Graph Communal Contrastive Learning

Bolian Li **Tianjin University** Tianjin, China libolian@tju.edu.cn

Baoyu Jing University of Illinois at Urbana-Champaign IL, USA baoyuj2@illinois.edu

Hanghang Tong University of Illinois at Urbana-Champaign IL, USA htong@illinois.edu

ABSTRACT

Graph representation learning is crucial for many real-world applications (e.g. social relation analysis). A fundamental problem for graph representation learning is how to effectively learn representations without human labeling, which is usually costly and time-consuming. Graph contrastive learning (GCL) addresses this problem by pulling the positive node pairs (or similar nodes) closer while pushing the negative node pairs (or dissimilar nodes) apart in the representation space. Despite the success of the existing GCL methods, they primarily sample node pairs based on the nodelevel proximity yet the community structures have rarely been taken into consideration. As a result, two nodes from the same community might be sampled as a negative pair. We argue that the community information should be considered to identify node pairs in the same communities, where the nodes insides are semantically similar. To address this issue, we propose a novel Graph Communal Contrastive Learning (gCooL) framework to jointly learn the community partition and learn node representations in an end-to-end fashion. Specifically, the proposed *qCooL* consists of two components: a Dense Community Aggregation (DeCA) algorithm for community detection and a Reweighted Self-supervised Cross-contrastive (ReSC) training scheme to utilize the community information. Additionally, the real-world graphs are complex and often consist of multiple views. In this paper, we demonstrate that the proposed gCooL can also be naturally adapted to multiplex graphs. Finally, we comprehensively evaluate the proposed *qCooL* on a variety of real-world graphs. The experimental results show that the *qCooL* outperforms the state-of-the-art methods.

CCS CONCEPTS

• Computing methodologies → Unsupervised learning; Learning latent representations; \bullet Mathematics of computing \rightarrow Information theory; Graph algorithms.

KEYWORDS

self-supervised learning, graph contrastive learning, community detection

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/10.1145/3485447.3512208

Bolian Li, Baoyu Jing, and Hanghang Tong. 2022. Graph Communal Contrastive Learning. In Proceedings of the ACM Web Conference 2022 (WWW 22), April 25-29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3485447.3512208

1 INTRODUCTION

ACM Reference Format:

Graph representation learning is crucial in many real-world applications, including predicting friendship in social relationships [76], recommending products to potential users [31] and detecting social opinions [60, 61]. However, traditional graph representation learning methods [19, 28, 58] demand labeled graph data of high quality, while such data is too expensive to obtain and sometimes even unavailable due to privacy and fairness concerns [25, 26]. To address this problem, recent Graph Contrastive Learning (GCL) [15, 16, 28, 33, 49, 71, 78] proposes to train graph encoders by distinguishing positive and negative node pairs without using external labels. Among them, node-level GCL methods regard all node pairs as negative samples, which pulls closer positive node pairs (representations of the same node in two views) and pushes away negative node pairs (representations of different nodes) [64]. The variants of GCL [16, 56, 72] employ mutual information maximization to contrast node pairs and have achieved state-of-the-art performance on many downstream tasks, such as node classification [2], link prediction [34] and node clustering [52].

Despite the success of these methods, most of the existing GCL methods primarily focus on the node-level proximity but fail to explore the inherent community structures in the graphs when sampling node pairs. For example, MVGRL [16], GCA [78] and GraphCL [71] regard the same nodes in different graph views as positive pairs; [16] regards the neighbors as positive node pairs; DGI [59] and HDI [21] produce negative pairs by randomly shuffling the node attributes; [14, 49, 74] regard negative samples as different sub-graphs. However, the structural information such as the community structures in the graphs has not been fully explored yet. The community structures can be found in many real graphs [66]. For example, in social graphs, people are grouped by their interests [53]. The group of people with the same interest tend to be densely connected by their interactions, while the people with different interests are loosely connected. Therefore, the people in the same interest community are graphically similar and treating them as negative pairs will introduce graphic errors to the node representations. To address this problem, we firstly propose a novel dense community aggregation (DeCA) algorithm to learn community partition based on structural information in the graphs. Next, we introduce a novel reweighted self-supervised contrastive (ReSC) training scheme to pull the nodes within the same community closer to each other in the representation space.

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.

Additionally, in the real graphs, nodes can be connected by multiple relations [22, 24, 44, 69], and each relation can be treated as a graph view (known as multiplex graphs [10]). One way for graph representation learning methods to be adapted to the multiplex graphs is to learn the representations of each graph view separately and then combining them with a fusion model [21]. However, this approach ignores the graphic dependence between different graph views. For example, IMDB [62] can be modeled as a multiplex movie graph connected by both co-actor and co-director relations. These two kinds of connections are not independent since personal relations between actors and directors would affect their participation in movies. To address this problem, we apply contrastive training that is consistent with the situation of the single-relation graph, applying pair-wise contrastive training on each pair of graph views and using a post-ensemble fashion to combine model outputs of all views in the down-stream tasks.

The main contributions of this paper are summarized as follows:

- We propose a novel Dense Community Aggregation (*DeCA*) algorithm to detect structurally related communities by end-to-end training.
- We propose a Reweighted Self-supervised Cross-contrastive (*ReSC*) training scheme, which employs community information to enhance the performance in down-stream tasks.
- We adapt our model to multiplex graphs by pair-wise contrastive training, considering the dependence between different graph views.
- We evaluate our model on various real-world datasets with multiple down-stream tasks and metrics. Our model achieves the state-of-the-art performance.

The rest of the paper is organized as: preliminaries (Sec. 2), methodology (Sec. 3), experiments (Sec. 4), related works (Sec. 5) and conclusion (Sec. 6).

2 PRELIMINARIES

In this section, we introduce the similarity measurement used in this paper (Sec. 2.1), the basic concepts of contrastive learning (Sec. 2.2) and community detection (Sec. 2.3). Then we give a brief definition of attributed multiplex graphs in Sec. 2.4. The notations used in this paper are summarized in Table 1.

2.1 Similarity Measurement

We adopt two classic similarity functions to determine the similarity between two nodes: **exponent cosine similarity** [50]

$$\delta_{c}(\mathbf{x}_{1}, \mathbf{x}_{2}) = exp\left\{\frac{\mathbf{x}_{1}^{T}\mathbf{x}_{2}/\tau}{||\mathbf{x}_{1}|| \cdot ||\mathbf{x}_{2}||}\right\},\tag{1}$$

in which we exponentiate the cosine similarity for non-negativity, and **Gaussian RBF similarity** [3]

$$\delta_e(\mathbf{x}_1, \mathbf{x}_2) = exp\left\{-||\mathbf{x}_1 - \mathbf{x}_2||^2/\tau^2\right\},$$
(2)

where τ is the sensitivity factor.

2.2 Contrastive Learning

Contrastive learning aims at distinguishing similar (positive) and dissimilar (negative) pairs of data points, encouraging the agreement of the similar pairs and the disagreement of the dissimilar



Figure 1: Overview of *gCooL*. We firstly generate two views by graph augmentations, secondly encode the node attributes with a shared encoder, and then detect communities with *DeCA* and learn node representations with *ReSC*.

ones [8]. The general process dating back to [1] can be described as: For a given data point x, let x^1 and x^2 be differently augmented data points derived from x. z^1 and z^2 are the representations by passing x^1 and x^2 through a shared encoder. The mutual information [9] (or its variants such as [43]) between the two representations is maximized. The motivation of contrastive learning is to learn representations that are invariant to perturbation introduced by the augmentation schemes [65], and thus are robust to inherent noise in the dataset.

A widely used objective for constrastive learning is InfoNCE [43], which applies instance-wise contrastive objective by sampling the same instances in different views as positive pairs and all other pairs between different views as negative pairs. The limitation of InfoNCE is that it uniformly samples negative pairs, ignoring the semantic similarities between some instances.

2.3 Community Detection

Communities in graphs are groups of nodes that are more strongly connected among themselves than others. Community structures widely exist in many real-world graphs, such as sociology, biology and transportation systems [41]. Detecting communities in a graph has become a primary approach to understand how structure relates to high-order information inherent to the graph [20].

Table 1: Notations.

Symbol	Description
G	attributed graph
\mathbb{G}	attributed multiplex graph
E	edge set
\widetilde{E}	augmented edge set
X	attribute matrix
\widetilde{X}	augmented attribute matrix
X^r	attribute matrix of view r
Α	adjacency matrix
\widetilde{A}	extended adjacency matrix
С	community
Φ	community centroid matrix
R	community assignment matrix
F	community density matrix
$E(\cdot)$	edge count function
$d(\cdot)$	edge density function
$H(\cdot)$	entropy
$I(\cdot; \cdot)$	mutual information
$\delta_c(\cdot, \cdot)$	exponent cosine similarity
$\delta_{e}(\cdot, \cdot)$	Gaussian RBF similarity
$f_{\theta}(\cdot)$	GCN encoder with parameters θ
$g_{w}(\cdot)$	linear classifier with weights w
Т	random augmentation
Dintra	intra-community density score, scalar
D_{inter}	inter-community density score, scalar

2.3.1 Formulation. For a graph G = (V, E), with its node set $V = \{v_1, v_2, ..., v_M\}$ and edge set $E = \{(v_{i_1}, v_{j_1}), (v_{i_1}, v_{j_1}), ..., (v_{i_M}, v_{j_M})\}$, we define one of its communities *C* as a sub-graph: $C = (V^C, E^C)$, where $V^C \subseteq V$ and $E^C = E \cap (V^C \times V^C)$.

2.3.2 Modularity. A widely used evaluation for community partitions is the modularity [42], defined as:

$$m = \frac{1}{2M} \sum_{i,j} \left[A[i,j] - \frac{d_i d_j}{2M} \right] r(i,j),$$
(3)

where *A* is the adjacency matrix, d_i is the degree of the *i*-th node, and r(i, j) is 1 if node *i* and *j* are in the same community, otherwise is 0. The modularity measures the impact of each edge to the local edge density $(d_i d_j / 2M$ represents the expected local edge density), and is easily disturbed by the changing of edges.

2.3.3 Other notations. We define the edge count function over the adjacency matrix:

$$E(C) = \sum_{i,j} \mathbb{1}\left\{ \mathbf{A}^{C}[i,j] \neq 0 \right\},\tag{4}$$

where A^C is the adjacency matrix for community *C*. The edge density function compares the real edge count to the maximal possible number of edges in the given community C_k :

$$d(k) = \frac{E(C_k)}{|C_k| (|C_k| - 1)}.$$
(5)

2.4 Attributed Multiplex Graph

Multiplex graphs are also known as multi-dimensional graphs [39] or multi-view graphs [12, 23, 55], and are comprised of multiple single-view graphs with shared nodes and attributes but different

graph structures (usually with different types of link) [6]. Formally, an attributed multiplex graph is $\mathbb{G} = \{G^1, G^2, \dots, G^R\}$, where $R \in \mathbb{N}_+$ and each $G^r = (V, E^r)$ is an attributed graph. If the number of views R = 1, $\mathbb{G} = \{G^1\}$ is equivalent to the attributed graph G^1 . We show an example of attributed multiplex graph in Fig. 2.



Figure 2: An example of attributed multiplex graph. It consists of two views with shared node attributes and different sets of edges.

3 METHODOLOGY

In this section, as important components of the proposed *gCooL* model, we first introduce Dense Community Detection (*DeCA*) algorithm in Sec. 3.1, and then in Sec. 3.2, we propose the Reweighted Self-supervised Cross-contrastive (*ReSC*) training scheme. Further discussion on adaption to multiplex graphs is presented in Sec. 3.3.

3.1 Dense Community Aggregation

Node-level GCL methods are vulnerable to the problem caused by pairing structurally close nodes as negative samples. For this issue, we propose a novel Dense Community Aggregation (*DeCA*) for detecting structurally related communities of a single graph view. Our method is inspired by the modularity [42] in graphs, which measures the local edge density in communities. However, modularity is easily perturbed by variations of edges [36], which limits its robustness in detecting communities (with edges dynamically changing in each epochs). Therefore, our goal is to enhance the robustness of modularity, and to further extend the modularity by maximizing edge density in each communities. The *DeCA* is carried out by end-to-end training and is illustrated in Fig. 3.



Figure 3: Dense Community Aggregation. Each node is softly assigned to community centroids with probabilities, and the community centroids are learned in end-to-end training with Eq. 10.

We learn the community partition by training a randomly initialized centroid matrix Φ in an end-to-end fashion, where each $\Phi[k,:]$ represents the center of the *k*-th community.

First, we assign each node in the graph softly to community centroids with probabilities (computed by similarity function Eq. 1 or Eq. 2). Specifically, we define a community assignment matrix R, where each R[i,:] is a normalized similarity vector measuring the distance between the *i*-th node and all community centroids. Formally, R is computed by

$$\mathbf{R} = normalize(\delta(f_{\theta}(\mathbf{X}, \mathbf{A}), \Phi)), \tag{6}$$

where $\delta(\cdot)$ is the similarity function defined in Sec. 2.1, $f_{\theta}(\cdot)$ is the graph encoder with parameters θ , and *normalize*(\cdot) normalizes the probability of each community by dividing it by the sum of all probabilities and holds $\sum_{i} \mathbf{R}[i, j] = 1$ for each *i*.

Second, we employ two objectives for training community partition (see Table 1 and Sec. 2.3.3 for symbol definitions): **intracommunity density**

$$D_{intra} = \frac{1}{N} \sum_{i,j} \sum_{k} \left[\mathbf{A}[i,j] - d(k) \right] \mathbf{R}[i,k] \mathbf{R}[j,k],$$
(7)

and inter-community density

$$D_{inter} = \frac{1}{N(N-1)} \sum_{i,j} \sum_{k_1 \neq k_2} \mathbf{A}[i,j] \mathbf{R}[i,k_1] \mathbf{R}[j,k_2].$$
(8)

These two objectives measure the impact of each edge on the community edge density (rather than the local edge density used in modularity [42]). Specifically, in Eq. 7 and Eq. 8, the term A[i, j] - d(k)and A[i, j] - 0 represent the gap between real local density (A[i, j]) and the expected density (d(k) for intra-community and 0 for intercommunity). The centroid matrix Φ will be updated to reach a reasonable community partition, by minimizing the joint objective:

$$l_{DeCA}(\mathbf{R}) = \lambda_{w} D_{inter} - D_{intra},$$
(9)

where λ_w is the co-efficient. Moreover, for computational efficiency, we slightly modify the form of l_{DeCA} in our actual implementation (shown in Appendix C). Finally, we combine the l_{DeCA} objectives for two graph views, and carry out dense community aggregation simultaneously for them:

$$L_{DeCA} = \frac{1}{2} \left[l_{DeCA}(\mathbf{R}^1) + l_{DeCA}(\mathbf{R}^2) \right].$$
(10)

3.2 Reweighted Self-supervised Cross-contrastive Training

We propose the Reweighted Self-supervised Cross-contrastive (*ReSC*) training scheme in this section. We firstly apply graph augmentations to generate two graph views and secondly apply **node contrast** and **community contrast** simultaneously to consider both node-level and community-level information. We introduce node-community pairs as additional negative samples to address the problem of pairing nodes in the same communities as negative samples.

3.2.1 Graph augmentation. We define the graph augmentation T as randomly masking attributes and edges with a specific probability [78]. First, we define the attribute masks as random noise vector m, where each dimension of it is independently drawn from a Bernoulli distribution: $m[i] \sim Bernoulli(1 - p_v)$. The augmented attribute matrix is computed by

$$\widetilde{X} = [X[1,:] \odot \boldsymbol{m}; X[2,:] \odot \boldsymbol{m}; \dots; X[N,:] \odot \boldsymbol{m}]', \quad (11)$$

where \odot is the Hadamard product [18] and *X* is the original attribute matrix. Second, we generate the augmented edge set \widetilde{E} by randomly dropping edges from the original edge set *E* [77] with probability

$$P\left\{(v_1, v_2) \in \widetilde{E}\right\} = 1 - p_e, \forall (v_1, v_2) \in E.$$
(12)

We apply the above graph augmentations (denoted as $t^1, t^2 \sim T$ respectively for 2 independent augmentations) to generate two graph views: $(X^1, A^1) = t^1(X, A)$ and $(X^2, A^2) = t^2(X, A)$ for $t^1, t^2 \sim T$. Then, we obtain their representations by a shared GCN encoder: $Z^1 = f_{\theta}(X^1, A^1)$ and $Z^2 = f_{\theta}(X^2, A^2)$.

3.2.2 Node contrast. After generating two graph views, we employ node contrast and community contrast simultaneously to learn the node representations. We introduce a contrastive loss based on the InfoNCE [43] to contrast node-node pairs:

$$I_{NCE}(Z^{1};Z^{2}) = -\log \sum_{i} \frac{\delta(Z^{1}[i,:],Z^{2}[i,:])}{\sum_{j} \delta(Z^{1}[i,:],Z^{2}[j,:])}.$$
 (13)

We apply the node contrast symmetrically for the two graph views:

$$L_{node} = \frac{1}{2} \left[I_{NCE}(Z^1; Z^2) + I_{NCE}(Z^2; Z^1) \right],$$
(14)

which distinguishes negative pairs in the two views, and enforces maximizing the consistency between positive pairs [35].

3.2.3 Community contrast. The community contrast is based on the result of *DeCA* in Sec. 3.1. First, we obtain the community centers by training the randomly initialized community centroids matrix Φ with Eq. 10. Second, we adopt a re-weighted cross-contrastive objective to contrast node representations of one view with the community centroids of the other view (a cross-contrastive fashion). We are inspired by the cross-prediction scheme in [4] and introduce the cross-contrast into InfoNCE objective [43] to maximize the community consistency between two views, enforcing that the node representations in one community are pulled close to those of the counterpart community in the other view. Formally, the community contrast is computed by

$$l_{com}(Z, \Phi) = -\log \sum_{i} \frac{\delta(Z[i, :], \Phi[k_{i}, :])}{\delta(Z[i, :], \Phi[k_{i}, :]) + \sum_{k_{i} \neq k} w(i, k) \cdot \delta(Z[i, :], \Phi[k, :])},$$
(15)

where the *i*-th node is assigned to the k_i -th community. Here, $w(i,k) = exp \{-\gamma || Z[i,:] - \Phi[k,:] ||^2\}$ is the RBF weight function (slightly different from Gaussian RBF similarity in Eq. 2), which gives more attention to the more similar community pairs. In this objective, the similarity of node representations within the same communities are maximized since they are contrasted positively to the same centroids, while in different communities, node representations are separated by negative contrast. Similarly, we compute the contrastive objective symmetrically for the two generated graph views:

$$L_{com} = \frac{1}{2} \left[l_{com}(Z^1, \Phi^2) + l_{com}(Z^2, \Phi^1) \right],$$
(16)

where Z^1 and Φ^2 are the node representation matrices of view 1 and view 2 respectively. 3.2.4 Joint objective. We propose the α -decay coefficient to combine L_{node} , L_{DeCA} and L_{com} into a joint objective:

$$L = L_{node} + \alpha(t)L_{DeCA} + [1 - \alpha(t)]L_{com},$$
(17)

where the coefficient $\alpha(t) = \exp \{-t/\eta\}$ would decay smoothly with the proceeding of training (*t* refers to the epoch). We observe that the community partition would be stabilized within a few hundreds of epochs by *DeCA* training (see Sec. 4.3), while the training the *gCooL* model usually takes thousands of epochs. To this end, we apply the α -decay to focus mainly on training community partition at first, and gradually divert the focus to learning node representations.

In summary, the ReSC training process is shown in Algorithm 1.

Algorithm 1: The <i>ReSC</i> training.
Inputs: attribute matrix X , adjacency matrix A ;
Output: graph encoder $f_{\theta}(\cdot)$;
for <i>t</i> = 1,2, do
Sample two random graph augmentations t^1 , $t^2 \sim T$;
Generate two graph views, $(X^1, A^1) = t^1(X, A)$ and
$(X^2, A^2) = t^2(X, A);$
Encode the two views by a shared GCN encoder,
$Z^1 = f_{\theta}(X^1, A^1)$ and $Z^2 = f_{\theta}(X^2, A^2)$;
Obtain L_{node} , L_{DeCA} and L_{com} to compute the joint loss:
$L = L_{node} + \alpha(t)L_{DeCA} + [1 - \alpha(t)]L_{com};$
Update parameters θ by stochastic gradient descent on
joint objective <i>L</i> .
end

3.3 Adaptation to Multiplex graphs

The proposed gCooL framework can be naturally adapted to multiplex graphs with a few modifications on the training and inference process. We apply contrastive training between different graph views, which consider the dependence between them.

3.3.1 Training. In multiplex graphs, we no longer need to generate graph views through graph augmentation, since the different views in a multiplex graph are naturally multi-viewed data. We propose to detect communities (*DeCA*) and learn node representations (*ReSC*) on each pair of views. The modified training process is shown in Algorithm 2.

3.3.2 Inference. At the inference time, we propose to combine classification results of each view by classifier fusion (a post-ensemble fashion): Given the results of R independent classifiers, we label the final prediction according to the confidence of each classifier (i.e. the maximal value of the output softmax distribution [17]). We choose the result with the strongest confidence as the final prediction.

4 EXPERIMENTS

In this section, we evaluate the proposed *gCooL* model with respect to diverse down-stream tasks on multiple real graphs. We introduce the specifications of experiments in Sec. 4.1, and provide detailed quantitative and visual results in Sec. 4.2 and Sec. 4.3 respectively.

Algorithm 2:	The ReSC	training on	multiplex	graphs.
				., .

Inputs: attribute matrix X, adjacency matrices $A^1...A^R$; Output: graph encoder $f_{\theta}(\cdot)$; for t = 1,2,.... do Encode the R graph views by a shared GCN encoder, $Z^1 = f_{\theta}(X^1, A^1),$ $Z^2 = f_{\theta}(X^2, A^2),$ $Z^R = f_{\theta}(X^R, A^R);$ Obtain L_{node}, L_{DeCA} and L_{com} to compute the joint loss for each pair of views (G^i, G^j) such that $i \neq j$; Update parameters θ by stochastic gradient descent on joint objective L in Eq. 17 (adding up objectives of each pair of views). end

4.1 Experimental Setup

4.1.1 Datasets. We use 6 real graphs, including Amazon-Computers, Amazon-Photo, Coauthor-CS, WikiCS, IMDB and DBLP, to evaluate the performance on node classification and node clustering. The detailed statistics of all datasets are summarized in Table 2.

- Amazon-Computers and Amazon-Photo [54] are two graphs of co-purchase relations, where nodes are products and there exists an edge between two products when they are bought together. Each node has a sparse bag-of-words attribute vector encoding product reviews.
- Coauthor-CS [54] is an academic network based on the Microsoft Academic Graph. The nodes are authors and edges are co-author relations. Each node has a sparse bag-of-words attribute vector encoding paper keywords of the author.
- WikiCS [40] is a reference network from Wikipedia. The nodes are articles and edges are hyperlinks between the articles. Node attributes are computed as the average of GloVe [47] word embedding of articles.
- IMDB [62] is a movie network with two types of links: coactor and co-director. The attribute vector of each node is a bag-of-words feature of movie plots. The nodes are labeled with Action, Comedy or Drama.
- **DBLP** [62] is a paper graph with three types of links: coauthor, co-paper and co-term. The attribute vector of each node is a bag-of-words feature vector of paper abstracts.

Table 2: Dataset statistics.

Dataset	Туре	Edges	Nodes	Attributes	Classes
Amazon-Computers ¹	co-purchase	245,778	13,381	767	10
Amazon-Photo ¹	co-purchase	119,043	7,487	745	8
Coauthor-CS ¹	co-author	81,894	18,333	6,805	15
WikiCS ²	reference	216,123	11,701	300	10
IMDB ³	co-actor	66,428	3 550	2 000	3
INIDB	co-director	13,788	5,550	2,000	5
	co-author	144,738			
$DBLP^4$	co-paper	90,145	7,907	2,000	4
	co-term	57,137,515			

4.1.2 Evaluation protocol. We follow the evaluations in [78], where each graph encoder is trained in a self-supervised way. Then, the encoded representations are used to train a classifier for node classification, and fit a K-means model for comparing baselines on node clustering. We train each graph encoder for five runs with randomly selected data splits (for WikiCS, we use the given 20 data splits) and report the average performance with the standard deviation.

For the node classification task, we measure the performance with Micro-F1 and Macro-F1 scores. For the node clustering task, we measure the performance with Normalized Mutual Information (NMI) score: $NMI = 2I(\hat{Y}; Y)/[H(\hat{Y}) + H(Y)]$, where \hat{Y} and Y refer to the predicted cluster indexes and class labels respectively, and Adjusted Rand Index (ARI): $ARI = RI - \mathbb{E}[RI]/(max\{RI\} - \mathbb{E}[RI])$, where RI is the Rand Index [51], which measures the similarity between two cluster indexes and class labels.

4.1.3 Baselines. We compare the gCooL methods ($gCooL_c$ for exponent cosine similarity and $gCooL_e$ for Gaussian RBF similarity) with two sets of baselines:

- On single-view graphs, we consider baselines as 1) traditional methods including node2vec [13] and DeepWalk [48], 2) supervised methods including GCN [28], and 3) unsupervised methods including MVGRL [16], DGI [59], HDI [21], graph autoencoders (GAE and VGAE) [29] and GCA [78].
- On multiplex graphs, we consider baselines as 1) methods with single-view representations including node2vec [13], DeepWalk [48], GCN [28] and DGI [59], and 2) methods with multi-view representations including CMNA [7], MNE [70], HAN [62], DMGI [44] and HDMI [21].

Additionally, we compare different clustering baselines including K-means, Spectral biclustering (SBC) [30] and modularity [42] to show the effectiveness of our proposed DeCA ($DeCA_c$ for exponent cosine similarity and $DeCA_e$ for Gaussian RBF similarity).

4.1.4 Implementation. We use a 2-layer GCN [28] as the graph encoder for each deep learning baseline. We use Adam optimizer [27] to optimize each model. Before applying the *ReSC* objective, we pass the representations of two graph views through a shard projection module (a 2-layer MLP) with the same input, hidden-layer and output dimensions. The number of communities is set according to the number of classes in each graph. More implementation details are listed in Appendix A.

4.2 Quantitative Results

4.2.1 Overall performance. We present the quantitative results of node classification on single-view graphs (Table 3), node clustering on single-view graphs (Table 4), and node classification on multiplex graphs (Table 5). *gCooL* outperforms the state-of-the-art methods on multiple tasks, and even outperforms the supervised methods.

Moreover, the proposed *DeCA* (used to cluster nodes based on its community partition) outperforms traditional clustering methods and the former modularity [42] on encoded representations (shown in Table 6), which illustrates that *DeCA* is more suitable for deep-learning-based clustering and addresses the instability of modularity.

4.2.2 Ablation study. We compare the different combinations of L_{node} (Eq. 14), L_{DeCA} (Eq. 10) and L_{com} (Eq. 16) on WikiCS in terms of node classification and node clustering. The results are listed in Table 7. We find that, first, applying node contrast is crucial to the performance (given that the combinations with L_{node} outperform others) since it makes the representations more distinguishable, second, community contrast is beneficial to the overall performance (given that $L_{node} + L_{com}$ outperforms L_{node}) since it considers community information in contrastive training, and third, the performance of community contrast is boosted significantly when it is used with DeCA (comparing L_{com} with $L_{DeCA} + L_{com}$). Moreover, we also show the visual ablation study in Table 8.

4.3 Visual Evaluations

We illustrate the significance of *DeCA* by visualizing the **edge density** and **class entropy** of the assigned communities. We evaluate each checkpoint for five times and show the means and deviations. We compare the results with traditional clustering methods (K-means and Spectral biclustering [30]) and the former modularity [42]. We also visualize the node representations for ablation study.

4.3.1 Edge density. The edge density is based on Eq. 5 and computed by the average density of all communities:

$$ED = \frac{1}{K} \sum_{k=1}^{K} d(k).$$
(18)

It is used to measure how *DeCA* learns the community partition that maximizes the intra-community density (see Sec. 3.1). Fig. 4 shows that, after a few hundreds of epochs, *DeCA* stably outperforms other clustering methods.



Figure 4: Visualization of edge density comparison.

4.3.2 *Class entropy.* The class entropy is a measure of the homogeneity of class labels in a community (the extent to which a community contains one major class, or has low entropy). We argue that a good community partition should distinguish structurally separated nodes, which is, in other words, distinguishing nodes

¹https://github.com/shchur/gnn-benchmark/raw/master/data/npz/

²https://github.com/pmernyei/wiki-cs-dataset/raw/master/dataset

³https://www.imdb.com/

⁴https://dblp.uni-trier.de/

		_						
Dataset	Amazon-(Computers	Amazon-Photo		Coaut	nor-CS	WikiCS	
Metric	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1
RawFeatures	73.82±0.01	70.10±0.04	78.45±0.04	76.10±0.01	90.40±0.02	89.01±0.06	72.00±0.03	70.28±0.09
Node2vec	84.38±0.08	82.65±0.08	89.72±0.08	87.39±0.07	85.11±0.06	82.93±0.11	71.84±0.09	70.44±0.03
DeepWalk	85.63±0.09	84.02±0.10	89.36±0.10	86.92±0.02	84.71±0.23	82.63±0.19	74.25±0.06	72.68±0.15
GCN	86.43±0.56	83.99±0.61	92.51±0.23	90.47±0.32	93.04±0.28	91.02±0.38	77.11±0.08	75.61±0.19
MVGRL	87.42±0.07	85.92±0.11	91.74±0.09	89.93±0.09	92.11±0.10	90.50±0.12	77.50±0.08	75.62±0.00
DGI	83.88±0.50	79.30±0.42	91.60±0.24	89.31±0.16	92.08±0.68	90.78±0.68	75.35±0.17	73.74±0.20
HDI	85.43±0.13	80.74±0.25	90.09±0.10	88.70±0.16	89.98±0.14	86.73±0.17	75.72±0.55	68.05±0.80
GAE	85.18±0.21	83.33±0.17	91.68±0.14	89.66±0.09	90.00±0.75	88.31±0.68	70.17±0.05	68.27±0.05
VGAE	86.44±0.25	83,72±0.12	92.24±0.08	90.04±0.17	92.08±0.08	90.11±0.06	75.56±0.20	74.12 ± 0.10
GCA	87.67±0.49	85.88±0.30	92.39±0.21	91.05±0.13	92.87±0.03	90.76±0.01	78.24±0.01	74.47 ± 0.02
gCooL _c	88.85±0.14	87.42±0.28	93.18±0.12	92.05±0.17	93.32±0.02	91.65±0.03	78.74±0.04	75.92±0.06
gCooLe	88.74±0.09	87.53±0.26	92.79±0.17	91.57±0.29	93.31±0.01	91.63±0.03	78.74±0.03	75.88±0.02

Table 3: Overall performance on node classification (in percentage).

Table 4: Overall performance on node clustering.

Dataset	Amazon-Computers		Amazon-Photo		Coaut	hor-CS	WikiCS	
Metric	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
MVGRL	0.244±0.000	0.141±0.001	0.344±0.040	0.239±0.039	0.740 ± 0.010	0.627±0.009	0.263±0.010	0.102 ± 0.011
DGI	0.318±0.020	0.165 ± 0.020	0.376±0.030	0.264 ± 0.030	0.747 ± 0.010	0.629±0.011	0.310 ± 0.020	0.131 ± 0.018
HDI	0.347±0.011	0.216 ± 0.006	0.429±0.014	0.307 ± 0.011	0.726 ± 0.008	0.607±0.016	0.238±0.002	0.105 ± 0.000
GAE	0.441 ± 0.000	0.258 ± 0.000	0.616±0.010	0.494 ± 0.008	0.731±0.010	0.614 ± 0.010	0.243±0.020	0.095 ± 0.018
VGAE	0.423±0.000	0.238±0.001	0.530±0.040	0.373 ± 0.041	0.733±0.000	0.618 ± 0.001	0.261±0.010	0.082 ± 0.008
GCA	0.426 ± 0.001	0.246 ± 0.001	0.614 ± 0.000	0.494 ± 0.000	0.735 ± 0.008	0.618 ± 0.010	0.299 ± 0.002	0121±0.003
$gCooL_c$	0.452±0.000	0.284±0.000	0.619±0.000	0.508 ± 0.000	0.750±0.021	0.624±0.083	0.318±0.011	0.145±0.024
gCooL _e	0.474±0.018	0.277 ± 0.024	0.632±0.000	0.524±0.001	0.753±0.005	0.621 ± 0.007	0.322±0.009	0.140 ± 0.010

Table 5: Overall classification performance on multiplex graphs.

Dataset	IM	DB	DBLP		
Metric	Micro-F1	Macro-F1	Micro-F1	Macro-F1	
node2vec	0.550	0.533	0.547	0.543	
DeepWalk	0.550	0.532	0.537	0.533	
GCN	0.611	0.603	0.717	0.734	
DGI	0.606	0.598	0.720	0.723	
CMNA	0.566	0.549	0.561	0.566	
MNE	0.574	0.552	0.562	0.566	
HAN	0.607	0.599	0.708	0.716	
DMGI	0.648	0.648	0.766	0.771	
HDMI	0.658	0.650	0.811	0.820	
gCooL _c	0.672	0.670	0.832	0.840	
gCooLe	0.671	0.668	0.832	0.839	

of different classes. The class entropy is computed as the average entropy of class labels in all communities:

$$CH = -\frac{1}{K} \sum_{k=1}^{K} \sum_{c} P_{k}(c) log P_{k}(c),$$
(19)

where $P_k(c)$ is the occurrence frequency of class c in the k-th community. Fig. 5 shows that, after a few hundreds of epochs, *DeCA* stably outperforms other clustering methods.

4.3.3 Visualization of node representations. To see how the node representations are distributed, we use t-SNE [57] to reduce the dimension of node representations for visualization. The node representations of each class are distributed more separately when applying L_{DeCA} as well as L_{com} , which illustrates the effectiveness of our proposed method. The results are shown in Table 8.



Figure 5: Visualization of class entropy comparison.

5 RELATED WORKS

In this section, we briefly review the related works, including graph contrastive learning (Sec. 5.1) and community detection (Sec. 5.2).

5.1 Graph Contrastive Learning

Graph contrastive learning [33] proposed to pull similar nodes into close positions in the representation space while pushing dissimilar nodes apart. [37, 64] comprehensively review the development of GCL. GRACE [77] proposed a simplified framework to maximize the agreement between two graph views. GCA [78] proposed an adaptive augmentation algorithm to generate graph views. GCC [49] designed a pre-training task to learn structural graph representations. HeCo [63] employed network schema and meta-path views WWW '22, April 25-29, 2022, Virtual Event, Lyon, France

		•	. 1° m .	1 4 4	1 1	1 1	1. 1	0.01 1
Table 6. Performance on nod	le clustering	comnarin	a different i	eliistering mei	node of	n randoml	V 10171211760	I I N encoder
Table 0. I citormance on nou	ie crusierme	, comparm	e unicicit,	ciusicime me	nous o	n rangomn	v mnnanzeu	OUT CHEOREI.
		,						

Dataset	Amazon-Computers		Amazon-Photo		Coaut	nor-CS	WikiCS	
Metric	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
K-means	0.192±0.019	0.086±0.013	0.225±0.012	0.121±0.006	0.497±0.004	0.315±0.008	0.047±0.010	0.022±0.002
SBC	0.160 ± 0.000	0.088 ± 0.000	0.100 ± 0.000	0.020 ± 0.000	0.415 ± 0.000	0.261 ± 0.000	0.039±0.000	0.024 ± 0.000
modularity	0.204±0.036	0.118 ± 0.038	0.233±0.027	0.124±0.032	0.447±0.017	0.280 ± 0.022	0.139±0.063	0.083 ± 0.052
DeCAc	0.213±0.009	0.114 ± 0.018	0.237 ± 0.024	0.126 ± 0.032	0.476 ± 0.015	0.305 ± 0.020	0.140±0.060	0.080 ± 0.055
DeCAe	0.363±0.009	0.262 ± 0.010	0.384±0.022	0.273±0.025	0.523±0.016	0.358±0.018	0.210±0.041	0.136±0.060

Table 7: Ablation study on node classification and node clustering, comparing combinations of objectives on WikiCS.

Lnode	L _{DeCA}	L _{com}	$\delta_c(\cdot)$	$\delta_{e}(\cdot)$	Micro-F1(%)	Macro-F1(%)	NMI	ARI
\checkmark			\checkmark		75.09±0.29	69.83±0.24	0.254 ± 0.005	0.120±0.007
\checkmark				\checkmark	75.71±0.58	68.52±0.55	0.228 ± 0.015	0.127 ± 0.030
		\checkmark	\checkmark		65.78±0.29	57.25±0.36	0.191 ± 0.001	0.102 ± 0.001
		\checkmark		\checkmark	63.48±0.67	53.78±0.90	0.188 ± 0.007	0.085 ± 0.005
\checkmark		\checkmark	\checkmark		77.29±0.20	74.19±0.23	0.273 ± 0.012	0.125 ± 0.020
\checkmark		\checkmark		\checkmark	77.47±0.28	74.43±0.30	0.274±0.003	0.139±0.001
	\checkmark	\checkmark	\checkmark		72.88±0.21	69.97±0.24	0.222 ± 0.020	0.120 ± 0.030
	\checkmark	\checkmark		\checkmark	68.18±0.14	64.88±0.16	0.227 ± 0.012	0.123±0.034
\checkmark	\checkmark	\checkmark	\checkmark		78.74±0.04	75.92±0.06	0.318±0.011	0.145±0.024
\checkmark	\checkmark	\checkmark		\checkmark	78.74±0.03	75.88±0.02	0.322±0.009	0.140 ± 0.010

Table 8: Ablation study on WikiCS by visualizing node representations with t-SNE.



to employ contrastive training. DGI [59], GMI [46] and HDI [21] randomly shuffle node attributes to generate negative samples. MV-GRL [16] employed graph diffusion to transform graphs.

For multiplex graphs, DGMI [44] and HDMI [21] sample negative pairs by randomly shuffling the attributes. MNE [73], mvn2vec [55] and GATNE [5] use random walk to sample negative pairs.

5.2 Community Detection

Community detection is a primary approach to understand how structure relates to high-order information inherent to the graph. [20] reviewed the recent progress in community detection. CommDGI [75] employed modularity for community detection. [32] introduced the identifiability of communities. RWM [38] used random walk to detect local communities. Destine [67, 68] detects dense subgraphs on multi-layered networks.

6 CONCLUSION

In this paper, we propose a novel Graph Communal Contrastive Learning (*gCooL*) framework to learn the community partition of structurally related communities by a Dense Community Aggregation (DeCA) algorithm, and to learn the graph representation with a reweighted self-supervised cross-contrastive (ReSC) training scheme considering the community structures. The proposed gCooL consistently achieves the state-of-the-art performance on multiple tasks and can be naturally adapted to multiplex graphs. We show that community information is beneficial to the overall performance in graph representation learning.

ACKNOWLEDGMENTS

B. Jing and H. Tong are partially supported by National Science Foundation under grant No. 1947135, the NSF Program on Fairness in AI in collaboration with Amazon under award No. 1939725, NIFA 2020-67021-32799, and Army Research Office (W911NF2110088). The content of the information in this document does not necessarily reflect the position or the policy of the Government or Amazon, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Graph Communal Contrastive Learning

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

REFERENCES

- Suzanna Becker and Geoffrey E Hinton. 1992. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature* 355, 6356 (1992), 161–163.
- [2] Smriti Bhagat, Graham Cormode, and S Muthukrishnan. 2011. Node classification in social networks. In *Social network data analytics*. Springer, 115–148.
- [3] Guido Bugmann. 1998. Normalized Gaussian radial basis function networks. Neurocomputing 20, 1-3 (1998), 97–110.
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS).
- [5] Yukuo Cen, Xu Zou, Jianwei Zhang, Hongxia Yang, Jingren Zhou, and Jie Tang. 2019. Representation learning for attributed multiplex heterogeneous network. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 1358–1368.
- [6] Shiyu Chang, Wei Han, Jiliang Tang, Guo-Jun Qi, Charu C Aggarwal, and Thomas S Huang. 2015. Heterogeneous network embedding via deep architectures. In Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining. 119–128.
- [7] Xiaokai Chu, Xinxin Fan, Di Yao, Zhihua Zhu, Jianhui Huang, and Jingping Bi. 2019. Cross-network embedding for multi-network alignment. In *The world wide web conference*. 273–284.
- [8] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. 2020. Debiased Contrastive Learning. In *NeurIPS*.
- [9] Thomas M Cover and Joy A Thomas. 1991. Information theory and statistics. Elements of Information Theory 1, 1 (1991), 279–335.
- [10] Manlio De Domenico, Albert Solé-Ribalta, Emanuele Cozzo, Mikko Kivelä, Yamir Moreno, Mason A Porter, Sergio Gómez, and Alex Arenas. 2013. Mathematical formulation of multilayer networks. *Physical Review X* 3, 4 (2013), 041022.
- [11] Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. arXiv preprint arXiv:1903.02428 (2019).
- [12] Dongqi Fu, Zhe Xu, Bo Li, Hanghang Tong, and Jingrui He. 2020. A View-Adversarial Framework for Multi-View Network Embedding. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 2025–2028.
- [13] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining. 855–864.
- [14] Hakim Hafidi, Mounir Ghogho, Philippe Ciblat, and Ananthram Swami. 2020. Graphcl: Contrastive self-supervised learning of graph representations. arXiv preprint arXiv:2007.08025 (2020).
- [15] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In Proceedings of the 31st International Conference on Neural Information Processing Systems. 1025–1035.
- [16] Kaveh Hassani and Amir Hosein Khasahmadi. 2020. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*. PMLR, 4116–4126.
- [17] Dan Hendrycks and Kevin Gimpel. 2016. A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint arXiv:1610.02136 (2016).
- [18] Roger A Horn. 1990. The hadamard product. In Proc. Symp. Appl. Math, Vol. 40. 87–169.
- [19] Fenyu Hu, Yanqiao Zhu, Shu Wu, Liang Wang, and Tieniu Tan. 2019. Hierarchical Graph Convolutional Networks for Semi-supervised Node Classification. In IJCAL
- [20] Xinyu Huang, Dongming Chen, Tao Ren, and Dongqi Wang. 2021. A survey of community detection methods in multilayer networks. *Data Mining and Knowledge Discovery* 35, 1 (2021), 1–45.
- [21] Baoyu Jing, Chanyoung Park, and Hanghang Tong. 2021. Hdmi: High-order deep multiplex infomax. In Proceedings of the Web Conference 2021. 2414–2424.
- [22] Baoyu Jing, Hanghang Tong, and Yada Zhu. 2021. Network of Tensor Time Series. In Proceedings of the Web Conference 2021. 2425–2437.
- [23] Baoyu Jing, Yuejia Xiang, Xi Chen, Yu Chen, and Hanghang Tong. 2021. Graph-MVP: Multi-View Prototypical Contrastive Learning for Multiplex Graphs. arXiv preprint arXiv:2109.03560 (2021).
- [24] Baoyu Jing, Zeyu You, Tao Yang, Wei Fan, and Hanghang Tong. 2021. Multiplex Graph Neural Network for Extractive Text Summarization. arXiv preprint arXiv:2108.12870 (2021).
- [25] Jian Kang, Jingrui He, Ross Maciejewski, and Hanghang Tong. 2020. Inform: Individual fairness on graph mining. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 379–389.
- [26] Jian Kang and Hanghang Tong. 2021. Fair Graph Mining. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management. 4849–4852.
- [27] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In ICLR (Poster).

- [28] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016).
- [29] Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. arXiv preprint arXiv:1611.07308 (2016).
- [30] Yuval Kluger, Ronen Basri, Joseph T Chang, and Mark Gerstein. 2003. Spectral biclustering of microarray data: coclustering genes and conditions. *Genome research* 13, 4 (2003), 703–716.
- [31] Chaoyi Li and Yangsen Zhang. 2020. A personalized recommendation algorithm based on large-scale real micro-blog data. *Neural Computing and Applications* 32, 15 (2020), 11245–11252.
- [32] Hui-Jia Li, Lin Wang, Yan Zhang, and Matjaž Perc. 2020. Optimization of identifiability for efficient community detection. New Journal of Physics 22, 6 (2020), 063035.
- [33] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. 2019. Graph matching networks for learning the similarity of graph structured objects. In International conference on machine learning. PMLR, 3835–3845.
- [34] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. Journal of the American society for information science and technology 58, 7 (2007), 1019–1031.
- [35] Shuai Lin, Pan Zhou, Zi-Yuan Hu, Shuojia Wang, Ruihui Zhao, Yefeng Zheng, Liang Lin, Eric Xing, and Xiaodan Liang. 2021. Prototypical Graph Contrastive Learning. arXiv preprint arXiv:2106.09645 (2021).
- [36] Fanzhen Liu, Shan Xue, Jia Wu, Chuan Zhou, Wenbin Hu, Cecile Paris, Surya Nepal, Jian Yang, and S Yu Philip. 2020. Deep learning for community detection: progress, challenges and opportunities. In 29th International Joint Conference on Artificial Intelligence, IJCAI 2020. International Joint Conferences on Artificial Intelligence, 4981–4987.
- [37] Yixin Liu, Shirui Pan, Ming Jin, Chuan Zhou, Feng Xia, and Philip S Yu. 2021. Graph self-supervised learning: A survey. arXiv preprint arXiv:2103.00111 (2021).
- [38] Dongsheng Luo, Yuchen Bian, Yaowei Yan, Xiao Liu, Jun Huan, and Xiang Zhang. 2020. Local community detection in multiple networks. In Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining. 266–274.
- [39] Yao Ma, Zhaochun Ren, Ziheng Jiang, Jiliang Tang, and Dawei Yin. 2018. Multidimensional network embedding with hierarchical structure. In Proceedings of the eleventh ACM international conference on web search and data mining. 387–395.
- [40] Péter Mernyei and Cătălina Cangea. 2020. Wiki-cs: A wikipedia-based benchmark for graph neural networks. arXiv preprint arXiv:2007.02901 (2020).
- [41] Mark Newman. 2018. Networks. Oxford university press.
- [42] Mark EJ Newman and Michelle Girvan. 2004. Finding and evaluating community structure in networks. *Physical review E* 69, 2 (2004), 026113.
- [43] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018).
- [44] Chanyoung Park, Donghyun Kim, Jiawei Han, and Hwanjo Yu. 2020. Unsupervised attributed multiplex network embedding. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34. 5371–5378.
- [45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems 32 (2019), 8026–8037.
- [46] Zhen Peng, Wenbing Huang, Minnan Luo, Qinghua Zheng, Yu Rong, Tingyang Xu, and Junzhou Huang. 2020. Graph representation learning via graphical mutual information maximization. In *Proceedings of The Web Conference 2020*. 259–270.
- [47] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). 1532–1543.
- [48] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. 701–710.
- [49] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. 2020. Gcc: Graph contrastive coding for graph neural network pre-training. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 1150–1160.
- [50] Faisal Rahutomo, Teruaki Kitasuka, and Masayoshi Aritsugi. 2012. Semantic cosine similarity. In The 7th International Student Conference on Advanced Science and Technology ICAST, Vol. 4. 1.
- [51] William M Rand. 1971. Objective criteria for the evaluation of clustering methods. Journal of the American Statistical association 66, 336 (1971), 846–850.
- [52] Satu Elisa Schaeffer. 2007. Graph clustering. Computer science review 1, 1 (2007), 27–64.
- [53] John Scott. 1988. Social network analysis. Sociology 22, 1 (1988), 109-127.
- [54] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Pitfalls of graph neural network evaluation. arXiv preprint arXiv:1811.05868 (2018).
- [55] Yu Shi, Fangqiu Han, Xinwei He, Xinran He, Carl Yang, Jie Luo, and Jiawei Han. 2018. mvn2vec: Preservation and collaboration in multi-view network embedding. arXiv preprint arXiv:1801.06597 (2018).

- [56] Fan-Yun Sun, Jordon Hoffman, Vikas Verma, and Jian Tang. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In International Conference on Learning Representations.
- [57] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. Journal of machine learning research 9, 11 (2008).
- [58] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In International Conference on Learning Representations.
- [59] Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2018. Deep Graph Infomax. In International Conference on Learning Representations.
- [60] Dong Wang, Jiexun Li, Kaiquan Xu, and Yizhen Wu. 2017. Sentiment community detection: exploring sentiments and relationships in social networks. *Electronic Commerce Research* 17, 1 (2017), 103–132.
- [61] Ruijie Wang, Zijie Huang, Shengzhong Liu, Huajie Shao, Dongxin Liu, Jinyang Li, Tianshi Wang, Dachun Sun, Shuochao Yao, and Tarek Abdelzaher. 2021. DyDiff-VAE: A Dynamic Variational Framework for Information Diffusion Prediction. 163–172.
- [62] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *The World Wide Web Conference*. 2022–2032.
- [63] Xiao Wang, Nian Liu, Hui Han, and Chuan Shi. 2021. Self-supervised Heterogeneous Graph Neural Network with Co-contrastive Learning. arXiv preprint arXiv:2105.09111 (2021).
- [64] Lirong Wu, Haitao Lin, Zhangyang Gao, Cheng Tan, Stan Li, et al. 2021. Selfsupervised on Graphs: Contrastive, Generative, or Predictive. arXiv preprint arXiv:2105.07342 (2021).
- [65] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. 2020. What Should Not Be Contrastive in Contrastive Learning. In International Conference on Learning Representations.
- [66] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In International Conference on Learning Representations.
- [67] Zhe Xu. 2021. Dense subgraph detection on multi-layered networks. Ph.D. Dissertation.
- [68] Zhe Xu, Si Zhang, Yinglong Xia, Liang Xiong, Jiejun Xu, and Hanghang Tong. 2021. DESTINE: Dense Subgraph Detection on Multi-Layered Networks. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management. 3558–3562.
- [69] Yuchen Yan, Lihui Liu, Yikun Ban, Baoyu Jing, and Hanghang Tong. 2021. Dynamic Knowledge Graph Alignment. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 35. 4564–4572.
- [70] Xu Yang, Cheng Deng, Feng Zheng, Junchi Yan, and Wei Liu. 2019. Deep spectral clustering using dual autoencoder network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4066–4075.
- [71] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. Advances in Neural Information Processing Systems 33 (2020), 5812–5823.
- [72] Hanlin Zhang, Shuai Lin, Weiyang Liu, Pan Zhou, Jian Tang, Xiaodan Liang, and Eric P Xing. 2020. Iterative graph self-distillation. arXiv preprint arXiv:2010.12609 (2020).
- [73] Hongming Zhang, Liwei Qiu, Lingling Yi, and Yangqiu Song. [n.d.]. Scalable Multiplex Network Embedding.
- [74] Shichang Zhang, Ziniu Hu, Arjun Subramonian, and Yizhou Sun. 2020. Motif-driven contrastive learning of graph representations. arXiv preprint arXiv:2012.12533 (2020).
- [75] Tianqi Zhang, Yun Xiong, Jiawei Zhang, Yao Zhang, Yizhu Jiao, and Yangyong Zhu. 2020. CommDGI: Community Detection Oriented Deep Graph Infomax. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 1843–1852.
- [76] Jiang Zhu, Bai Wang, Bin Wu, and Weiyu Zhang. 2017. Emotional community detection in social network. *IEICE Transactions on Information and Systems* 100, 10 (2017), 2515–2525.
- [77] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2020. Deep graph contrastive representation learning. arXiv preprint arXiv:2006.04131 (2020).
- [78] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2021. Graph contrastive learning with adaptive augmentation. In Proceedings of the Web Conference 2021. 2069–2080.

Graph Communal Contrastive Learning

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

IMPLEMENTATION DETAILS Α

We conduct all of our experiments with PyTorch Geometric 1.7.2 [11] and PyTorch 1.8.1 [45] on a NVIDIA Tesla V100 GPU (with 32 GB available video memory). All single-viewed graph datasets used in experiments are available in PyTorch Geometric libraries⁵, and all multiplex graph datasets are available in ⁶. Our hyperparameter specifications are listed in Table 9. Those hyperparameters are determined empirically by gird search based on the settings of [78].

Table 9: Hyperparameter specifications.

Hyperparameter	Amazon- Computers	Amazon -Photo	Coauthor -CS	WikiCS	IMDB	DBLP
Training epochs	2900	2900 /1400	600	1000	50	50
Hidden-layer dimension	2558	256	7105	1368	8096	4256
Representation dimension	512	128	300	384	2048	128
Learning rate	0.01	0.1	0.0005	0.01	0.0001	0.001
Activation function	RReLU	ReLU	RReLU	PReLU	ReLU	RReLU
η	1000	500	500	500	50	50
τ	0.2	0.3	0.4	0.4	1.0	1.0
λ_w	0.9	0.1	0.9	1.0	0.5	0.5
Y	2e-5	8e-5	8e-5	8e-5	2e-4	2e-4
p_v	0.25	0.1	0.3	0.1	0.1	0.1
p_e	0.45	0.4	0.2	0.2	0.0	0.0

B NUMBER OF COMMUNITIES

We compare varying numbers of communities in the vicinity of the numbers of classes on each single-viewed graph dataset (refer to Fig. 6). Although the performances remain stable for nearly all settings, we recommend to set the number of communities according to the actual number of classes.



Figure 6: Comparison of community number on node classification.

C PARALLELIZATION OF DENSE COMMUNITY AGGREGATION

In DeCA, The intra-community density objective is computationally costly and hard to vectorize since the expected edge density is dependent on the community choice. To address this problem, we derive a lower bound of *D_{intra}*:

$$D_{intra} \geq \frac{1}{N} \sum_{i,j} \sum_{k} \left[\boldsymbol{A}[i,j] - \max_{\kappa} \{d(\kappa)\} \right] \boldsymbol{R}[i,k] \boldsymbol{R}[j,k]$$
$$= \frac{1}{N} \sum_{i,j} \sum_{k} \widetilde{\boldsymbol{A}}[i,j] \boldsymbol{R}[i,k] \boldsymbol{R}[j,k]$$
$$= \widetilde{D}_{intra}, \tag{20}$$

where $\widetilde{A} = A - \max_{\kappa} d(\kappa)I$ is the extended adjacency matrix, and $\tilde{D}_{intra} = \inf D_{intra}$ is the lower bound. We use \tilde{D}_{intra} to replace Dintra in Eq. 9.

Next, we employ the community density matrix F = R'AR and $\widetilde{F} = R'\widetilde{A}R$ to vectorize the objective D_{inter} and \widetilde{D}_{intra} . The entry of *F* is $F[u, v] = \sum_{i} R[i, u] \cdot (AR)[i, v] = \sum_{i,j} A[i, j]R[i, u]R[j, v]$, which naturally accords with the form of Eq. 7 and Eq. 8. Therefore, these two objectives can be re-formalized as

$$\widetilde{D}_{intra} = \frac{1}{N} tr(\widetilde{F}),$$

(21)

$$D_{inter} = \frac{1}{N(N-1)} \left[\sum_{i,j} F[i,j] - tr(F) \right].$$
(22)

Finally, the vectorized DeCA objective is

~

 $l_{DeCA}(\mathbf{R}) = \lambda_w D_{inter} - \widetilde{D}_{intra}$

and

$$=\frac{\lambda_{w}}{N(N-1)}\left[\sum_{i,j}F[i,j]-tr(F)\right]-\frac{1}{N}tr(\widetilde{F}), \qquad (23)$$

which can be directly computed by the CUDA operators [45].

⁵https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html ⁶https://www.dropbox.com/s/480e7shjq0ih151/data.tar.gz?dl=0