to the anchor instance as the hard negatives [20]–[22], [24], [25]. Although achieving improved CL performance on image data, these hard negative mining approaches often perform poorly on graph data, as shown in some recent studies [23], [26] and our experiments. This is mainly because the learned graph representations are not sufficiently discriminative due to i) the non-i.i.d. (non-independent and identically distributed) nature of graph data, e.g., nodes with the same label tend to be densely connected in graph data, and ii) the over-smooth graph representations resulting from the iterative message passing mechanism. Consequently, for graph data, the most similar negatives to the anchor can be false negatives with high probability. To address this issue, the very recent method ProGCL [23] imposed a beta mixture model on the pairwise similarities between the negatives and the anchor to

Chaoxi Niu and Ling Chen are with the Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, NSW 2007, Australia. (email: Chaoxi.Niu@student.uts.edu.au; Ling.Chen@uts.edu.au).

Guansong Pang is with School of Computing and Information Systems, Singapore Management University, 178902, Singapore.

Corresponding author: G. Pang (email: pangguansong@gmail.com).

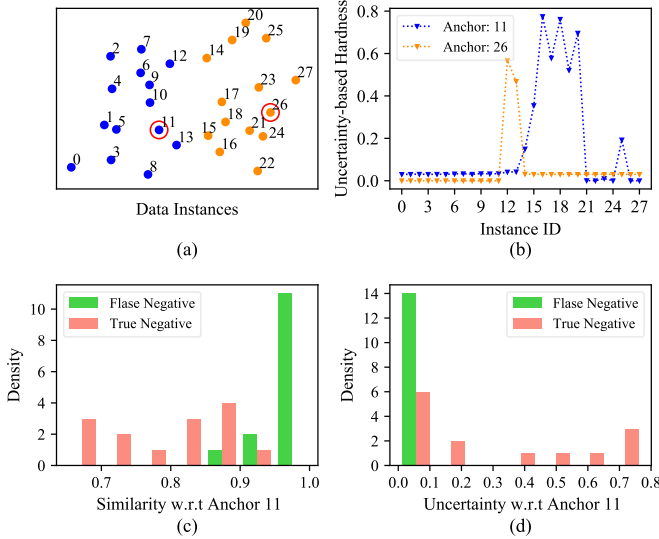


Fig. 1. (a): Two groups of data instances in blue and orange. (b): The affinity uncertainty-based hardness results learned by our approach using instance 11 or 26 as the anchor instance. Instances with a larger uncertainty are more likely to be hard negative samples w.r.t. the anchor instance. (c): The histograms of the similarity of the instances to the anchor instance 11. It is clear that treating the most similar instances to the anchor as the hard negatives can lead to many false negatives. (d): The uncertainty results learned by our approach for the instances w.r.t the anchor instance 11, where true negatives including hard negatives have large uncertainty values (and thus large hardness weights) while false negative cases receive very small uncertainty values.

estimate the probability of a negative being true one, and it subsequently combined the estimated probability and the pairwise similarity to measure the hardness of the negatives. The method relies on the prior that the similarity distribution of negatives w.r.t. positive is bimodal and works well in node classification tasks. It fails to work when its prior is not fully met. As shown in our experiments (Table III in Sec. IV-B), such failure cases occur in most graph classification datasets where ProGCL has very marginal improvement, or even worse performance, compared to the original GCL methods. Besides, AdaS [27] emphasized the contrastive learning on the negatives that are neither hard nor easy to discriminate from the anchor. HSAN [28] exploited hard samples by utilizing high-confidence clustering and comprehensive similarity information. However, HSAN is primarily proposed for node clustering. This work aims to eliminate the need to specify the similarity distribution prior as in ProGCL for better applicability in practice.

To this end, this paper introduces a novel Affinity Uncertainty-based hard negative mining for Graph Contrastive Learning, termed as AUGCL, to tackle this problem. AUGCL learns a data-driven, affinity-based uncertainty estimator to evaluate the hardness of negative instances relative to each anchor instance, meaning that the hardness of an instance is dependent on the given anchor instance, as shown by an example in Fig. 1(a-b). Particularly, AUGCL builds a discriminative model on *collective affinity* information (i.e., two sets of pairwise affinities between the negative instances and the anchor instance) to evaluate how confident/uncertain the discriminative model is about the affinity of each negative

instance to the anchor instance. Instances that have a larger affinity uncertainty would be more likely to be hard negatives, and they are subsequently assigned with a larger hard-negative weight to receive more attention from the GCL models. By doing so, AUGCL learns discriminative affinity uncertainties for the negative instances relative to each anchor instance, as shown by the results of the anchor instance 11 in Fig. 1(b) and (d), where small and large uncertainty-based hardness values are assigned to false negatives and true negatives, respectively. By contrast, the current similarity-based methods that regard the most similar negative instances to the anchor instance as hard negatives fail to identify the truly hard negatives but lead to many false negatives, as shown in Fig. 1(c). Those learned hardness results can then be seamlessly incorporated into popular GCL models (e.g., InfoNCE-based models [29]) as a hardness weight to enhance their performance. AUGCL addresses a similar issue as ProGCL, but it eliminates the prior information posited in ProGCL, enabling AUGCL to work more effectively on diverse node-level and graph-level datasets.

In summary, this work makes the following three main contributions.

- We propose a novel approach AUGCL that utilizes the modeling of collective affinities to take account of the non-i.i.d. and over-smooth representations issues in graph data via the learning of an uncertainty-based hardness measure. To the best of our knowledge, it is the first work that addresses the problem using an uncertainty learning framework.
- We show theoretically that our approach transforms popular GCL losses such as InfoNCE into a triplet loss with an adaptive hardness-based margin, enforcing a large margin for hard negatives while pulling false negatives close to anchor instances.
- Extensive experiments on 10 graph datasets demonstrate the superiority of AUGCL in consistently enhancing different state-of-the-art GCL methods in both graph and node classification tasks (having maximal classification accuracy improvement by $\sim 2\%$ and $\sim 1.5\%$, respectively), and the robustness against graph adversarial attacks (maximal improvement by $\sim 8\%$).

II. RELATED WORKS

A. Graph Contrastive Learning

Recently, contrastive learning [29]–[32] has become a prominent technique in self-supervised learning. It has been successfully adapted into diverse domains, including the graph domain. A number of GCL methods [9]–[15], [33]–[36] have been proposed. DGI [9] is an early attempt that obtained node representations by maximizing the mutual information between node embeddings and high-level graph information. MVGRL [10] improved DGI by introducing different structural views to learn node and graph-level representations. InfoGraph [11] performed contrastive learning by directly maximizing the consistency between sampled subgraphs and pooled graph representations. DiGCL [33] performed direct

graph contrastive learning based on laplacian perturbation. Additionally, GraphCL [12] systematically explored the influence of different augmentations on graph-level contrastive learning. GRACE [34] designed a node-level contrastive framework by maximizing the agreement of node embeddings between two corrupted graph views. GCA [13] conducted node-level contrastive learning with adaptive augmentation on the topology and node attribute level. G-Zoom [19] learned node representation by establishing contrastiveness on different scales progressively. PA-GCL [37] explored channel-level contrastive learning among three generated proximity graphs and updated the generated graphs during the training process to improve their contrast effect. Besides, some studies have been put forward to enhance the GCL by automating data augmentations [14] or discarding explicit data augmentations [15] [36]. InfoGCL [35] further built principles to build the optimal GCL model based on Information Bottleneck [38]. The main differences among these methods lie in the way they obtain positive pairs. By contrast, our approach AUGCL focuses on hard negative mining, which is orthogonal to these GCL methods and can be plugged into their loss function to improve their performance on graph/node-level tasks.

B. Hard Negative Mining in Contrastive Learning

Hard negative mining refers to generating or mining the negatives that are difficult to discriminate from the positive. Various methods have been proposed to perform hard negative mining to facilitate contrastive learning, including employing mixup strategy [39] to mix the anchor instance and negative instance to synthesize hard negatives [20], [25], [40], [41], and developing unsupervised sampling methods for selecting hard negative samples [21], [22]. Specifically, DCL [22] adopted positive-unlabeled learning [42] to avoid sampling false negatives. HCL [21] improved DCL [22] by further up-weighting the negatives close to the anchor. These methods are mainly focused on image data and they often treat negative instances that are most similar to the anchor instance as the hard negatives. However, for graph data, the similar negatives could be false negatives relative to the anchor, and the GCL performance would be degraded by employing these hard negative mining methods [23], [26]. To address this issue, ProGCL [23] exploited a two-component beta mixture model on the similarity distribution of negatives to estimate the probability of negative instances being true for an anchor and then measured the hardness of negative instances by integrating the estimated posterior probability and the similarity between the negative and the anchor. AdaS [27] devised an adaptive sampling strategy to emphasize the negatives that are neither too hard nor too easy to discriminate from the anchor, and employed a polarization regularizer to enhance the influence of them. HSNAN [28] explored hard positives and negatives by introducing a weight-modulating function. Specifically, this function adjusts the weights of sample pairs in contrastive learning by calculating the disparity between high-confidence pseudo labels and the similarity obtained from both attribute and structure information. Instead of measuring the hardness of negatives, HCHSM [43] selected hard samples

based on a mutual information agreement gap between positive and negative pairs and performed hierarchically contrastive learning for hard samples to exploit multilevel intrinsic graph features. Similar to ProGCL, our method also measures the hardness of negatives for each anchor instance. However, we employ the uncertainty estimation model to directly learn the negative instance hardness. The learned hardness is then incorporated into the contrastive loss via a weighting term, resulting in an anchor-instance-adaptive contrastive learning framework with good theoretical support.

C. Uncertainty Estimation

Numerous methods and theories have been introduced to measure the prediction uncertainty, e.g., by using the maximum of predicted probabilities [44]–[46], the prediction entropy/energy [45], [47]–[49], or an extra (void/background) class [49]–[51]. These methods focus on calibrating prediction confidence in supervised learning, whereas we utilize uncertainty estimation under the self-supervised setting to empower contrastive learning. Our work is motivated by the observation that hard samples are typically the instances at the decision boundary between the positive and negative instances, which are also the samples that learning models are uncertain about. Thus, uncertainty estimation offers an effective way to measure the hardness of negative instances. To be applicable in graph contrastive learning, AUGCL is designed in a novel way by using an anchor-instance-dependent uncertainty learning approach.

III. AUGCL: AFFINITY UNCERTAINTY-BASED GRAPH CONTRASTIVE LEARNING

A. Preliminaries

Self-supervised graph representation learning has demonstrated promising performance in empowering diverse graph learning tasks. This work focuses on node-level and graph-level tasks. Particularly, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote a graph where \mathcal{V} and \mathcal{E} denote the set of nodes and edges respectively, then for a node-level task, the goal of self-supervised graph representation learning is to leverage a single graph \mathcal{G} to learn an encoder $\psi(\mathcal{V}, \mathcal{E})$ without using the labels of nodes so that $\psi(\mathcal{V}, \mathcal{E})$ can yield an expressive low-dimensional embedding z_i for each node in \mathcal{V} . The resulting node embeddings $\mathcal{Z} = \{z_i\}_{i=1}^{|\mathcal{V}|}$ can then be used in various downstream node-level tasks, such as node classification. For a graph-level task, the goal instead is to learn a graph encoder $\psi(\mathcal{V}_i, \mathcal{E}_i)$ given a set of N graphs $\{\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i)\}_{i=1}^N$, where the encoder $\psi(\mathcal{V}_i, \mathcal{E}_i)$ outputs a low-dimensional embedding z_i for each graph \mathcal{G}_i , and the graph embeddings $\mathcal{Z} = \{z_i\}_{i=1}^N$ can then be used in various downstream graph-level tasks, e.g., graph classification. Our approach can be used to improve the self-supervised learning of graph representations and node representations, as shown in Sec. IV. Without loss of generality, we use the graph-level tasks to introduce our approach below.

B. Popular GCL Methods and Their Weaknesses

Graph contrastive learning is one of the most popular approaches for self-supervised graph representation learning.

As an instance-wise discriminative approach, it aims to pull two different augmentations of the same graph closer and push augmentations of different graphs apart [10], [12]. InfoNCE [29] is among the most popular contrastive learning loss functions to achieve this goal. Specifically, given a minibatch of randomly sampled graphs $\{\mathcal{G}_i\}_{i=1}^N$, two augmentation functions t_1 and t_2 are first sampled from the augmentation pool \mathcal{T} which consists of all possible augmentations. Then, two graph views $\{\tilde{\mathcal{G}}_i\}_{i=1}^N$ and $\{\hat{\mathcal{G}}_i\}_{i=1}^N$ are generated by applying t_1, t_2 to each graph. The embeddings $\{\tilde{z}_i\}_{i=1}^N$ and $\{\hat{z}_i\}_{i=1}^N$ of the augmented graphs are obtained by feeding the augmented graphs into a shared GNN encoder $\psi(\cdot)$, followed by a projection head (2-layer perceptron) [31]. For an anchor instance \mathcal{G}_i – a graph augmented from \mathcal{G}_i using t_1 , the positive is $\hat{\mathcal{G}}_i$ – a graph augmented from the same graph \mathcal{G}_i but using a different augmentation t_2 , while the source of the negative instances is $\{\tilde{\mathcal{G}}_j\}_{j=1}^N$, from which negative instances are sampled. To enforce the maximization of the consistency between positive embeddings, the pairwise objective for a positive pair (\tilde{z}_i, \hat{z}_i) is formulated as:

$$\ell_{\text{InfoNCE}}(\tilde{z}_i, \hat{z}_i) = -\log \frac{e^{(h(\tilde{z}_i, \hat{z}_i)/\tau)}}{e^{(h(\tilde{z}_i, \hat{z}_i)/\tau)} + \sum_{j, j \neq i}^N e^{(h(\tilde{z}_i, \hat{z}_j)/\tau)}}, \quad (1)$$

where τ denotes the temperature parameter and $h(\tilde{z}_i, \hat{z}_j)$ is the cosine similarity function measuring similarity between \tilde{z}_i and \hat{z}_j .

Although these graph contrastive learning methods have achieved great success in graph representation learning, they often fail to consider the semantics of negatives in $\{\tilde{\mathcal{G}}_j\}_{j=1}^N$. Consequently, instances that share the same semantics with the positive can be sampled and treated as negatives in (1). This false negative sampling issue, also known as sampling bias in [22], would hinder the learning of contrastive representations between positive instances and negative instances. More importantly, the contrastive learning cannot exploit *hard negatives*, i.e., instances that are similar to but semantically different from the anchor, which are driving force for contrastive learning to learn substantially enhanced discriminative representations, as shown in the literature empirically and theoretically [20], [21], [23].

C. The Proposed Approach AUGCL

To address the negative sampling weaknesses discussed in Sec. III-B, we present a novel framework for learning an affinity uncertainty-based hardness measure for enhancing current state-of-the-art graph contrastive learning methods. The key idea is to first learn the hardness of a negative instance relative to each anchor instance by comparing the affinity between them to the affinities of the anchor instance to the other instances. The hardness results can then be plugged into a contrastive loss, e.g., InfoNCE, to improve the effectiveness of current GCL methods in utilizing the hard negatives.

Overview of AUGCL. Since the hardness of a negative instance varies largely w.r.t. different anchor instances, our approach AUGCL aims to learn a hardness measure based on the relative affinity between the negative instance and each

anchor instance. That is, for an anchor instance \tilde{z}_i and its negative instance candidate set $\hat{\mathcal{Z}}_i = \{\hat{z}_j\}_{j=1}^N$, we learn a single hardness measure function $\phi(\hat{z}_j|\tilde{z}_i; \Theta) : \hat{\mathcal{Z}}_i \rightarrow \mathbb{R}$ that yields a hardness value for each $\hat{z} \in \hat{\mathcal{Z}}_i$ relative to \tilde{z}_i . Note that the function ϕ parameterized by Θ is trained across all anchor instances; yet the hardness it yields for the negative instance \hat{z}_j is dependent on the anchor \tilde{z}_i . For brevity, $\phi(\hat{z}_j|\tilde{z}_i; \Theta)$ is denoted as $\phi_i(\hat{z}_j; \Theta)$ hereafter.

Unlike current hardness measures that define the hardness of a negative instance based on its individual relation to the anchor instance (e.g., the similarity between them), one key novelty in AUGCL is that it defines the hardness based on two groups of pairwise affinities between the negative instances and the anchor instance. More specifically, we introduce the concept of affinity uncertainty below to achieve this goal:

Definition 1 (Affinity Uncertainty). *Given an anchor instance \tilde{z}_i and its negative instance candidate set $\hat{\mathcal{Z}}_i = \{\hat{z}_j\}_{j=1}^N$, and let \mathcal{C}_1^i and \mathcal{C}_2^i be two disjoint groups of instances in $\hat{\mathcal{Z}}_i$ such that: one group \mathcal{C}_1^i includes the instances that are closely aligned and distributed around the anchor \tilde{z}_i , while the other group \mathcal{C}_2^i contains the rest of other instances; and $\hat{\mathcal{Z}}_i = \mathcal{C}_1^i \cup \mathcal{C}_2^i$. Then the affinity uncertainty of each $\hat{z} \in \hat{\mathcal{Z}}_i$ w.r.t. \tilde{z}_i is defined as:*

$$\phi_i(\hat{z}) = g(\hat{z}, \mathcal{C}_1^i, \mathcal{C}_2^i), \quad (2)$$

where g is an uncertainty estimator that evaluates how confident the estimator is about the affinity of \hat{z} to the instances in the anchor instance-centered group \mathcal{C}_1^i compared to the other group \mathcal{C}_2^i .

The affinity uncertainty in (2) takes a holistic approach that considers diverse affinities of the negative instances within and across the two groups \mathcal{C}_1^i and \mathcal{C}_2^i to learn an accurate hardness for each negative instance \hat{z} . As shown in the literature (e.g., [51]) and Fig. 1, instances that are ambiguous to distinguish are assigned to large uncertainty values. These instances typically have a poor affinity to both groups \mathcal{C}_1^i and \mathcal{C}_2^i , such as those located on the boundary between the two groups. By contrast, if the instances are coherently aligned within \mathcal{C}_1^i or \mathcal{C}_2^i , their uncertainty would be small. Thus, this type of uncertainty can be naturally used to define the hardness of the negative instances.

The obtained hardness can then be easily plugged into existing contrastive losses, such as the InfoNCE loss, via a weighting term for the negative instances. Particularly, the AUGCL-enhanced InfoNCE is given as follows:

$$\ell_{\text{AUGCL}}(\tilde{z}_i, \hat{z}_i) = -\log \frac{e^{(h(\tilde{z}_i, \hat{z}_i)/\tau)}}{e^{(h(\tilde{z}_i, \hat{z}_i)/\tau)} + \sum_{j, j \neq i}^N w_{ij} e^{(h(\tilde{z}_i, \hat{z}_j)/\tau)}}, \quad (3)$$

where $w_{ij} = \alpha \phi_i(\hat{z}_j; \Theta)$ is the hardness-based weight added to \hat{z}_j relative to \tilde{z}_i . $\phi_i(\hat{z}_j; \Theta)$ is the hardness learned by AUGCL for the negative instance \hat{z}_j w.r.t. the anchor instance \tilde{z}_i and α is a hyperparameter. This enables effective exploitation of the hard negatives, as large weights are expected for hard negatives while small weights are expected for the other instances, e.g., false negatives.

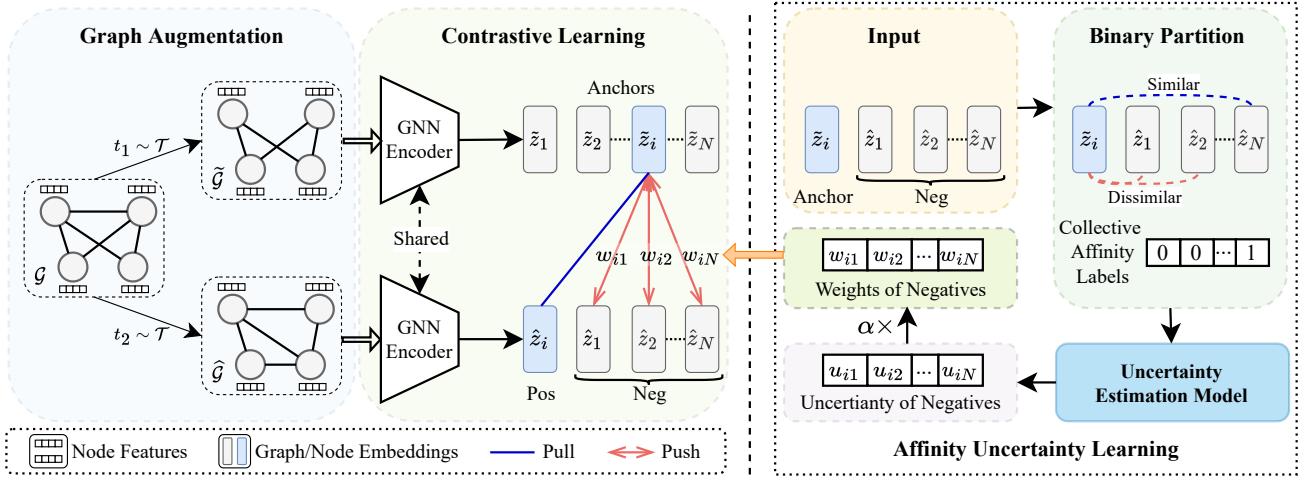


Fig. 2. Overview of our approach AUGCL. *Left*: AUGCL-based graph contrastive learning. The objective and the general procedures are the same as existing GCL methods, but AUGCL leverages affinity uncertainty to learn anchor-instance-dependent hardness-based instance weights $\{w_{i1}, w_{i2}, \dots, w_{iN}\}$ for all negative instances to improve existing GCL methods. *Right*: The proposed affinity uncertainty learning approach to obtain the weights. For an anchor \tilde{z}_i , AUGCL first obtains collective affinity information (i.e., pairwise affinity across the instances) via binary partition of its negative instances. It then utilizes those affinity information to learn an uncertainty estimator that evaluates how confident the estimator is about the affinity of each negative instance \hat{z}_j relative to the anchor instance \tilde{z}_i . A larger affinity uncertainty value u_{ij} indicates more likely of \hat{z}_j being a hard negative, and thus, a larger weight w_{ij} ($w_{ij} = \alpha u_{ij}$ where α is a hyperparameter).

The overall procedure of AUGCL is illustrated in Fig. 2. It follows the standard graph contrastive learning in the graph augmentation and contrastive learning except that we incorporate the affinity uncertainty-based hardness through a weighting term into the contrastive loss as in (3). The right panel in Fig. 2 shows the steps of learning an anchor-dependent hardness measure ϕ for each anchor \tilde{z}_i , consisting of instance partition and uncertainty estimation as indicated in Def. 1. Before introducing the details of these two components in Sec. III-D, below we demonstrate the theoretical motivation of the proposed method.

Theoretical Motivation. We show below that (3) is equivalent to a triplet loss with an adaptive margin exponentially proportional to the learned hardness-based weight $\phi_i(\hat{z}_j; \Theta)$. This provides a more straightforward explanation of the working mechanism of the proposed weighting method.

Theorem 1. Let $u_{ij} = \phi_i(\hat{z}_j; \Theta)$ be the affinity uncertainty-based hardness of a negative instance \hat{z}_j w.r.t. the anchor instance \tilde{z}_i . When the projection function is an identity function and assumes the positive instance is more similar to the anchor than the negative instances, then minimizing the proposed objective in (3) is equivalent to minimizing a modified triplet loss with an adaptive margin $m_{ij} = \frac{\tau}{2} \log(\alpha u_{ij})$, i.e.,

$$\ell_{\text{AUGCL}}(\tilde{z}_i, \hat{z}_i) \propto \frac{1}{2\tau} \sum_{j, j \neq i}^N \left(\|\tilde{z}_i - \hat{z}_i\| - \|\tilde{z}_i - \hat{z}_j\| + m_{ij} \right), \quad (4)$$

where \tilde{z}_i is the normalized embedding.

The proof for this theorem is detailed in the Appendix. From the theorem, we can see that the optimal embeddings to (4) should satisfy the following inequality:

$$\|\tilde{z}_i - \hat{z}_i\| \ll \|\tilde{z}_i - \hat{z}_j\| - m_{ij}, \quad (5)$$

where $m_{ij} = \frac{\tau}{2} \log(\alpha u_{ij})$ and $u_{ij} = \phi_i(\hat{z}_j; \Theta)$. Thus, m_{ij} is equivalent to a transformed affinity uncertainty-based hardness of the negative instance \hat{z}_j relative to the anchor \tilde{z}_i , satisfying:

$$\begin{cases} m_{ij} \geq 0, & \text{if } \alpha u_{ij} \geq 1; \\ m_{ij} < 0, & \text{otherwise.} \end{cases} \quad (6)$$

If \hat{z}_j is a hard negative for \tilde{z}_i , the large uncertainty u_{ij} between \tilde{z}_i and \hat{z}_j makes the inequality (5) hard to satisfy through $m_{ij} > 0$, enforcing better representation learning. On the contrary, if the uncertainty u_{ij} is small, (5) can be easily satisfied with $m_{ij} \ll 0$, reducing the impact of the possible false negative instances.

D. Instantiation of AUGCL

We introduce an instantiation of our AUGCL framework in this subsection. As demonstrated in Def. 1, the affinity uncertainty-based hardness function ϕ parameterized with Θ can be decomposed into two modules, including a binary clustering function $f: \{\hat{z}_j\}_{j=1}^N \rightarrow \{0, 1\}$ parameterized by Θ_f and an uncertainty estimation function $g: \{\hat{z}_j\}_{j=1}^N \times \{0, 1\} \rightarrow \mathbb{R}$ parameterized by Θ_g , i.e., $\Theta = \{\Theta_f, \Theta_g\}$. AUGCL is a generic framework. Different clustering and uncertainty estimation methods can be adopted in AUGCL to implement a specific model, as shown by our empirical results in Sec. IV-D. Below we describe the two modules of the best instantiated AUGCL model based on our experiments.

1) *Anchor-dependent Binary Partition of Negatives*: Given an anchor \tilde{z}_i , binary clustering is used to partition the negative samples into two coherent groups – C_1^i and C_2^i – for subsequent affinity uncertainty estimation. Without having access to label information, clustering is often adopted on the full dataset to divide instances into several clusters [52]–[58], and instances from clusters other than the anchor-based cluster are directly treated as negatives.

Our clustering differs from these existing methods in two main ways. First, we perform an anchor-dependent binary partition on only the negative instances in each batch of instances rather than the full dataset. Specifically, given a batch of node/graph embeddings $\{\hat{z}_j\}_{j=1}^N$, for each anchor $\tilde{z}_i \in \{\tilde{z}_i\}_{i=1}^N$, we perform a binary partition on the negative instance candidates $\{\hat{z}_j\}_{j=1}^N$ using an existing clustering method (e.g., k -means), i.e., $f_{k\text{-means}} : \{\hat{z}_j\}_{j=1}^N \rightarrow \{0, 1\}$. As a result, the negative instances are partitioned into two clusters. We denote the cluster distributed around the anchor instance \tilde{z}_i as C_1^i and assign the label of 1 to the negative instances within C_1^i , i.e., $C_1^{ij} = 1$ if $\hat{z}_j \in C_1^i$. Simultaneously, we refer to the other cluster as C_2^i and assign the label of 0 to the negative instances within it.

The second difference is that the obtained partitions are used to gain a sense of the affinity of an instance to the other instances, rather than being the direct negative sample clusters. The affinity information would be used to evaluate the hardness of each negative instance through an uncertainty estimation model in AUGCL.

2) *Affinity Uncertainty Estimation*: For an anchor \tilde{z}_i , the binary cluster labels $\{C_{12}^{ij}\}_{j=1}^N$ carry the affinity semantics of the instances $\{\hat{z}_j\}_{j=1}^N$ w.r.t. the anchor instance \tilde{z}_i . We further propose to perform an uncertainty estimation upon these affinity semantic-based labels for each anchor $\tilde{z}_i \in \{\tilde{z}_i\}_{i=1}^N$, and use this uncertainty to measure the hardness of instances $\{\hat{z}_j\}_{j=1}^N$. By doing so, a large uncertainty-based hardness is assigned to fringe instances that are located around the boundary between the two clusters; these instances are typically hard negatives w.r.t. \tilde{z}_i . A small hardness is assigned otherwise.

Different uncertainty estimation methods can be used to specify this component. We found that the recently proposed method Deep Gambler (DG) [51] worked best in our experiments, so DG is used in AUGCL by default. Specifically, DG extends a multi-class classification task to a problem that learns an extra class to represent the uncertainty of instances, in addition to guaranteeing the classification of the original classes. For an anchor instance \tilde{z}_i , given its associated negative instance candidates $\{\hat{z}_j\}_{j=1}^N$ and their affinity labels $\{C_{12}^{ij}\}_{j=1}^N$, the DG-based uncertainty estimation is trained by minimizing the following loss:

$$\ell_i^{DG} = - \sum_j^N \log(p_{C_{12}^{ij}} * o + u_{ij}), \quad (7)$$

where $p_{C_{12}^{ij}}$ is the predicted class probability on class C_{12}^{ij} from a multi-layer perceptrons-based (MLP-based) DG model $g(\hat{z}, C_1^i, C_2^i; \Theta_g)$ parameterized by Θ_g , u_{ij} is the uncertainty that the model g generates for the instance \hat{z}_j . o is the reward parameter which is an important and necessary parameter for DG. A larger o encourages DG to be more confident in inferring, i.e., producing a lower uncertainty value, and vice versa. The final loss of DG is computed across all anchor instances $\tilde{z}_i \in \{\tilde{z}_i\}_{i=1}^N$.

After the DG model is trained, for each anchor \tilde{z}_i , we calculate u_{ij} for its negative instance \hat{z}_j and obtain a uncertainty matrix $\mathbf{U} \in \mathbb{R}^{N \times (N-1)}$ where each row \mathbf{u}_i contains the

uncertainty of all negative instances w.r.t. the anchor \tilde{z}_i . These uncertainty values are then used in (3) to improve contrastive learning.

E. Time Complexity Analysis

We take k -means and Deep Gambler [51] as the partition and uncertainty estimation methods, respectively, to analyze the additional time complexity introduced by AUGCL. Specifically, let L be the number of MLP layers in DG and d be the number of hidden units for all layers. For the graph classification task, given a graph dataset with N graphs and the batch size is set as B , the time complexities of partition and the uncertainty modeling are $\mathcal{O}(2(\frac{N}{B})B^2T)$ and $\mathcal{O}(KL(\frac{N}{B})B^2d^2)$ respectively, where T is the number of iterations for k -means and K is the number of training epochs for the uncertainty estimation model. For the node classification task, given a graph with N nodes, in order to reduce the computation cost, we only sample M ($M \ll N$) negatives for an anchor when training AUGCL. The resulting time complexities of the partition and training uncertainty model are $\mathcal{O}(2NMT)$ and $\mathcal{O}(KLNMd^2)$ respectively. In experiments, we use the well-established k -means clustering implementation from scikit-learn [59], as it runs very fast in practice. Besides, the values of K , L , M and d are relatively small and the uncertainty estimation model is only trained once. Therefore, the computational overhead over the base model is not significant.

IV. EXPERIMENTS

A. Experimental Setup

1) *Datasets*: Seven commonly used graph classification datasets are used in our experiments. They come from two popular application domains: bioinformatics (MUTAG, DD, NCI1, and PROTEINS) and social networks (COLLAB, REDDIT-BINARY, and IMDB-BINARY). For the node classification task, we use three widely used datasets, i.e., Wiki-CS [60], Amazon-Computers, and Amazon-Photo [61]. Wiki-CS is a reference network constructed based on Wikipedia. Amazon-Computers and Amazon-Photo are two co-purchase networks constructed from Amazon. The statistics of the datasets are summarized in Table I.

TABLE I
KEY STATISTICS OF DATASETS USED.

Task	Dataset	Graphs	Avg.Nodes	Avg.Edges
Graph Classification	NCI1	4,110	29.87	32.30
	PROTEINS	1,113	39.06	72.82
	DD	1,178	284.32	715.66
	MUTAG	188	17.93	19.79
	COLLAB	5,000	74.49	2,457.78
	RD-B	2,000	429.63	497.75
Node Classification	IMDB-B	1,000	19.77	96.53
	Wiki-CS	1	11,701	216,123
	Aamazon-Computers	1	13,752	245,861
	Aamazon-Photo	1	7,650	119,081

2) *Implementation Details and Evaluation Protocol*: For the graph classification task, GraphCL [12], a recent SOTA InfoNCE-based contrastive learning method for graph classification, is used as our base, into which our affinity uncertainty-based hardness learning method is plugged. For a fair comparison, the network backbone, the graph augmentation methods

and the hyper-parameters of our AUGCL-enabled GraphCL are kept exactly the same as the original GraphCL. We follow a widely-used two-stage evaluation protocol in the literature [11], [12], [62], [63], in which we first learn graph representations in a self-supervised manner and then use the representations to train a downstream SVM classifier. The 10-fold evaluation is adopted in classification, and it is repeated five times with the mean accuracy (%) and standard variation reported.

For the node classification task, we adopt GCA [13] as the base model and plug our AUGCL-based affinity uncertainty hardness into it. The evaluation protocol for node classification follows DGI [9] where the model is first trained in an unsupervised manner and then the learned node representations are used to train and test a simple ℓ_2 -regularized logistic regression classifier. On each dataset, the experiment is repeated for 20 runs with different data splits, and the average classification accuracy, together with the standard variation, is reported.

For graph and node classification, we use the same architecture in our affinity uncertainty estimation model, i.e., a three-layer multi-layer-perceptrons (MLP) architecture, containing 128 units per layer with *ReLU* activation. We adopt the Stochastic Gradient Descent (SGD) optimizer for the uncertainty estimation model and the learning rate is set to 0.01 across all the datasets. The uncertainty scaling parameter α is set to the reciprocal of the mean of uncertainties. The training epoch number of the uncertainty estimation model is set to 10 for all datasets. For the reward parameter in the uncertainty estimation model, it is selected through a grid search, and the search space is $\{1.5, 1.6, 1.7, 1.8, 1.9\}$.

TABLE II
COMPARISON BETWEEN THE BASELINES AND THEIR AUGCL-ENABLED COUNTERPARTS. THE BASELINES ARE GRAPHCL [12] AND GCA [13] FOR GRAPH AND NODE CLASSIFICATION TASKS, RESPECTIVELY. “↑” REFERS TO THE IMPROVEMENT COMPARED TO THE BASELINES.

Task	Dataset	GraphCL	GraphCL _{AUGCL}
Graph Classification	NCI	78.26	80.16 (↑ 1.90)
	PROTEINS	74.36	75.76 (↑ 1.40)
	DD	78.01	80.14 (↑ 2.13)
	MUTAG	87.15	89.20 (↑ 2.05)
	COLLAB	71.53	72.10 (↑ 0.57)
	RDT-B	90.09	91.19 (↑ 1.10)
	IMDB-B	71.20	72.46 (↑ 1.26)
Task	Dataset	GCA	GCA _{AUGCL}
Node Classification	Wiki-CS	78.08	78.59 (↑ 0.51)
	Amazon-Computers	87.80	88.94 (↑ 1.14)
	Amazon-Photo	91.99	93.43 (↑ 1.44)

3) *Competing Methods*: We evaluate the effectiveness of AUGCL in both graph and node classification tasks. In both tasks, AUGCL is evaluated against three state-of-the-art hard negative mining-based contrastive learning methods, including DCL [22], HCL [21] and ProGCL [23]. For the node classification task, we further compare the proposed method with the recently proposed hard sample aware method HSAN [28]. While HSAN was initially designed for node clustering, we adapt it to node classification by replacing the clustering operation with the node-level evaluation protocol mentioned above.

For both tasks, we also include a set of other relevant state-of-the-art competing methods, including non-contrastive methods and other contrastive methods. Particularly, for graph classification, the non-contrastive methods include Graphlet Kernel (GK) [64], Weisfeiler-Lehman Sub-tree Kernel (WL) [65], Deep Graph Kernels (DGK) [62], node2vec [66], sub2vec [67] and graph2vec [63], while the GCL methods include InfoGraph [11], JOAOv2 [14], SimGRACE [15] and GraphCL [12]. For the node classification task, non-contrastive methods include node2vec [66], DeepWalk (DW) [68], and Graph AutoEncoders (GAE and VGAE) [69]. Contrastive methods include DGI [9], GMI [70], MVGRL [10], and GCA [13].

Note that ProGCL proposed two strategies to utilize the estimated hardness results, i.e., weighting and mixup. The results reported are based on the weighting strategy of ProGCL to have a direct comparison to our weighting-based AUGCL.

B. Enabling Different GCL Methods on Graph and Node Classification

1) *Performance Improvement over Baselines*: We first compare the performance of our method with the baselines on graph and node classification tasks. The results are shown in Table II. It is clear that, by incorporating our affinity uncertainty-based hardness measure, the two baselines – GraphCL [12] and GCA [13] – are substantially and consistently boosted on all datasets from different domains for both graph and node classification tasks. This demonstrates that our method AUGCL can enable these baselines to effectively attend to hard negative instances and learn better representations of graphs/nodes.

2) *Comparison to State-of-the-art Methods*: We then compare AUGCL to diverse advanced graph embedding learning methods.

Graph Classification. The results on graph classification are reported in Table III. We can observe that graph contrastive learning methods generally obtain better performance than non-contrastive methods. Our method further improves the performance by learning and feeding the affinity uncertainty-based hardness into contrastive learning, substantially outperforming SOTA GCL methods on 6 out of 7 datasets.

Compared to the three recent hardness-aware methods DCL, HCL and ProGCL, our method AUGCL performs much better across all seven datasets. Particularly, DCL, HCL and ProGCL improve GraphCL on some datasets such as PROTEINS, MUTAG, and IMDB-B, but they fail on the other ones. By contrast, our method improves over GraphCL by a large margin across all the seven datasets, indicating the superiority of our affinity uncertainty-based hardness learning method over its recent counterparts.

Node Classification. The node classification results are reported in Table IV. In general, the trends here are similar to the results in Table III: i) contrastive methods are generally more effective than the non-contrastive ones, ii) the competing hardness-aware methods DCL, HCL and ProGCL further improve over the contrastive methods on part of the datasets, while our method AUGCL achieves consistently better performance on all the three datasets, and iii) although

TABLE III

RESULTS OF GRAPH CLASSIFICATION ACCURACY (%). WE OBTAIN THE RESULTS OF GRAPHCL AND ITS FOUR HARDNESS-AWARE VARIANTS UNDER THE SAME EXPERIMENTAL SETTING AS [12], FROM WHICH THE RESULTS OF THE OTHER METHODS ARE TAKEN. THE RESULTS OF JOAOv2 [14] AND SIMGRACE [15] ARE TAKEN FROM THEIR OWN PAPER.

Type	Method	NCII	PROTEINS	DD	MUTAG	COLLAB	RDT-B	IMDB-B
Non-Contrastive Methods	GK	-	-	-	81.66±2.11	-	77.34±0.18	65.87±0.98
	WL	80.01±0.50	72.92±0.56	-	80.72±3.00	-	68.82±0.41	72.30±3.44
	DGK	80.31±0.46	73.30±0.82	-	87.44±2.72	-	78.04±0.39	66.96±0.56
	node2vec	54.89±1.61	57.49±3.57	-	72.63±10.20	-	-	-
	sub2vec	52.84±1.47	53.03±5.55	-	61.05±15.80	-	71.48±0.41	55.26±1.54
	graph2vec	73.22±1.81	73.30±2.05	-	83.15±9.25	-	75.78±1.03	71.10±0.54
Contrastive Methods	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	73.03±0.87
	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	70.83±0.25
	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	71.30±0.77
	GraphCL	78.26±0.39	74.36±0.44	78.01±0.77	87.15±0.86	71.53±1.03	90.09±0.70	71.20±0.72
Hardness-aware Methods	GraphCL _{DCL}	77.62±0.67	74.73±0.39	76.84±1.24	88.28±1.75	70.36±0.97	89.88±0.72	70.62±0.58
	GraphCL _{HCL}	78.16±0.53	74.39±0.77	76.83±1.15	88.94±1.22	70.37±0.33	90.05±0.47	71.38±0.62
	GraphCL _{ProGCL}	76.93±0.47	74.48±0.37	79.22±0.90	88.73±1.40	70.46±0.28	90.51±0.49	71.58±0.59
	GraphCL _{AUGCL} (Ours)	80.16±0.34	75.76±0.39	80.14±0.54	89.20±1.08	72.10±0.65	91.19±0.44	72.46±0.80

TABLE IV

RESULTS OF NODE CLASSIFICATION ACCURACY (%). WE OBTAIN THE RESULTS OF GCA AND ITS FOUR HARDNESS-AWARE VARIANTS UNDER THE SAME EXPERIMENTAL SETTING AS [13], FROM WHICH THE RESULTS OF THE OTHER METHODS ARE TAKEN.

Type	Method	Wiki-CS	Amazon-Computers	Amazon-Photo
Non-Contrastive Methods	Raw feature	71.98±0.00	73.81±0.00	78.53±0.00
	node2vec	71.79±0.05	84.39±0.08	89.67±0.12
	DW	74.35±0.06	85.68±0.06	89.44±0.11
	DW+feature	77.21±0.03	86.28±0.07	90.05±0.08
	GAE	70.15±0.01	85.27±0.19	91.62±0.13
	VGAE	75.63±0.19	86.37±0.21	92.20±0.11
Contrastive Methods	DGI	75.35±0.14	83.95±0.47	91.61±0.22
	GMI	74.85±0.08	82.21±0.31	90.68±0.17
	MVGRL	77.52±0.08	87.52±0.11	91.74±0.07
	GCA	78.08±0.58	87.80±0.42	91.99±0.39
Hardness-aware Methods	HSAN	78.55±0.51	81.23±1.05	89.15±0.66
	GCA _{DCL}	78.12±0.60	86.79±0.48	91.29±0.32
	GCA _{HCL}	78.19±0.64	87.64±0.34	91.79±0.29
	GCA _{ProGCL}	78.33±0.64	88.68±0.35	93.01±0.29
	GCA _{AUGCL} (Ours)	78.59±0.56	88.94±0.44	93.43±0.32

HSAN [28] showed impressive performance in enabling node clustering tasks compared to GCA [13], it does not demonstrate consistently promising performance in our evaluation tasks. This discrepancy may arise from the fact that HSAN and GCA use different backbones and HSAN was not originally designed for node classification tasks.

From the results on graph and node classification tasks, we can see that general hardness-aware methods DCL and HCL do not achieve satisfactory performance on both tasks. This can be attributed to that they treat negatives similar to the anchor as hard negatives. Different from independent data, the non-i.i.d nature of graph data and the message-passing mechanism of graph neural networks result in similar negatives having the same label as the anchor with high probability, making DCL and HCL less effective for graph contrastive learning. To address this issue, ProGCL measured the hardness of negatives based on the similarity between negatives and the anchor and the learned probability of a negative being true from a two-component beta mixture model. However, the prior employed by ProGCL that the negative similarity distribution is bimodal limits the applicability of ProGCL. Thus, ProGCL is not generalizable to the graph classification task, i.e., ProGCL fails to work as effectively as the baseline

GraphCL on some datasets as shown in Table III, while our method AUGCL consistently outperforms the baseline. This is due to that our method learns a data-driven affinity uncertainty estimation model without the prior assumption, resulting in better applicability and flexibility of AUGCL on different graph mining tasks than ProGCL.

C. Improving Robustness against Graph Adversarial Attacks

Self-supervised learning has shown effective performance in defending against adversarial perturbations [72], [73]. This subsection investigates whether AUGCL can further improve over the GCL methods on this important property. In this experiment, following [74], three different types of graph adversarial attacks: RandSampling, GradArgmax and RL-S2V are used, where RandSampling randomly adds or deletes edges from graphs, GradArgmax performs edge modification based on gradient information, and RL-S2V is a reinforcement learning based attack method that learns a generalizable attack policy. We also use the widely-used evaluation protocol as in [74] where the graph task is to classify the component numbers in synthetic graphs and structure2vec [71] is adopted as the graph encoder, with different depths of structure2vec considered in the experiments. Both the original structure2vec trained from scratch (i.e., no pre-training) and the pre-trained structure2vec [71] using GraphCL [12] are used as baselines. The experimental results of our method are obtained by incorporating our affinity uncertainty-based hardness into GraphCL to pre-train the structure2vec. The best-competing method ProGCL is adopted in the same way. The results are reported in Table V.

From the table, we can observe that: i) all three GCL methods GraphCL, ProGCL, and AUGCL can largely improve the robustness against all three graph adversarial attacks, particularly the more advanced attacks GradArgmax and RL-S2V, on different network layers, compared with the original model, ii) the robustness can be further improved by exploiting hard negative mining techniques used in ProGCL and AUGCL, compared to GraphCL, and iii) compared with ProGCL, the better hard negative mining in our method AUGCL generally results in more remarkably and stably improved robustness

TABLE V

CLASSIFICATION ACCURACY UNDER THREE EVASION ATTACKS ON THREE DIFFERENT LAYERS OF STRUCTURE2VEC [71]. PROGCL AND AUGCL BELOW REPRESENT GRAPHCL METHODS WITH THEIR RESPECTIVE HARD NEGATIVE MINING COMPONENT.

Attacks	2-Layer				3-Layer				4-Layer			
	Original	GraphCL	ProGCL	AUGCL	Original	GraphCL	ProGCL	AUGCL	Original	GraphCL	ProGCL	AUGCL
Unattack	93.20	94.73	94.13	94.20	98.20	98.33	98.67	98.87	98.87	99.00	99.47	99.20
RandSampling	78.73	80.68	82.47	82.67	92.27	92.60	93.93	94.67	95.13	97.40	97.13	97.93
GradArgmax	69.47	69.26	74.80	77.53	64.60	89.33	94.07	93.27	95.80	97.00	95.67	96.47
RL-S2V	42.93	42.20	42.13	42.47	41.93	61.66	62.07	63.73	70.20	84.86	86.67	87.33

TABLE VI

ABLATION STUDY RESULTS OF GRAPHCL_{AUGCL} USING DIFFERENT CLUSTERING METHODS f OR UNCERTAINTY ESTIMATION METHODS g .

Method	f	g	NCII	PROTEINS	DD	MUTAG	COLLAB	RDT-B	IMDB-B
GraphCL (Baseline)	×	×	78.26±0.39	74.36±0.44	78.01±0.77	87.15±0.86	71.53±1.03	90.09±0.70	71.20±0.72
GraphCL _{AUGCL}	Spectral	ExtraClass	79.79±0.52	75.57±0.79	79.05±0.44	88.86±2.22	71.74±0.88	91.01±0.56	71.86±0.47
	k -means	ExtraClass	80.16±0.34	75.76±0.39	80.14±0.54	89.20±1.08	72.10±0.65	91.19±0.44	72.46±0.80
	k -means	Distance	78.88±0.26	75.48±1.06	79.22±0.25	88.64±0.83	71.27±0.43	90.41±0.51	72.20±0.55
	k -means	Softmax	79.42±0.43	75.34±0.57	78.10±0.65	86.81±1.65	71.89±0.92	90.38±0.33	71.72±0.57
	k -means	Entropy	79.98±0.50	75.13±0.39	79.58±0.49	88.98±1.78	71.71±0.37	90.37±0.50	71.76±0.57

over the GraphCL. Overall, the proposed method AUGCL increases the classification accuracy by up to 8% over GraphCL and up to 2.7% over ProGCL and performs very competitively to the two methods (i.e., around 0.2%-0.8% difference) in the limited cases where AUGCL is not the best performer.

D. Ablation Studies

This subsection evaluates the impact of using different clustering and uncertainty estimation methods in f and g , respectively. The GraphCL_{AUGCL} method is used, with GraphCL as the baseline.

Partition Methods in f . An important module of our proposed method is the instance-wise partition function f . k -means is used by default to implement f . Here we also examine the use of spectral clustering [75] to perform the binary partition. The results are shown in Table VI. We can see that AUGCL using either spectral clustering or k -means achieves similar improvement over GraphCL, suggesting the stability of our method w.r.t. the generation of the partition labels. AUGCL with k -means clustering performs consistently better than spectral clustering. Hence, k -means clustering is used by default in our experiments and recommended in practice.

Uncertainty Estimation Methods in g . The uncertainty estimation method g is another important module in AUGCL. In addition to the extra class-based method used by default in AUGCL, two alternative approaches are used, including the maximum prediction probability-based method Softmax-Response [44] and the entropy-based method Predictive Entropy [45]. We also include a distance-based method as another simplified variant of AUGCL. The detailed descriptions of these uncertainty estimation methods are presented in Appendix.

The results are reported in Table VI. It is clear that regardless of the specific uncertainty estimation method used, all variants of AUGCL can generally improve the baseline GraphCL on nearly all datasets. This provides further evidence for the effectiveness of our approach. Additionally,

the uncertainty estimation method matters: the default method (a recently proposed extra class-based method [49], [51]), a more effective uncertainty estimation model than the other three methods, shows consistently better performance than the other three variants, implying that the hardness can be better captured by more advanced uncertainty estimation methods.

E. Hyperparameter Analysis


We examine the sensitivity of AUGCL w.r.t two key hyperparameters, i.e., the uncertainty parameter α in (3) and the reward parameter o in ϕ (particularly in (7)). Without loss of generality, one graph dataset from biochemical molecules and social networks respectively, i.e., PROTEINS and IMDB-B, are used.

Uncertainty Parameter α . α is adaptively set, depending on the uncertainty matrix \mathbf{U} , to enable stable performance of AUGCL. Particularly, given \mathbf{U} , we can calculate the mean μ and standard deviation δ of \mathbf{U} , based on which α is set to $\alpha = \frac{1}{\mu}$. We vary the parameter α in the range of $\{\frac{1}{\mu-\delta}, \frac{1}{\mu-0.5\delta}, \frac{1}{\mu}, \frac{1}{\mu+0.5\delta}, \frac{1}{\mu+\delta}\}$. The mean classification accuracy (%) under different α are shown in Fig. 3(a) where the labels in the x-axis denote the coefficient of δ when calculating α . It is clear that the performance of our model is generally stable with varying α , and $\alpha = \frac{1}{\mu}$ is a recommended setting.

Reward Parameter o . We further examine the reward parameter o in the uncertainty estimation model [51]. With o varying in $\{1.5, 1.6, 1.7, 1.8, 1.9\}$, we report the mean classification accuracy (%) in Fig. 3(b). The results also show that AUGCL can achieve reasonably stable performance for a wide range of the o settings.

F. Training Time

This subsection demonstrates the computation time of the proposed method. Following the experimental setting in the hyperparameter analysis section, we compare the training time on PROTEINS and IMDB-B datasets in a single run with the same setting. The comparison results of our method AUGCL

Fig. 3. Sensitivity analysis of hyperparameters α and o .TABLE VII
TRAINING TIME (SECOND) ON PROTEINS AND IMDB-B.

Dataset	GraphCL	DCL	HCL	ProGCL	AUGRL
PROTEINS	25.30	25.82	26.11	27.15	29.02
IMDB-B	14.22	14.68	14.78	16.44	18.60

to the baseline GraphCL [12] and all the hardness-aware methods are reported in Table VII.

From the table, we can see that the utilization of hardness-aware methods leads to some increases in computation time. Although AUGCL has the highest computation time, the overall computational overhead compared to the baseline is very minor.

G. Convergence

This subsection demonstrates the convergence of the proposed method. We illustrate the loss curves on PROTEINS and IMDB-B datasets in Fig IV-G.




Fig. 4. Loss curve of the proposed method.

From the figure, we can see that the proposed method achieves fast convergence. Note that there are abrupt drops for both datasets at epoch 10 instead of gradual decreases of loss values. This is attributed to the affinity uncertainty learning performed at epoch 10, in which the learned hardness-based weights effectively capture the hardness of the negatives, i.e., large weights for hard negatives while small weights for the other instances (false and easy negatives). By incorporating the learned weights, the loss from false and easy negatives is significantly reduced, resulting in the learning focusing on the hard negatives and facilitating the convergence of graph contrastive learning.

V. CONCLUSION

This paper proposes the idea of affinity uncertainty and utilizes it to measure the hardness of negative samples to


improve popular GCL models. To this end, we introduce the affinity uncertainty-based hardness learning approach AUGCL that synthesizes binary partition and uncertainty estimation to learn anchor-instance-dependent hardness for all negative instances, i.e., their hardness results are relative to each anchor instance. AUGCL is a data-driven approach that eliminates the prior assumption made in very recent hardness-aware GCL methods like ProGCL [23], resulting in better applicability and flexibility on different graph mining tasks, as well as better robustness to diverse graph adversarial attacks. It also shows better performance in enabling different GCL loss functions, compared to a wide range of other state-of-the-art graph representation methods on graph and node classification tasks. We also show theoretically that the resulting contrastive loss in AUGCL is equivalent to a triplet loss with an adaptive margin that adaptively exploits the hard negatives with a large margin, with a small margin assigned to the other negative instances.

REFERENCES


- [1] F. Xia, K. Sun, S. Yu, A. Aziz, L. Wan, S. Pan, and H. Liu, "Graph learning: A survey," *IEEE Transactions on Artificial Intelligence*, vol. 2, no. 2, pp. 109–127, 2021.
- [2] J. Ren, F. Xia, I. Lee, A. Noori Hoshyar, and C. Aggarwal, "Graph learning for anomaly analytics: Algorithms, applications, and challenges," *ACM Transactions on Intelligent Systems and Technology*, vol. 14, no. 2, pp. 1–29, 2023.
- [3] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [4] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *International Conference on Learning Representations*, 2019.
- [5] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations*, 2017.
- [6] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *International Conference on Learning Representations*, 2018.
- [7] Y. Xie, Z. Xu, J. Zhang, Z. Wang, and S. Ji, "Self-supervised learning of graph neural networks: A unified review," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 2, pp. 2412–2429, 2023.
- [8] Y. Liu, M. Jin, S. Pan, C. Zhou, Y. Zheng, F. Xia, and P. S. Yu, "Graph self-supervised learning: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 6, pp. 5879–5900, 2023.
- [9] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax," *International Conference on Learning Representations*, vol. 2, no. 3, p. 4, 2019.
- [10] K. Hassani and A. H. Khasahmadi, "Contrastive multi-view representation learning on graphs," in *International Conference on Machine Learning*. PMLR, 2020, pp. 4116–4126.
- [11] F.-Y. Sun, J. Hoffman, V. Verma, and J. Tang, "Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization," in *International Conference on Learning Representations*, 2019.
- [12] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5812–5823, 2020.
- [13] Y. Zhu, Y. Xu, F. Yu, Q. Liu, S. Wu, and L. Wang, "Graph contrastive learning with adaptive augmentation," in *Proceedings of the Web Conference 2021*, 2021, pp. 2069–2080.
- [14] Y. You, T. Chen, Y. Shen, and Z. Wang, "Graph contrastive learning automated," in *International Conference on Machine Learning*. PMLR, 2021, pp. 12121–12132.
- [15] J. Xia, L. Wu, J. Chen, B. Hu, and S. Z. Li, "Simgrace: A simple framework for graph contrastive learning without data augmentation," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 1070–1079.

- [16] J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang, "Gcc: Graph contrastive coding for graph neural network pre-training," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1150–1160.
- [17] L. Wan, Z. Fu, L. Sun, X. Wang, G. Xu, X. Yan, and F. Xia, "Self-supervised teaching and learning of representations on graphs," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 489–498.
- [18] Z. Chen, Y. Peng, S. Yu, C. Cao, and F. Xia, "Subgraph adaptive structure-aware graph contrastive learning," *Mathematics*, vol. 10, no. 17, p. 3047, 2022.
- [19] Y. Zheng, M. Jin, S. Pan, Y.-F. Li, H. Peng, M. Li, and Z. Li, "Toward graph self-supervised learning with contrastive adjusted zooming," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2022.
- [20] Y. Kalantidis, M. B. Sariyildiz, N. Pion, P. Weinzaepfel, and D. Larlus, "Hard negative mixing for contrastive learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 798–21 809, 2020.
- [21] J. Robinson, C.-Y. Chuang, S. Sra, and S. Jegelka, "Contrastive learning with hard negative samples," *International Conference on Learning Representations*, 2021.
- [22] C.-Y. Chuang, J. Robinson, Y.-C. Lin, A. Torralba, and S. Jegelka, "Debiased contrastive learning," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [23] J. Xia, L. Wu, G. Wang, J. Chen, and S. Z. Li, "Progcl: Rethinking hard negative mining in graph contrastive learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 24 332–24 346.
- [24] M. Wu, M. Mosse, C. Zhuang, D. Yamins, and N. Goodman, "Conditional negative sampling for contrastive learning of visual representations," in *International Conference on Learning Representations*, 2020.
- [25] K. Lee, Y. Zhu, K. Sohn, C.-L. Li, J. Shin, and H. Lee, "\$i\$-mix: A domain-agnostic strategy for contrastive representation learning," in *International Conference on Learning Representations*, 2021.
- [26] Y. Zhu, Y. Xu, Q. Liu, and S. Wu, "An empirical study of graph contrastive learning," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [27] S. Wan, Y. Zhan, S. Chen, S. Pan, J. Yang, D. Tao, and C. Gong, "Boosting graph contrastive learning via adaptive sampling," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–13, 2023.
- [28] Y. Liu, X. Yang, S. Zhou, X. Liu, Z. Wang, K. Liang, W. Tu, L. Li, J. Duan, and C. Chen, "Hard sample aware network for contrastive deep graph clustering," in *Proc. of AAAI*, 2023.
- [29] A. Van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv e-prints*, pp. arXiv–1807, 2018.
- [30] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3733–3742.
- [31] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1597–1607.
- [32] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.
- [33] Z. Tong, Y. Liang, H. Ding, Y. Dai, X. Li, and C. Wang, "Directed graph contrastive learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 19 580–19 593, 2021.
- [34] Y. Zhu, Y. Xu, F. Yu, Q. Liu, S. Wu, and L. Wang, "Deep graph contrastive representation learning," *arXiv preprint arXiv:2006.04131*, 2020.
- [35] D. Xu, W. Cheng, D. Luo, H. Chen, and X. Zhang, "Infogcl: Information-aware graph contrastive learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 30 414–30 425, 2021.
- [36] H. Li, J. Cao, J. Zhu, Q. Luo, S. He, and X. Wang, "Augmentation-free graph contrastive learning of invariant-discriminative representations," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–11, 2023.
- [37] W. Zhuo and G. Tan, "Graph contrastive learning with adaptive proximity-based graph augmentation," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2023.
- [38] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," *arXiv preprint physics/0004057*, 2000.
- [39] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018.
- [40] V. Verma, T. Luong, K. Kawaguchi, H. Pham, and Q. Le, "Towards domain-agnostic contrastive learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 10 530–10 541.
- [41] S. Ge, S. Mishra, C.-L. Li, H. Wang, and D. Jacobs, "Robust contrastive learning using negative samples with diminished semantics," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [42] X. Chen, W. Chen, T. Chen, Y. Yuan, C. Gong, K. Chen, and Z. Wang, "Self-pu: Self boosted and calibrated positive-unlabeled training," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1510–1519.
- [43] W. Tu, S. Zhou, X. Liu, C. Ge, Z. Cai, and Y. Liu, "Hierarchically contrastive hard sample mining for graph self-supervised pretraining," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2023.
- [44] Y. Geifman and R. El-Yaniv, "Selective classification for deep neural networks," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [45] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [46] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," in *International Conference on Learning Representations*, 2018.
- [47] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1050–1059.
- [48] W. Liu, X. Wang, J. Owens, and Y. Li, "Energy-based out-of-distribution detection," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 464–21 475, 2020.
- [49] Y. Tian, Y. Liu, G. Pang, F. Liu, Y. Chen, and G. Carneiro, "Pixel-wise energy-biased abstention learning for anomaly segmentation on complex urban driving scenes," in *European Conference on Computer Vision*, 2022.
- [50] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," in *International Conference on Learning Representations*, 2018.
- [51] Z. Liu, Z. Wang, P. P. Liang, R. R. Salakhutdinov, L.-P. Morency, and M. Ueda, "Deep gamblers: Learning to abstain with portfolio theory," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [52] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [53] J. Li, P. Zhou, C. Xiong, and S. Hoi, "Prototypical contrastive learning of unsupervised representations," in *International Conference on Learning Representations*, 2020.
- [54] S. Lin, C. Liu, P. Zhou, Z.-Y. Hu, S. Wang, R. Zhao, Y. Zheng, L. Lin, E. Xing, and X. Liang, "Prototypical graph contrastive learning," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–12, 2022.
- [55] H. Zhao, X. Yang, Z. Wang, E. Yang, and C. Deng, "Graph debiased contrastive learning with joint representation clustering," in *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2021, pp. 3434–3440.
- [56] M. Xu, H. Wang, B. Ni, H. Guo, and J. Tang, "Self-supervised graph-level representation learning with local and global structure," in *International Conference on Machine Learning*. PMLR, 2021, pp. 11 548–11 558.
- [57] B. Li, B. Jing, and H. Tong, "Graph communal contrastive learning," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 1203–1213.
- [58] X. Luo, W. Ju, M. Qu, Y. Gu, C. Chen, M. Deng, X.-S. Hua, and M. Zhang, "Clear: Cluster-enhanced contrast for self-supervised graph representation learning," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2022.
- [59] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [60] P. Mernyei and C. Cangea, "Wiki-cs: A wikipedia-based benchmark for graph neural networks," *arXiv preprint arXiv:2007.02901*, 2020.
- [61] O. Shchur, M. Mumme, A. Bojchevski, and S. Günnemann, "Pitfalls of graph neural network evaluation," *arXiv preprint arXiv:1811.05868*, 2018.
- [62] P. Yanardag and S. Vishwanathan, "Deep graph kernels," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 1365–1374.

- [63] A. Narayanan, M. Chandramohan, R. Venkatesan, L. Chen, Y. Liu, and S. Jaiswal, “graph2vec: Learning distributed representations of graphs,” *arXiv preprint:1707.05005*, 2017.
- [64] N. Shervashidze, S. Vishwanathan, T. Petri, K. Mehlhorn, and K. Borgwardt, “Efficient graphlet kernels for large graph comparison,” in *Artificial Intelligence and Statistics*. PMLR, 2009, pp. 488–495.
- [65] N. Shervashidze, P. Schweitzer, E. J. Van Leeuwen, K. Mehlhorn, and K. M. Borgwardt, “Weisfeiler-lehman graph kernels,” *Journal of Machine Learning Research*, vol. 12, no. 9, 2011.
- [66] A. Grover and J. Leskovec, “node2vec: Scalable feature learning for networks,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 855–864.
- [67] B. Adhikari, Y. Zhang, N. Ramakrishnan, and B. A. Prakash, “Sub2vec: Feature learning for subgraphs,” in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2018, pp. 170–182.
- [68] B. Perozzi, R. Al-Rfou, and S. Skiena, “Deepwalk: Online learning of social representations,” in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014, pp. 701–710.
- [69] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” *arXiv preprint arXiv:1611.07308*, 2016.
- [70] Z. Peng, W. Huang, M. Luo, Q. Zheng, Y. Rong, T. Xu, and J. Huang, “Graph representation learning via graphical mutual information maximization,” in *Proceedings of The Web Conference 2020*, 2020, pp. 259–270.
- [71] H. Dai, B. Dai, and L. Song, “Discriminative embeddings of latent variable models for structured data,” in *International Conference on Machine Learning*. PMLR, 2016, pp. 2702–2711.
- [72] D. Hendrycks, M. Mazeika, S. Kadavath, and D. Song, “Using self-supervised learning can improve model robustness and uncertainty,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [73] M. Kim, J. Tack, and S. J. Hwang, “Adversarial self-supervised contrastive learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 2983–2994, 2020.
- [74] H. Dai, H. Li, T. Tian, X. Huang, L. Wang, J. Zhu, and L. Song, “Adversarial attack on graph structured data,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 1115–1124.
- [75] A. Ng, M. Jordan, and Y. Weiss, “On spectral clustering: Analysis and an algorithm,” *Advances in Neural Information Processing Systems*, vol. 14, 2001.



Ling Chen (Senior Member, IEEE) received the Ph.D. degree from Nanyang Technological University, Singapore, and undertook postdoctoral training at Leibniz University Hannover (L3S Research Centre), Germany. She is a Professor in the School of Computer Science at University of Technology of Sydney, Sydney, Australia. She leads the Data Science and Knowledge Discovery Laboratory (The DSKD Lab) within the Australian Artificial Intelligence Institute (AAIL) at UTS. Her papers appear in major conferences and journals, including SIGKDD, AAAI, ICLR, ICDM, NeurIPS, and TNNLS. Her research interests mainly include discovering regularities and irregularities from various types of data, data representation learning, social media and social network mining, and dialogue and interactive systems.



Chaoxi Niu received the B.S. and M.S. degrees in electronic information science and technology from Lanzhou University, Lanzhou, China. He is currently pursuing the Ph.D. degree with the University of Technology Sydney, Sydney, NSW, Australia. His current research interests primarily focus on deep learning on graphs.



Guansong Pang is a tenure-track Assistant Professor of Computer Science in the School of Computing and Information Systems at Singapore Management University (SMU), Singapore. Before joining SMU, he was a Research Fellow with the Australian Institute for Machine Learning (AIML). He received a Ph.D. degree from University of Technology Sydney, Australia. His research interests lie in machine learning techniques and their applications, with a focus on handling abnormal and unknown data.