

---

# Efficient Contrastive Learning for Fast and Accurate Inference on Graphs

---

Teng Xiao<sup>1</sup> Huaisheng Zhu<sup>1</sup> Zhiwei Zhang<sup>1</sup> Zhimeng Guo<sup>1</sup>  
Charu C. Aggarwal<sup>2</sup> Suhang Wang<sup>1</sup> Vasant G Honavar<sup>1</sup>

## Abstract

Graph contrastive learning has made remarkable advances in settings where there is a scarcity of task-specific labels. Despite these advances, the significant computational overhead for representation inference incurred by existing methods that rely on intensive message passing makes them unsuitable for latency-constrained applications. In this paper, we present GraphECL, a simple and efficient contrastive learning method for fast inference on graphs. GraphECL does away with the need for expensive message passing during inference. Specifically, it introduces a novel coupling of the MLP and GNN models, where the former learns to computationally efficiently mimic the computations performed by the latter. We provide a theoretical analysis showing why MLP can capture essential structural information in neighbors well enough to match the performance of GNN in downstream tasks. The extensive experiments on widely used real-world benchmarks that show that GraphECL achieves superior performance and inference efficiency compared to state-of-the-art graph contrastive learning (GCL) methods on homophilous and heterophilous graphs. Code is available at: <https://github.com/tengxiao1/GraphECL>.

## 1. Introduction

Over the past decade, there has been considerable interest in graph learning problems such as node classification, link prediction, and graph classification (Grover & Leskovec, 2016; Cui et al., 2018; Xu et al., 2019; Wu et al., 2020; Xiao et al., 2021; Chen et al., 2022; Xiao et al., 2024). Graph contrastive learning (GCL) has recently emerged as an attractive approach to graph representation learning in settings where

there is a scarcity of task-specific labels (Zhu et al., 2021; Thakoor et al., 2021). On many key benchmarks, GCL has been shown to achieve performance that is competitive with or superior to that of state-of-the-art methods trained using ground truth labels (Zhu et al., 2021; Thakoor et al., 2021).

However, the significant computational overhead incurred by existing GCL methods, which rely on message passing for representation inference, limits their usefulness in latency-constrained applications. In particular, we observe that the state-of-the-art GCL methods achieve their superior performance using a graph neural network (GNN) encoder (see Figure 1). The message passing in GNN involves fetching the topology and features of numerous neighboring nodes to perform inference on a target node, which is computation-intensive during inference. Hence, there is an urgent need for inference-efficient alternatives to the state-of-the-art GCL methods. Recent work has begun to tackle the inference latency of GNN (Zheng et al., 2021; Zhang et al., 2021b; Tian et al., 2022; Wu et al., 2023), e.g., using knowledge distillation (KD) (Hinton et al., 2015) to learn an inference-efficient *student* MLP to mimic the output of a *teacher* GNN. However, they require task-specific labels to first train a good teacher GNN, limiting their applicability to GCL in settings where there is a lack of task-specific labels.

To the best of our knowledge, the following critical question, with important implications for real-world latency-constrained applications of GCL, remains unanswered: *How can we design a new GCL algorithm that outperforms state-of-the-art GCL methods on downstream tasks while avoiding high inference latency?* To answer this question, we present GraphECL, a *simple, effective, and efficient* contrastive regime on graphs. Specifically, to capture the graph structure of the nodes and achieve fast inference, GraphECL introduces a cross-model contrastive architecture in which positive examples consist of cross-model pairs (e.g., MLP-GNN) directly derived from neighborhood relations extracted from the graph. These positive samples are obtained from the representations of MLP and GNN of central nodes and their neighbors, respectively. This simple architecture allows GraphECL to benefit from the graph structure during training via GNN while using MLP to avoid relying on the graph structure during inference. Based on this cross-model architecture, we introduce a novel *general-*

---

<sup>1</sup>Penn State University, USA <sup>2</sup>IBM T. J. Watson Research Center, USA. Correspondence to: Teng Xiao <tengxiao@psu.edu>.

Figure 1. Inference latency v.s Accuracy on Pubmed. For GCL baselines, we test them on using MLP and GNN backbones. We can observe that SOTA methods require GNN as the encoder to achieve good performance, which is computation-intensive during inference. GraphECL is fast with even higher accuracy. GraphECL uses regularized contrastive loss which facilitates the learning of a computationally efficient MLP encoder, allowing the resulting model to effectively capture graph structural information so as to match the performance of state-of-the-art GCL methods on benchmark datasets, but without their prohibitive computational cost during inference.

**Key Contributions.** (i) We identify the key limitations of current GCL methods that limit their applicability to representation learning in latency-constrained real-world applications that require fast inference. (ii) We design GraphECL, a novel coupling of MLP and GNN models where the former learns to computationally efficiently mimic the computations performed by the latter. We show that the resulting model can effectively learn graph structural information and conduct fast inference with a simple MLP. (iii) We present theoretical analyses that offer insight into how GraphECL gradually encodes useful structural information using an MLP. Specifically, we show how GraphECL can theoretically achieve good generalization performance on downstream tasks. (iv) We demonstrate through extensive experiments that GraphECL can achieve ultra-fast inference speed and superior performance on downstream tasks simultaneously. Specifically, we show that the proposed GraphECL can run significantly faster than baselines on graphs, making it especially useful for latency-constrained applications where fast inference is a key requirement.

## 2. Related Work

**Graph Contrastive Learning.** GCL has emerged as a class of effective methods for learning useful representations from unlabeled graph data (Večerković et al., 2018; Zhu et al., 2021; Guo et al., 2024; Suresh et al., 2021; Zhang et al., 2021a; Zhu et al., 2024; Meng et al., 2019). Some authors have proposed freeing GCL from the need for negative samples (Thakoor et al., 2021; Zhang et al., 2021a) or even the need for graph augmentation (Xiao et al., 2022; Xiao &

Wang, 2021; Lee et al., 2022; Zhang et al., 2022). Others have explored approaches to accelerate GCL training (Zheng et al., 2022b; Yang et al., 2022a; Han et al., 2022). Despite this progress, current methods of GCL incur significant computational overhead during inference, which limits their usefulness in latency-constrained real-world applications. This is largely due to their need to fetch neighbors and their associated features for a target node while performing inference (Zhang et al., 2021b). In this paper, we aim to address this limitation by avoiding the need for expensive message passing during inference by coupling MLP and GNN models so that the former can learn to computationally efficiently mimic the computations performed by the latter.

**Learning MLPs on Graphs.** Our work is also related to graph-regularized MLP (Yang et al., 2016; Hu et al., 2021; Yang et al., 2021b; Liu et al., 2020), which incorporates graph structure into MLPs through various auxiliary regularization terms inspired by traditional network embedding methods (Hamilton et al., 2017b; Grover & Leskovec, 2016; Tang et al., 2015). By implicitly encoding structural information into MLPs, one can enhance the representational power of MLP encoders while maintaining fast inference. It is worth noting that these methods, despite their differences, share a reliance on the strong homophily assumption (McPherson et al., 2001), which posits that one-hop neighbors of nodes that are linked should exhibit similar latent representations. Consequently, graph-regularized MLP significantly falls short of the performance achievable by GCL methods (Večerković et al., 2018; You et al., 2020), as shown in our experiments. GraphECL on the other hand, is designed to match the performance of state-of-the-art GCL methods, while achieving significantly faster inference.

**Knowledge Distillation on Graphs.** Knowledge distillation on graphs, which aims to distill pre-trained teacher GNNs into smaller student MLPs, has recently garnered significant attention (Yang et al., 2020; 2022b; Yan et al., 2020; Yang et al., 2021a; Joshi et al., 2022). Since student GNNs still require time-consuming message passing in inference, recent studies (Yang et al., 2023; Zhang et al., 2021b; Zheng et al., 2021; Tian et al., 2022; Wu et al., 2023) have shifted their focus towards GNN-MLP distillation. This involves learning an inference-efficient student MLP by distilling knowledge from the teacher GNN. However, these methods typically rely on task-specific labels to train the teacher GNN, which can be challenging in real-world scenarios where labels are often inaccessible. In contrast, our work aims to develop an inference-efficient and structure-aware MLP for faster inference in settings where task-specific labels are unavailable.

## 3. Preliminaries

**Notations and Problem Setup.** The input graph is denoted as  $G = (V; E)$ , where  $V = \{v_1; \dots; v_{|V|}\}$  is a set of  $|V|$

Figure 2. Existing contrastive schemes and GraphECL. (a) and (c) rely on invariant assumptions, aiming to learn augment-invariant representations of the same node. (b) is based on homophily assumptions, forcing neighboring nodes to exhibit same representations. In contrast, (d) showcases GraphECL, which achieves significant inference efficiency and strong performance using learned

nodes and  $E$  denotes the set of edges. Each edge  $e \in E$  denotes a link between  $v_i$  and  $v_j$ . We use  $X \in \mathbb{R}^{V \times D}$  to denote the node attributes, where the row of  $X$ , i.e.,  $x_i$ , is the attribute vector of node  $v_i$ . The graph structure can be characterized by its adjacency matrix  $A \in \{0, 1\}^{V \times V}$ , where  $A_{i,j} = 1$  if there exists an edge  $e_{i,j} \in E$ , and  $A_{i,j} = 0$  otherwise. Then, the graph  $G$  can be also denoted as a tuple of matrices:  $G = (X; A)$ . Given  $G$ , our goal is to learn an efficient MLP encoder denoted by  $f_M$  with only attributes  $X$  as input, so that the inferred representation for node  $v$ :  $v = f_M(X)[v] \in \mathbb{R}^K$  is useful for downstream tasks. For brevity, in what follows, we omit the input and use  $f_M(v)$  to denote  $v$ 's representation from MLP.

Graph Contrastive Learning (GCL) with Augmentations. GCL aims to learn representations (Trivedi et al., 2022; Velicković et al., 2018; Zhu et al., 2021; You et al., 2020; Suresh et al., 2021) by contrasting augmented views of nodes as presented in Figure 2 (a). Thus, for a given node  $v$ , its representation in an augmented view is trained to be similar to the representation of the same node from another augmented view, while being distinct from the representations of other nodes, which serve as negative samples. Given two views  $G_1$  and  $G_2$ , a widely-used contrastive objective is:

$$L_{GCL} = \frac{1}{|V|} \sum_{v \in V} \log \left( \frac{\exp(f_G(v^1) \cdot f_G(v^2) / \tau)}{\exp(f_G(v^1) \cdot f_G(v^2) / \tau) + \sum_{v' \in V} \exp(f_G(v^1) \cdot f_G(v') / \tau)} \right); \quad (1)$$

Here  $f_G(v^1) = f_G(G_1)[v]$  and  $f_G(v^2) = f_G(G_2)[v]$  are GNN representations of the same node from two views, where  $f_G$  denote the GNN encoder.  $V$  is the set of negative samples from inter- or intra- augmented view (Zhu et al., 2020).  $\tau$  is the temperature. Although GCL with augmentations has achieved remarkable success, we note that such methods predominantly rely on the GNN encoder to capture structural invariances in different augmented views of the graph. This reliance is further discussed in Section 4.1 and results in substantial computational overhead during inference compared to MLP, as shown in Figure 1.

Graph-regularized MLP. Graph-MLP (Yang et al., 2016; Hu et al., 2021; Yang et al., 2021b; Liu et al., 2020) proposes to avoid the need for GNN neighbor fetching by learn-

ing an inference-efficient MLP model with a neighbor contrastive loss inspired by traditional graph embedding methods (Hamilton et al., 2017b; Grover & Leskovec, 2016; Tang et al., 2015). All such methods essentially minimize the following contrastive loss over neighbors in the graph:

$$L_{GR} = \frac{1}{|V|} \sum_{v \in V} \frac{1}{|N(v)|} \sum_{u \in N(v)} \log \left( \frac{\exp(f_M(v) \cdot f_M(u) / \tau)}{\exp(f_M(v) \cdot f_M(u) / \tau) + \sum_{v' \in V} \exp(f_M(v) \cdot f_M(v') / \tau)} \right); \quad (2)$$

where  $f_M(v)$ ,  $f_M(u)$  and  $f_M(v')$  are projected representations by MLP of nodes  $v$ ,  $u$  and  $v'$ , respectively.  $N(v)$  denotes the set of positive examples containing local neighborhoods of the node  $v$  and  $V$  denotes the set of negative examples that are randomly sampled from  $V$ . This approach is illustrated in Figure 2 (b). Despite its improved inference efficiency due to its exclusive use of MLP, this approach takes a significant hit in performance compared to GCL with augmentations (See Figure 4). Moreover, this scheme over-emphasizes homophily, assuming that nodes that are linked in the graph should have similar representations in the latent space, at the expense of structural information (You et al., 2020), making it difficult to generalize to graphs with heterophily (Lim et al., 2021). Table 6 in Appendix A details comparisons between current graph contrastive schemes and our GraphECL in terms of design assumptions, effectiveness (or representational power) and inference efficiency.

## 4. Efficient Graph Contrastive Learning

We proceed to introduce GraphECL, which aims to dramatically speed up inference while matching the performance of GCL. GraphECL adopts a cross-model contrastive architecture, wherein we design an asymmetric GNN-MLP architecture for the nodes and their neighbors to extract effective positive and negative examples in contrastive training (Section 4.1). To capture structural information in the graph, we introduce a generalized contrastive loss that extends the classic InfoNCE loss (Chen et al., 2020) from independent instance discrimination over augmentations to neighborhood contrast over graph structures, taking into account meaningful distance between neighboring nodes. Finally, we provide a theoretical analysis to show

Figure 3. Current GCL methods that employ GNN-MLP architecture (where MLP is used for inference) exhibit a significant performance decay compared to those using GNN-GNN architecture (where GNNs are used for inference). We illustrate Pubmed as an example, though we observe the similar trend in other datasets, that GraphECL can encode structural information to ensure good performance on downstream tasks (Section 4.2).

#### 4.1. Simple Cross-model Contrastive Learning

As shown in Figure 1, the state-of-the-art GCL methods incur substantial computational overhead and hence inference latency due to the layer-wise message passing in GNN encoders. A straightforward idea is to replace GNN encoders with the MLP encoders to speed up inference. However, as seen in Figure 1, the performance of current GCL methods using MLP instead of GNN is significantly worse than that of those using GNN. These results are consistent with what our expectations in that, while replacing GNN with MLP eliminates message passing and hence speeds up inference, does so by ignoring critical structural information from the graph. Thus, the encoder GNN plays a critical role in the success of GCL based on graph augmentations.

**Cross-model Contrastive Architecture** To address this limitation, we first introduce a simple cross-model architecture of GraphECL. As Figure 1 shows, using MLP as the encoder for GCL achieves substantial speedup in inference but does so at the cost of substantial drop in performance. Conversely, utilizing GNN as the encoder yields superior performance but at the cost of significant slowdown during inference. Our solution to resolving this dilemma is elegantly simple, yet, as we will demonstrate, remarkably effective. Specifically, we employ a cross-model architecture with two encoders, one of which is GNN and the other an MLP. GNN in this architecture is exclusively dedicated to extracting and encoding structural information from the graph during the learning phase, whereas MLP is used during the inference process to circumvent the need for computationally expensive message passing.

Can we directly apply this architecture to speed up inference in GCL? For instance, can we minimize the following cross-model InfoNCE-style loss, referred to as MLP-Augmented

GCL loss for the architecture shown in Figure 2 (c)?

$$L_{MA} = \frac{1}{|V|} \sum_{v \in V} \log \hat{\rho}(v); \text{ where } \hat{\rho}(v) = \frac{\exp(f_G(v^1) \cdot f_M(v^2))}{\exp(f_G(v^1) \cdot f_M(v^2)) + \sum_{v' \in V} \exp(f_G(v^1) \cdot f_M(v'))}; \quad (3)$$

Unfortunately, the answer to this question is negative. As shown in Figure 3, even with the cross-model architecture, current GCL methods take a significant performance loss compared to those that use GNN encoders in Equation (3). In what follows, we introduce an effective and efficient contrastive learning loss that addresses this problem.

**Efficient Contrastive Learning Loss on Graphs.** Before proceeding to introduce the proposed contrastive loss, we first motivate it. Existing state-of-the-art GCL methods adopt graph augmentations that emphasize similarities in the encoding of a same node in different "augmented views", using GNN as the encoder (See Figure 2). In contrast, we want GraphECL to avoid relying on graph augmentations, but instead, learn an MLP representation of node by extracting and encoding its neighborhood structure and features from GNN. In particular, positive pairs in GCL are generated by random graph augmentations of the same node. In contrast, positive examples of GraphECL are cross-model pairs (e.g. MLP-GNN) directly provided by neighborhood relations present in graphs. These positive examples are obtained through MLP and GNN representations, respectively, of nodes and their neighbors. The preceding suggests a contrastive loss for GraphECL, minimizing which has the effect of pushing the MLP representation of each node closer to the GNN representations of its neighbors:

$$L_{ECL} = \frac{1}{|V|} \sum_{v \in V} \frac{1}{|N(v)|} \sum_{u \in N(v)} \log \hat{\rho}(v); \text{ where } \hat{\rho}(v) = \frac{\exp(f_M(v) \cdot f_G(u))}{\exp(f_G(u) \cdot f_G(v)) + \sum_{v' \in V} \exp(f_M(v) \cdot f_G(v'))}; \quad (4)$$

where  $f_M(v)$  and  $f_G(u)$  are the L2-normalized representations obtained from the MLP and GNN encoders, respectively, of node  $v$  and its neighbor  $u$ . Here,  $(f_G(u), f_G(v))$  and  $(f_M(v), f_G(v'))$  represent intra-model and inter-model negative pairs, respectively.  $\rho$  is independently sampled as a negative example and serves as a hyperparameter to control the balance between the two types of negative pairs. For large graphs, we randomly sample negative pairs for each node as an efficient approximation.

We highlight the benefits of using this simple alignment loss. First, the MLP encoder  $f_M$  can effectively preserve the local neighborhood distribution captured by GNN encoder  $f_G$  without the need for graph augmentation. GraphECL encodes the latent distributions (representations of GNN) of neighborhoods into the representation of central node from

Figure 4. (a) The toy graph where the color denotes node's semantic class. (b) Uni-model contrastive learning in GR-MLP encourages one-hop neighbors to have similar representations (c) Multi-model contrastive objective GraphECL is not based on the one-hop homophily assumption but automatically captures graph structures based on different graphs beyond homophily. Thus GraphECL exhibits robustness and generalizability on both homophilic and heterophilic graphs (See section 5).

MLP, enabling MLP to implicitly encode the structural information captured by GNN. As latent neighbors' representations encode high-order information through multi-layer message passing of GNN, MLP effectively distills the high-order structural information from the GNN. Second, in contrast to the GR-MLP model, which performs uni-model contrastive learning shown in Figure 4 (b), GraphECL aims to push cross-model contrastive learning, pushing the representation of a node  $f_M(v)$  and that of its neighbors  $f_G(u)$  close to each other. In particular, GraphECL does not necessarily imply the learned MLP representation  $f_M(v); f_M(u)$  become identical (homophily assumption). In other words, the multi-model signals ensure that node pairs with the same same neighborhood context serve as positive pairs for contrastive learning as illustrated in Figure 4 (c).

Interpretation.  $L_{ECL}$  is a simple yet very effective generalization of the popular InfoNCE loss in Equation (3) from uni-model instance discrimination over augmentations to cross-model contrast over graph neighbors. During the learning process, cross-model positive pairs of neighbors  $f_M(v), f_G(u)$  are pulled together in the latent space, while inter-modal  $(f_M(v), f_G(v))$  and intra-model negative pairs  $(f_G(u), f_G(v))$  are pushed apart. We empirically demonstrate in Section 5 that GraphECL generalizes as well as state-of-the-art GCL methods during the inference, with the additional benefit of significantly faster inference.

#### 4.2. Theoretical Analysis

In this section, we provide theoretical evidence to support the design of our simplified GraphECL. All detailed proofs can be found in Appendix B. We denote the normalized adjacency matrix  $D^{-1}A$ , with  $D$  being the diagonal degree matrix. We define the two representation metrics  $M$  and  $G$  where the  $v$ -th row  $(M)_v = f_M(v)$  and the  $u$ -th row  $(G)_u = f_G(u)$  represent the corresponding encoded representations from MLP and GNN respectively. Let  $A = \exp(MG^T)$  is the estimated adjacency matrix based on representation similarity  $D = \text{deg}(A)$  is the diagonal matrix, whose element  $(A)_{ij}$  is the sum of the

of  $A$ . Next, we reveal the stationary point of the learning dynamics of GraphECL, which implies the equilibrium as:

Theorem 4.1. The learning dynamics of the MLP encoder  $f_M$  with efficient contrastive loss ( $\beta = 1$ ) in Equation (4) saturates when the true normalized adjacency and the estimated normalized adjacency matrices agree:  $D^{-1}A = D^{-1}\hat{A}$ , implying that, for  $\forall v; u \in V$ , we have:

$$P_n(u|j|v) = P_f(u|j|v), \quad P = \frac{\exp(f_M(v)^T f_G(u))}{\sum_{v' \in V} \exp(f_M(v)^T f_G(v'))}; \quad (5)$$

where  $P_n(u|j|v)$  is the 1-hop neighborhood distribution (i.e., the  $v$ -th row of the normalized adjacency matrix) and  $P_f(u|j|v)$  is the estimated neighborhood distribution.

Theorem 4.1 implies that GraphECL essentially learns a probabilistic model based on cross-modal encoders to predict the conditional 1-hop neighborhood distribution. Specifically, our assumption is more general than the homophily assumption. Even in heterophilic graphs, two nodes of the same semantic class tend to share similar structural roles, i.e., the 1-hop neighborhood context as shown in (Ma et al., 2021; Xiao et al., 2023) and statistics in Appendix C.2.

We also establish formal guarantees for the generalization of GraphECL on downstream tasks for learned MLP and GNN encoders. Without loss of generality, we use the linear probing task as an example. In this task, we train a linear classifier to predict class labels  $Y$  based on the MLP representation  $f_M$  using  $g_{f;W}(v) = \arg \max_{c \in [C]} (f_M(v)^T W)_c$ , where  $W \in \mathbb{R}^{K \times C}$  represents the weight matrix.

Theorem 4.2. Let  $f_M$  be the global minimum of generalized contrastive loss ( $\beta = 1$ ) in Equation (4) and  $\lambda_1(v)$  denote the label of  $v$ .  $\lambda_1, \dots, \lambda_N$  are the eigenvalues with descending order of the normalized adjacency matrix  $D^{-1}A$ . Then, the linear probing error  $\text{err}_M$  is upper-bounded by:

$$E(f_M), \min_W \frac{1}{|V|} \sum_{v \in V} \mathbb{1}[g_{f;W}(v) \neq y(v)] \leq \frac{1}{1 - \lambda_{K+1}}; \quad (6)$$

where  $\lambda_{K+1} = \frac{1}{|V|} \sum_{v \in V} \mathbb{1}[N(v) \neq N(v)]$  and  $K$  is the dimensionality of the representation.

This theorem establishes a significant relationship between the downstream error in learned representations and two crucial factors: the parameter  $\beta$  and the  $(K+1)$ -th largest eigenvalue. Remarkably,  $\lambda_{K+1}$  coincides precisely with the node homophily ratio metric (Pei et al., 2019; Lim et al., 2021). This metric calculates the proportion of a node's neighbors that share the same class label and then averages these values across all nodes within the graph. Homophilous graphs ( $\lambda_{K+1} \rightarrow 1$ ), exhibit a tendency for nodes to connect with others of the same class, while heterophilic graphs ( $\lambda_{K+1} \rightarrow 0$ ), display a preference for connections across different classes. This theorem shows that graphs characterized by a low homophily value (i.e., heterophilic graphs) may require a larger

Table 1. Node classification results (%) under the transductive setting on benchmarking homophilic and heterophilic graphs.

Datasets	Cora	Citeseer	Pubmed	Photo	WikiCS	Flickr	Cornell	Wisconsin	Texas	Actor
Graph-MLP	76.70 $\pm$ 0.18	70.30 $\pm$ 0.27	78.70 $\pm$ 0.33	89.59 $\pm$ 0.45	71.75 $\pm$ 0.15	41.33 $\pm$ 0.25	42.65 $\pm$ 2.21	57.96 $\pm$ 1.11	60.22 $\pm$ 1.76	25.66 $\pm$ 0.77
VGAE	76.30 $\pm$ 0.21	66.80 $\pm$ 0.23	75.80 $\pm$ 0.40	91.50 $\pm$ 0.20	72.19 $\pm$ 0.31	40.71 $\pm$ 0.22	48.73 $\pm$ 4.19	55.67 $\pm$ 1.37	50.27 $\pm$ 2.21	26.99 $\pm$ 1.56
DGI	82.30 $\pm$ 0.60	71.80 $\pm$ 0.70	76.80 $\pm$ 0.60	91.61 $\pm$ 0.22	75.35 $\pm$ 0.14	44.70 $\pm$ 0.26	45.33 $\pm$ 6.11	55.21 $\pm$ 1.02	58.53 $\pm$ 2.98	28.30 $\pm$ 0.76
GCA	82.93 $\pm$ 0.42	72.19 $\pm$ 0.31	80.79 $\pm$ 0.45	91.70 $\pm$ 0.10	78.35 $\pm$ 0.05	46.10 $\pm$ 0.19	52.31 $\pm$ 1.09	59.55 $\pm$ 0.81	52.92 $\pm$ 0.46	28.77 $\pm$ 0.29
SUGRL	83.40 $\pm$ 0.50	73.00 $\pm$ 0.40	81.90 $\pm$ 0.30	93.07 $\pm$ 0.15	79.83 $\pm$ 0.31	46.22 $\pm$ 0.31	50.18 $\pm$ 0.30	61.31 $\pm$ 2.07	57.88 $\pm$ 2.21	30.31 $\pm$ 0.82
BGRL	82.70 $\pm$ 0.60	71.10 $\pm$ 0.80	79.60 $\pm$ 0.50	92.90 $\pm$ 0.30	79.98 $\pm$ 0.10	45.33 $\pm$ 0.19	50.33 $\pm$ 2.29	51.23 $\pm$ 1.17	52.77 $\pm$ 1.98	28.80 $\pm$ 0.54
CCA-SSG	84.00 $\pm$ 0.40	73.10 $\pm$ 0.30	81.00 $\pm$ 0.40	93.14 $\pm$ 0.14	79.31 $\pm$ 0.21	47.54 $\pm$ 0.14	52.17 $\pm$ 1.04	58.46 $\pm$ 0.96	59.89 $\pm$ 0.78	27.82 $\pm$ 0.60
GGD	83.90 $\pm$ 0.40	73.00 $\pm$ 0.60	81.30 $\pm$ 0.30	92.50 $\pm$ 0.60	78.72 $\pm$ 0.61	46.33 $\pm$ 0.20	51.46 $\pm$ 0.33	58.93 $\pm$ 0.65	60.17 $\pm$ 0.52	28.27 $\pm$ 0.23
SGCL	82.97 $\pm$ 0.20	72.58 $\pm$ 0.25	81.25 $\pm$ 0.32	93.46 $\pm$ 0.30	79.85 $\pm$ 0.53	46.35 $\pm$ 0.18	53.28 $\pm$ 1.37	59.93 $\pm$ 0.75	61.26 $\pm$ 0.65	26.51 $\pm$ 0.47
GraphACL	84.20 $\pm$ 0.31	73.63 $\pm$ 0.22	82.02 $\pm$ 0.15	93.31 $\pm$ 0.19	78.75 $\pm$ 0.36	46.95 $\pm$ 0.27	59.33 $\pm$ 1.48	69.22 $\pm$ 0.40	71.08 $\pm$ 0.34	30.03 $\pm$ 0.13
AF-GCL	83.16 $\pm$ 0.13	71.96 $\pm$ 0.42	79.16 $\pm$ 0.73	92.49 $\pm$ 0.31	79.01 $\pm$ 0.51	46.95 $\pm$ 0.33	52.29 $\pm$ 1.21	60.12 $\pm$ 0.39	59.81 $\pm$ 1.33	28.94 $\pm$ 0.69
AFGRL	81.30 $\pm$ 0.20	68.70 $\pm$ 0.30	80.60 $\pm$ 0.40	93.22 $\pm$ 0.28	77.62 $\pm$ 0.49	46.81 $\pm$ 0.20	55.37 $\pm$ 3.56	63.21 $\pm$ 1.55	60.35 $\pm$ 1.05	30.31 $\pm$ 0.95
GraphECL	84.25 $\pm$ 0.05	73.15 $\pm$ 0.41	82.21 $\pm$ 0.05	94.22 $\pm$ 0.11	80.17 $\pm$ 0.15	48.49 $\pm$ 0.15	69.19 $\pm$ 6.86	79.41 $\pm$ 2.19	75.95 $\pm$ 5.33	35.80 $\pm$ 0.89

Table 2. Node classification results on large-scale graphs.

Datasets	snap-patents	ogbn-arxiv	ogbn-papers100M
BGRL	24.33 $\pm$ 0.13	71.64 $\pm$ 0.24	58.75 $\pm$ 0.31
CCA-SSG	25.51 $\pm$ 0.46	71.21 $\pm$ 0.20	57.31 $\pm$ 0.18
GraphACL	26.18 $\pm$ 0.39	71.72 $\pm$ 0.26	59.35 $\pm$ 0.27
SUGRL	25.11 $\pm$ 0.32	69.30 $\pm$ 0.20	60.31 $\pm$ 0.22
SGCL	24.91 $\pm$ 0.46	70.99 $\pm$ 0.09	59.96 $\pm$ 0.37
GraphECL	27.22 $\pm$ 0.06	71.75 $\pm$ 0.22	61.45 $\pm$ 0.31

representation dimension, i.e., smaller  $k$  ( $k+1$ )-th largest  $k_{k+1}$  to effectively bound the downstream error.

## 5. Experiments

**Datasets.** We use established benchmarks for homophilic graphs: Cora, Citeseer, Pubmed, Photo, WikiCS, and Flickr, and for heterophilic graphs: Cornell, Wisconsin, Texas, and Actor. Additionally, we evaluate GraphECL on large-scale graphs, specifically the heterophilic Snap-patents, and homophilic Ognb-arxiv and Ognb-papers100M. In all datasets, we use the standard splits used in prior studies (Zhang et al., 2021a). The dataset details, splits, and statistics are in C.3.

**Baselines.** We compare GraphECL with the following graph contrastive learning methods: Graph-MLP (Hu et al., 2021), VGAE (Kipf & Welling, 2016), DGI (Velicković et al., 2018), GCA (Zhu et al., 2021), SUGRL (Mo et al., 2022), BGRL (Thakoor et al., 2021), CCA-SSG (Zhang et al., 2021a), AF-GCL (Wang et al., 2022), AFGRL (Lee et al., 2022), GGD (Zheng et al., 2022b), GraphACL (Xiao et al., 2023), and SGCL (Sun et al., 2024).

**Evaluation Protocol.** Following (Velicković et al., 2018; Thakoor et al., 2021), we consider three downstream tasks: node classification and graph classification. We use standard linear-evaluation protocol, where a linear classifier is trained on top of the frozen node or graph representations, and test accuracy is used as a proxy for representation quality.

**Transductive vs. Inductive.** Evaluation of node representations obtained using unsupervised learning through transductive node classification is a prevalent practice in the GCL literature. However, such evaluation neglects the scenarios of inferring representations for previously unseen nodes. Thus, it can not evaluate the real-world applicability of a deployed model, which often requires the inference of representations of novel nodes. Hence, following (Zhang et al., 2021b), we consider evaluation of learned representations under two settings: transductive (tran) and inductive (ind). The details about these two settings are in Appendix C.4.

**Setup.** For a fair comparison, we employ a standard GCN model (Kipf & Welling, 2017) as the GNN encoder for full-batch training on small graphs. For large-scale graphs, we use Graphsage (Hamilton et al., 2017a) with the sub-graph sampling strategy in a mini-batch manner. For ogbn paper datasets, we utilize all-roberta-large-v1 (Reimers & Gurevych, 2019; Duan et al., 2023) as the feature extractor. We conduct experiments using several random seeds and report both the average performance and standard deviation. We select the optimal hyperparameters solely based on accuracy on the validation set. In cases where publicly available and standardized data splits were used in the original paper, we adopt their reported results. For baselines that deviated from standardized data splits, we either reproduce the results using the authors' of cial code. The details of hyperparameter search are provided in Appendix C.5.

### 5.1. Main Results and Comparison on both Transductive and Inductive Settings

In this section, we evaluate the node representations from the MLP encoder learned by GraphECL.

**Transductive Setting.** We first consider the standard transductive setting in the task of node classification. We provide the results of other tasks in Table 8 in Appendix, respectively. Table 1 reports the average accuracy on both heterophilic

Table 3. Node classification results in a real-world scenario with both inductive and transductive nodes. tran denotes the accuracy on seen test transductive nodes and ind indicates the accuracy on unseen test inductive nodes.

Methods	Citeseer		Pubmed		Photo		Actor		Flickr		ogbn-arxiv	
	tran	ind	tran	ind	tran	ind	tran	ind	tran	ind	tran	ind
DGI	63.82 <sup>±1.69</sup>	66.25 <sup>±2.54</sup>	70.33 <sup>±2.61</sup>	70.48 <sup>±2.41</sup>	87.11 <sup>±1.65</sup>	88.14 <sup>±0.45</sup>	28.07 <sup>±2.19</sup>	28.08 <sup>±1.96</sup>	37.84 <sup>±0.22</sup>	39.71 <sup>±0.30</sup>	65.21 <sup>±0.35</sup>	63.91 <sup>±0.37</sup>
GCA	66.33 <sup>±1.16</sup>	69.02 <sup>±2.08</sup>	81.16 <sup>±0.80</sup>	81.52 <sup>±0.56</sup>	90.54 <sup>±0.54</sup>	90.59 <sup>±0.51</sup>	27.94 <sup>±1.62</sup>	27.72 <sup>±1.51</sup>	41.25 <sup>±0.33</sup>	42.95 <sup>±0.18</sup>	67.15 <sup>±0.29</sup>	66.95 <sup>±0.25</sup>
BGRL	67.04 <sup>±1.44</sup>	67.62 <sup>±1.24</sup>	78.36 <sup>±0.41</sup>	79.55 <sup>±0.40</sup>	87.95 <sup>±0.68</sup>	88.30 <sup>±0.45</sup>	29.04 <sup>±1.06</sup>	29.07 <sup>±0.65</sup>	40.78 <sup>±0.20</sup>	41.75 <sup>±0.15</sup>	68.57 <sup>±0.31</sup>	67.11 <sup>±0.29</sup>
SUGRL	69.16 <sup>±0.63</sup>	71.24 <sup>±1.06</sup>	81.07 <sup>±0.76</sup>	80.52 <sup>±1.21</sup>	89.88 <sup>±0.64</sup>	89.11 <sup>±0.24</sup>	28.95 <sup>±1.37</sup>	28.68 <sup>±1.18</sup>	40.37 <sup>±0.20</sup>	41.33 <sup>±0.25</sup>	69.96 <sup>±0.39</sup>	68.21 <sup>±0.37</sup>
CCA-SSG	68.84 <sup>±1.05</sup>	70.05 <sup>±2.70</sup>	79.76 <sup>±2.32</sup>	80.34 <sup>±2.32</sup>	88.60 <sup>±1.95</sup>	88.77 <sup>±1.85</sup>	28.52 <sup>±1.11</sup>	28.06 <sup>±2.69</sup>	42.16 <sup>±0.25</sup>	43.22 <sup>±0.27</sup>	68.34 <sup>±0.17</sup>	67.72 <sup>±0.29</sup>
GraphECL	69.96 <sup>±0.10</sup>	72.87 <sup>±1.30</sup>	81.71 <sup>±0.91</sup>	82.47 <sup>±1.00</sup>	92.18 <sup>±0.15</sup>	89.42 <sup>±0.03</sup>	36.18 <sup>±1.29</sup>	37.17 <sup>±1.84</sup>	45.43 <sup>±0.14</sup>	43.50 <sup>±0.20</sup>	70.58 <sup>±0.23</sup>	70.12 <sup>±0.12</sup>

Table 4. Ablation studies on Flickr dataset.

Ablation	Accuracy (%)
A1 w/o inter-model negative pair	42.34 <sup>±0.03</sup>
A2 w/o intra-model negative pair	42.24 <sup>±0.01</sup>
A3 w/o both types of negative pairs	40.25 <sup>±0.05</sup>
A4 w/ only MLP encoder	44.83 <sup>±0.06</sup>
A1 & A4	42.34 <sup>±0.04</sup>
A2 & A4	42.32 <sup>±0.10</sup>
A3 & A4	41.28 <sup>±0.02</sup>
GraphECL	48.49 <sup>±0.15</sup>

and homophilic graphs. As shown in the table, across different datasets, GraphECL can learn representations that outperform other methods. These results are indeed remarkable, given that GraphECL exclusively employs the learned MLP representations for inference without any reliance on input graph structures. This demonstrates that the MLP learned by GraphECL is able to capture meaningful structural information that is beneficial and generalized to downstream tasks. Since the inference time is much more heavily weighted in the large-scale graphs, we compare GraphECL with baselines in large-scale graphs in Table 2.5.3. Ablation Studies and Further Model Analysis

From the table, we can observe that our efficient GraphECL can still achieve better performance on large-scale graphs. Inductive Setting. To gain a better understanding of GraphECL's effectiveness, we evaluate the representations in a realistic production scenario that encompasses both transductive and inductive settings. In inductive evaluation, we set aside certain test nodes (20%) from test nodes in the transductive setting to create an inductive set (see Section C.4). We adopt GraphSAGE (Hamilton et al., 2017a) as the encoder during training for all methods. As shown in Table 3, GraphECL still achieves superior or competitive performance compared to elaborate methods employing GNNs as the inference encoder. These results support the deployment of MLP learned by GraphECL as a significantly faster model, with minimal or no performance degradation.

## 5.2. Inference Time Comparison on Large-Scale Graphs

To demonstrate the inference efficiency of GraphECL, we compare the inference time on large-scale graphs with BGRL (GNN-L2W256), CCA-SSG (GNN-L2W128), and SUGRL (GNN-L1W128), where GNN-LiWj indicates the method achieving the best performance with  $i$  layers of GraphSAGE with dimension  $j$ s as shown in Table 1. The

Figure 5. Inference time comparison of different methods on large-scale graphs. Note that time axes are log-scaled.

The MLP learned by GraphECL has 128 hidden dimensions.

Our results in Table 1 and Figure 5 indicate that GraphECL achieves the highest accuracy while attaining significant speedups in inference. In large-scale graphs, the MLP learned by GraphECL is about 200x faster than CCA-SSG with the same number of layers, which demonstrates the superior inference efficiency of GraphECL.

### 2.5.3. Ablation Studies and Further Model Analysis

Ablation Studies. We study the effects of intra-model and inter-model negative losses. We consider three ablations:

- (A1) Removing the inter-model negative pairs;
- (A2) Removing the intra-model negative pairs; and
- (A3) Removing both intra-model and inter-model negative pairs. We also explore the effects of the cross-model contrastive architecture in GraphECL by removing the asymmetric GNNMLP architecture and consider other two ablations GraphECL: (A4) using only the MLP as the encoder. Table 4 lists the

results. We also find that the performance of GraphECL suffers in the absence of negative examples, which shows that the information provided by negative examples is crucial for good generalization. Additionally, we observe that GraphECL using only MLP as the encoder can not match the performance of SOTA methods, although it can achieve fast inference. These results collectively underscore the importance of each of the components of GraphECL's GNNMLP architecture in achieving SOTA performance on downstream tasks while achieving substantial inference speedup.

Effectiveness of Generalized Contrastive Loss. We maintain the cross-model contrastive architecture while replacing our generalized contrastive loss with the vanilla InfoNCE loss, as shown in Equation (3). Table 5 summarizes the

Figure 6. (Left) The training dynamics and inference performance on Cora. (Right two) The pairwise cosine similarity of representations for randomly sampled node pairs, one-hop neighbors, and two-hop neighbors on Cora and Actor. More results are in Appendix D.5.

Table 5. Ablation study on the effectiveness of using cross-model contrastive loss in GraphECL.

	Cora	Citeseer	Pubmed	Photo	Actor
GraphECL (InfoNCE)	74.55 $\pm$ 0.45	67.15 $\pm$ 0.25	76.50 $\pm$ 1.20	91.44 $\pm$ 0.61	33.39 $\pm$ 0.37
GraphECL (Generalized)	84.25 $\pm$ 0.05	73.15 $\pm$ 0.41	82.21 $\pm$ 0.05	94.22 $\pm$ 0.11	35.80 $\pm$ 0.89

results across all six datasets, demonstrating consistent improvements when using the generalized contrastive loss (Equation (4)) in our GraphECL formulation.

**Size of Negative Pairs** We investigate the influence of different number of negative pairs (i.e.,  $M$ ) in GraphECL (Appendix D.3). While a proper range can lead to certain gains (Figure 9), a small number of negative samples (e.g.,  $M = 5$ ) is enough to achieve good performance.

**Dimensionality of Representation.** We study the effects of dimensionality representation (i.e.,  $K$ ) in Appendix D.4. Not surprisingly, we find that larger dimensions often yield better results, with performance leveling off or decreasing when dimensionality becomes very large, for both homophilic and heterophilic graphs. This observation is consistent with Theorem 4.2, showing that a larger dimension can effectively reduce the upper bound of downstream errors.

**The Parameter  $\beta$ .** We explore the effects of the trade-off parameter of  $\beta$  and message passing layers during training. As shown in Figure 8, while a specific value can lead to certain gains, GraphECL is robust to different choices of the value  $\beta$  on different graphs and GraphECL is not very sensitive to  $\beta$  used in training.

**Training Dynamics.** We also investigate the training process of GraphECL. Figure 6 (left) shows the curves of training losses and downstream performance using GNN and MLP, respectively. We find that: (1) GraphECL exhibits training stability, consistently improving performance as training losses decrease; (2) As the training proceeds, GraphECL gradually and dynamically acquires knowledge from GNN, facilitating the dynamic exchange of information between cross-model GNN and MLP in GraphECL.

**Visualization.** In addition to quantitative analysis, we visualize pairwise cosine similarities among randomly sampled nodes, one-hop neighbors, and two-hop neighbor pairs

based on learned representations. Figure 6 shows that, in the homophilic graph (i.e., Cora), nodes exhibit representations that are similar to those of their neighbors. GraphECL enhances similarities between neighbor nodes compared to randomly sampled node pairs, demonstrating its ability to effectively preserve one-hop neighborhood contexts. In addition, in the heterophilic graph (i.e., Actor), GraphECL strives to bring two-hop neighbor nodes closer together. This observation is consistent with our analytical insights, showing that GraphECL is effective at automatically capturing regularities in graph structures beyond just homophily.

## 6. Conclusion

In this paper, we introduced GraphECL, a simple, novel, effective and inference-efficient GCL framework for learning effective node representations from graph data. GraphECL introduces a cross-model contrastive architecture and a generalized contrastive loss to train MLP encoder. GraphECL is faster, often by orders of magnitude, than GCL methods using the GNN encoder, while also achieving superior performance! We demonstrate theoretically that GraphECL leverages neighborhood distribution as an inductive bias. Extensive experiments on real-world small and large-scale graphs demonstrate its advantages over current methods, including the vastly superior inference efficiency and generalization on both homophilic and heterophilic graphs.

While our approach relies on node attributes, assuming access to these attributes is a modest assumption. In real-world applications, node attributes are often high-dimensional and rich in information. Therefore, a MLP optimized with our GraphECL can effectively distill structural information from the GNN even without any labels. Addressing the challenge of dealing with graphs lacking node attributes is a promising future direction, potentially involving an early layer or other pre-processing network embedding methods.



## Acknowledgements

This work was funded in part by grants from the National Science Foundation (2226025, 2041759), the National Center for Advancing Translational Sciences, and the National Institutes of Health (UL1 TR002014).

## Impact Statement

This paper advances contrastive learning on graphs. By improving the inference efficiency and accuracy of contrastive learning on graphs, this research has the potential to contribute to advancements in latency-constrained applications.

This work shares the societal implications of machine learning broadly, without raising any additional concerns.

## References

- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning* pp. 1597–1607, 2020.
- Chen, Z., Xiao, T., and Kuang, K. Ba-gnn: On learning bias-aware graph neural network. *2022 IEEE 38th International Conference on Data Engineering*, pp. 3012–3024, 2022.
- Chien, E., Chang, W.-C., Hsieh, C.-J., Yu, H.-F., Zhang, J., Milenkovic, O., and Dhillon, I. S. Node feature extraction by self-supervised multi-scale neighborhood prediction. In *International Conference on Learning Representations* 2021.
- Cui, P., Wang, X., Pei, J., and Zhu, W. A survey on network embedding. *IEEE transactions on knowledge and data engineering* pp. 833–852, 2018.
- Duan, K., Liu, Q., Chua, T.-S., Yan, S., Ooi, W. T., Xie, Q., and He, J. Simteg: A frustratingly simple approach improves textual graph learning. *arXiv preprint arXiv:2308.02565* 2023.
- Dwivedi, V. P., Ramprasad, L., Galkin, M., Parviz, A., Wolf, G., Luu, A. T., and Beaini, D. Long range graph benchmark. *Advances in Neural Information Processing Systems* 35:22326–22340, 2022.
- Eckart, C. and Young, G. The approximation of one matrix by another of lower rank. *Psychometrika* 1(3):211–218, 1936.
- Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 855–864, 2016.
- Guo, X., Wang, Y., Wei, Z., and Wang, Y. Architecture matters: Uncovering implicit mechanisms in graph contrastive learning. In *Thirty-seventh Conference on Neural Information Processing Systems* 2023.
- Guo, X., Wang, Y., Wei, Z., and Wang, Y. Architecture matters: Uncovering implicit mechanisms in graph contrastive learning. *Advances in Neural Information Processing Systems* 36, 2024.
- Hamilton, W. L., Ying, R., and Leskovec, J. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 2017a.
- Hamilton, W. L., Ying, R., and Leskovec, J. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584* 2017b.
- Han, X., Zhao, T., Liu, Y., Hu, X., and Shah, N. Mlpinit: Embarrassingly simple gnn training acceleration with mlp initialization. In *The Eleventh International Conference on Learning Representations* 2022.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *arXiv preprint arXiv:2106.04156* 2021.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* 2015.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, pp. 22118–22133, 2020.
- Hu, Y., You, H., Wang, Z., Wang, Z., Zhou, E., and Gao, Y. Graph-mlp: Node classification without message passing in graph. *arXiv preprint arXiv:2106.04051* 2021.
- Joshi, C. K., Liu, F., Xun, X., Lin, J., and Foo, C. S. On representation knowledge distillation for graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 2022.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* 2014.
- Kipf, T. N. and Welling, M. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* 2016.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations* 2017.
- Lee, N., Lee, J., and Park, C. Augmentation-free self-supervised learning on graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 7372–7380, 2022.

- Lim, D., Hohne, F., Li, X., Huang, S. L., Gupta, V., Bhalerao, O., and Lim, S. N. Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods. *Advances in neural information processing systems*, pp. 20887–20902, 2021.
- Liu, Q., Zhang, H., and Liu, Z. Simplification of graph convolutional networks: A matrix factorization-based perspective. *arXiv preprint arXiv:2007.09036*, 2020.
- Liu, Y., Zheng, Y., Zhang, D., Lee, V. C., and Pan, S. Beyond smoothing: Unsupervised graph representation learning with edge heterophily discriminating. *Proceedings of the AAAI conference on artificial intelligence* pp. 4516–4524, 2023.
- Ma, Y., Liu, X., Shah, N., and Tang, J. Is homophily a necessity for graph neural networks? *International Conference on Learning Representations*, 2021.
- McAuley, J., Targett, C., Shi, Q., and Van Den Hengel, A. Image-based recommendations on styles and substitutes. *In ACM SIGIR conference on research and development in information retrieval* pp. 43–52, 2015.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27(1):415–444, 2001.
- Meng, E. and Liu, Y. Graph contrastive learning with graph info-min. *In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management* pp. 4195–4199, 2023.
- Meng, Z., Liang, S., Fang, J., and Xiao, T. Semi-supervisedly co-embedding attributed networks. *Advances in neural information processing systems*, 2019.
- Mernyei, P. and Cangea, C. Wiki-cs: A wikipedia-based benchmark for graph neural networks. *arXiv preprint arXiv:2007.02901*, 2020.
- Mo, Y., Peng, L., Xu, J., Shi, X., and Zhu, X. Simple unsupervised graph representation learning. *AAAI Conference on Artificial Intelligence*, 2022.
- Pei, H., Wei, B., Chang, K. C.-C., Lei, Y., and Yang, B. Geom-gcn: Geometric graph convolutional networks. *In International Conference on Learning Representations* 2019.
- Reimers, N. and Gurevych, I. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- Sun, W., Li, J., Chen, L., Wu, B., Bian, Y., and Zheng, Z. Rethinking and simplifying bootstrapped graph latents. *ACM international conference on web search and data mining* 2024.
- Suresh, S., Li, P., Hao, C., and Neville, J. Adversarial graph augmentation to improve graph contrastive learning. *Advances in Neural Information Processing Systems*, pp. 15920–15933, 2021.
- Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., and Mei, Q. Line: Large-scale information network embedding. *In The Web Conference* pp. 1067–1077, 2015.
- Thakoor, S., Tallec, C., Azar, M. G., Azabou, M., Dyer, E. L., Munos, R., Velicković, P., and Valko, M. Large-scale representation learning on graphs via bootstrapping. *In International Conference on Learning Representations* 2021.
- Tian, Y., Zhang, C., Guo, Z., Zhang, X., and Chawla, N. Learning mlps on graphs: A unified view of effectiveness, robustness, and efficiency. *The Eleventh International Conference on Learning Representations*, 2022.
- Trivedi, P., Lubana, E. S., Heimann, M., Koutra, D., and Thiragarajan, J. Analyzing data-centric properties for graph contrastive learning. *Advances in Neural Information Processing Systems*, pp. 14030–14043, 2022.
- Velicković, P., Fedus, W., Hamilton, W. L., W. L. P., Bengio, Y., and Hjelm, R. D. Deep graph infomax. *International Conference on Learning Representations*, 2018.
- Wang, H., Zhang, J., Zhu, Q., and Huang, W. Augmentation-free graph contrastive learning with performance guarantee. *arXiv preprint arXiv:2204.04874*, 2022.
- Wu, L., Lin, H., Huang, Y., and Li, S. Z. Quantifying the knowledge in gnns for reliable distillation into mlps. *In International Conference on Machine Learning*, 2023.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* pp. 4–24, 2020.
- Xia, J., Wu, L., Wang, G., Chen, J., and Li, S. Z. Progl: Rethinking hard negative mining in graph contrastive learning. *In International Conference on Machine Learning*, pp. 24332–24346, 2022.
- Xiao, T. and Wang, D. A general of ine reinforcement learning framework for interactive recommendation. *In Proceedings of the AAAI Conference on Artificial Intelligence* pp. 4512–4520, 2021.
- Xiao, T., Chen, Z., Wang, D., and Wang, S. Learning how to propagate messages in graph neural networks. *In*

- Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining 2021.
- Xiao, T., Chen, Z., Guo, Z., Zhuang, Z., and Wang, S. Decoupled self-supervised learning for graph neural networks. *Advances in Neural Information Processing Systems* pp. 620–634, 2022.
- Xiao, T., Zhu, H., Chen, Z., and Wang, S. Simple and asymmetric graph contrastive learning without augmentations. *Advances in neural information processing systems* 2023.
- Xiao, T., Cui, C., Zhu, H., and Honavar, V. G. Molbind: Multimodal alignment of language, molecules, and proteins. *arXiv preprint arXiv:2403.08167* 2024.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? *International Conference on Learning Representations* 2019.
- Yan, B., Wang, C., Guo, G., and Lou, Y. Tinygnn: Learning efficient graph neural networks. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* pp. 1848–1856, 2020.
- Yang, C., Liu, J., and Shi, C. Extract the knowledge of graph neural networks and go beyond it: An effective knowledge distillation framework. *Proceedings of the web conference 2021* pp. 1227–1237, 2021a.
- Yang, C., Wu, Q., Wang, J., and Yan, J. Graph neural networks are inherently good generalizers: Insights by bridging gnns and mlps. *The Eleventh International Conference on Learning Representations* 2022a.
- Yang, C., Wu, Q., and Yan, J. Geometric knowledge distillation: Topology compression for graph neural networks. *Advances in Neural Information Processing Systems* pp. 29761–29775, 2022b.
- Yang, H., Ma, K., and Cheng, J. Rethinking graph regularization for graph neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence* pp. 4573–4581, 2021b.
- Yang, L., Tian, Y., Xu, M., Liu, Z., Hong, S., Qu, W., Zhang, W., Cui, B., Zhang, M., and Leskovec, J. Vqgraph: Graph vector-quantization for bridging gnns and mlps. *arXiv preprint arXiv:2308.02117* 2023.
- Yang, Y., Qiu, J., Song, M., Tao, D., and Wang, X. Distilling knowledge from graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 7074–7083, 2020.
- Yang, Z., Cohen, W., and Salakhudinov, R. Revisiting semi-supervised learning with graph embeddings. *International conference on machine learning* pp. 40–48. PMLR, 2016.
- You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., and Shen, Y. Graph contrastive learning with augmentation. *Advances in neural information processing systems* pp. 5812–5823, 2020.
- You, Y., Chen, T., Wang, Z., and Shen, Y. Bringing your own view: Graph contrastive learning without prefabricated data augmentations. *ACM international conference on web search and data mining* pp. 1300–1309, 2022.
- Zeng, H., Zhou, H., Srivastava, A., Kannan, R., and Prasanna, V. K. Graphsaint: Graph Sampling Based Inductive Learning Method. *International Conference on Learning Representations* 2020.
- Zhang, H., Wu, Q., Yan, J., Wipf, D., and Yu, P. S. From canonical correlation analysis to self-supervised graph neural networks. *Advances in Neural Information Processing Systems* 2021a.
- Zhang, S., Liu, Y., Sun, Y., and Shah, N. Graph-less neural networks: Teaching old mlps new tricks via distillation. In *International Conference on Learning Representations* 2021b.
- Zhang, Y., Zhu, H., Song, Z., Koniusz, P., and King, I. Costa: covariance-preserving feature augmentation for graph contrastive learning. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* pp. 2524–2534, 2022.
- Zheng, W., Huang, E. W., Rao, N., Katariya, S., Wang, Z., and Subbian, K. Cold brew: Distilling graph node representations with incomplete or missing neighborhoods. In *International Conference on Learning Representations* 2021.
- Zheng, X., Liu, Y., Pan, S., Zhang, M., Jin, D., and Yu, P. S. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082* 2022a.
- Zheng, Y., Pan, S., Lee, V., Zheng, Y., and Yu, P. S. Rethinking and scaling up graph contrastive learning: An extremely efficient approach with group discrimination. *Advances in Neural Information Processing Systems* pp. 10809–10820, 2022b.
- Zhu, H., Xiao, T., and Honavar, V. G. 3m-diffusion: Latent multi-modal diffusion for text-guided generation of molecular graphs. *arXiv preprint arXiv:2403.07179* 2024.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Deep graph contrastive representation learning. *arXiv preprint arXiv:2006.04131* 2020.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Graph contrastive learning with adaptive augmentation. *The Web Conference* pp. 2069–2080, 2021.

Table 6. Comparison to previous contrastive schemes on graphs. Our approach does not rely on the invariant assumption that the augmentation can preserve the semantic nature of samples or the homophily assumption that connected nodes have similar representations.

Contrastive Schemes	Invariant Assumption	Homophily Assumption	Task Effective	Inference Efficient
GCL with Augmentations	3	7	3	7
Graph-MLP	7	3	7	3
MLP-Augmented GCL	3	7	7	3
Our GraphECL	7	7	3	3

## A. Additional Related Work

In this section, we provide a detailed comparison of characteristics with previous contrastive schemes on graphs in Table 6.

Comparisons with GCL methods with augmentations (Trivedi et al., 2022; Veličković et al., 2018; Zhu et al., 2021; You et al., 2020; 2022; Suresh et al., 2021; Xia et al., 2022; Meng & Liu, 2023; Guo et al., 2023): This contrastive learning scheme relies on graph augmentations and is built based on the invariant assumption that the augmentation can preserve the semantic nature of samples, i.e., the augmented samples have invariant semantic labels with the original ones. This scheme requires a GNN as the encoder to achieve good performance, which is computation-intensive during inference. However, our GraphECL is not based on graph augmentations but directly captures the 1-hop neighborhood distribution.

Comparisons with Graph-MLP (Yang et al., 2016; Hu et al., 2021; Yang et al., 2021b; Liu et al., 2020): Despite its inference efficiency due to the exclusive use of MLP, this approach exhibits significantly lower performance compared to GCL, as depicted in Figure 1. Moreover, this scheme over-emphasizes homophily (You et al., 2020; Xiao et al., 2022), making it difficult to generalize to graphs with heterophily (Liu et al., 2023; Xiao et al., 2022; Zheng et al., 2022a). In contrast, GraphECL enjoys good downstream performance with fast inference speed for both homophilic and heterophilic graphs.

Comparisons with MLP-augmented GCL: This method mentioned in Section 4.1 also relies on the invariant assumption that augmentation can preserve the semantic nature of samples, i.e., the augmented samples have invariant semantic labels with the original ones. In addition, it suffers from significant performance degradation on downstream tasks when using an MLP encoder for inference as shown in the results in Section 4.1.

## B. Proofs

### B.1. Proofs of Theorem 4.1

Theorem 4.1. The learning dynamics of the MLP encoder  $f_M$  with the generalized contrastive loss  $\beta = 1$  in Equation (4) saturates when the true normalized adjacency and the estimated normalized adjacency matrices  $D^{-1}A$  and  $\hat{D}^{-1}\hat{A}$ , which implies that, for  $\forall u, v \in V$ , we have:

$$P_n(u, j, v) = P_f(u, j, v), \quad P = \frac{\exp(f_M(v) \cdot f_G(u))}{\sum_{v' \in 2V} \exp(f_M(v') \cdot f_G(v'))}; \quad (7)$$

where  $P_n(u, j, v)$  is the 1-hop neighborhood distribution (i.e., the  $j$ -th row of the normalized adjacency matrix) and  $P_f(u, j, v)$  is the estimated neighborhood distribution.

Proof. We first show that minimizing GraphECL objective with  $\beta = 1$  is approximately to minimizing the losses of the positive and negative pairs of MLP representations.

$$\begin{aligned} L_{ECL} &= \frac{1}{|V|} \sum_{v \in 2V} \frac{1}{|N(v)|} \sum_{u \in 2N(v)} \log P \frac{\exp(f_M(v) \cdot f_G(u))}{\sum_{v' \in 2V} \exp(f_G(u) \cdot f_G(v')) + \exp(f_M(v) \cdot f_G(v'))}; \\ &= \frac{1}{|V|} \sum_{v \in 2V} \left( \frac{1}{|N(v)|} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) + \log \frac{\exp(f_M(v) \cdot f_G(v))}{\sum_{v' \in 2V} \exp(f_M(v') \cdot f_G(v'))} \right) \\ &= \underbrace{\frac{1}{|V|} \sum_{v \in 2V} \frac{1}{|N(v)|} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u)}_{L_{pos}} + \underbrace{\frac{1}{|V|} \sum_{v \in 2V} \log \frac{\exp(f_M(v) \cdot f_G(v))}{\sum_{v' \in 2V} \exp(f_M(v') \cdot f_G(v'))}}_{L_{neg}}; \quad (8) \end{aligned}$$

Then, we consider the unfolded iterations of descent steps of MLP representation  $f_M(v)$ . Specifically, we first consider

taking the derivatives  $\frac{\partial L_{\text{pos}}}{\partial \mathbf{f}_M(v)}$  and  $\frac{\partial L_{\text{neg}}}{\partial \mathbf{f}_M(v)}$  on  $\mathbf{f}_M(v)$ :

$$\frac{\partial L_{\text{pos}}}{\partial \mathbf{f}_M(v)} = \frac{1}{j|v|} \frac{1}{j|N(v)|} \sum_{u \in N(v)} \mathbf{f}_G(u); \quad \frac{\partial L_{\text{neg}}}{\partial \mathbf{f}_M(v)} = \frac{1}{j|v|} \sum_{v_2 \in V} \frac{\exp(\mathbf{f}_M(v)^\top \mathbf{f}_G(v_2)) \mathbf{f}_G(v_2)}{\sum_{v_2 \in V} \exp(\mathbf{f}_M(v)^\top \mathbf{f}_G(v_2))}. \quad (9)$$

As we denote representation matrix  $\mathbf{M}$  with  $\mathbf{f}_M(v)$  as the  $v$ -th row, the gradients in the Equation (9) can be written as the following matrix forms for simplicity and clarity:

$$\frac{\partial L_{\text{pos}}}{\partial \mathbf{M}} = \frac{1}{j|v|} \mathbf{D}^{-1} \mathbf{A} \mathbf{G}; \quad \frac{\partial L_{\text{neg}}}{\partial \mathbf{M}} = \frac{1}{j|v|} \mathbf{D}^{-1} \mathbf{A} \mathbf{G}; \quad (10)$$

where  $\mathbf{A} = \exp(\mathbf{M} \mathbf{G}^\top)$  is the affinity matrix based on feature similarity,  $\mathbf{D} = \text{deg}(\mathbf{A})$  is the diagonal matrix, whose element in the  $v$ -th row and  $v$ -th column is the sum of the  $v$ -th row of  $\mathbf{A}$ . In order to reduce the losses on positive pairs and negative pairs, we take a step by performing a gradient descent, which is to update representations  $\mathbf{M}$  as follows:

$$\mathbf{M}^{(t+1)} = \mathbf{M}^{(t)} - \frac{\partial (L_{\text{pos}} + L_{\text{neg}})}{\partial \mathbf{M}} = \mathbf{M}^{(t)} - (\mathbf{D}^{-1} \mathbf{A} - \mathbf{D}^{-1} \mathbf{A}) \mathbf{G}; \quad (11)$$

where  $\mathbf{M}^{(t+1)}$  and  $\mathbf{M}^{(t)}$  denote the representations before and after the update, respectively, and  $\eta$  is the step size of the gradient descent. Note that the constant  $\frac{1}{j|v|}$  has been absorbed in  $\eta$ . We can easily notice the updating in Equation (11) reveals the global minimum of the learning:  $\mathbf{D}^{-1} \mathbf{A} = \mathbf{D}^{-1} \mathbf{A}$ . Combining this with Equation (8) completes the proof.  $\square$

## B.2. Proofs of Theorem 4.2

To prove Theorem 4.2, we first present the following lemma:

**Lemma B.1.** (Theorem B.3 (page 32) in (HaoChen et al., 2021)). Let  $\mathbf{f}$  be a minimizer of the spectral contrastive loss:  $L_{\text{SCL}} = \sum_{x, x^0 \in \mathcal{X}} w_{xx^0} (f(x)^\top f(x^0) + w_x w_{x^0} (f(x)^\top f(x^0))^2)$ , where  $w_{x, x^0} = w(x)w(x^0)$  is the probability of a random positive pair being  $(x, x^0)$  while  $w_x$  is the probability of a randomly selected data point being  $x$ , we have:

$$\mathbb{E}(f), \min_W \sum_{x \in \mathcal{X}} w_x \mathbb{1}[g_{k+1}(w(x)) \neq y(x)] \leq \frac{\gamma}{k+1}; \quad (12)$$

where  $\gamma = \sum_{x, x^0 \in \mathcal{X}} w_{xx^0} \mathbb{1}[y(x) \neq y(x^0)]$  and  $W$  is the downstream linear classifier.  $\lambda_{k+1}$  is the  $(k+1)$ -th smallest eigenvalues of the matrix:  $\mathbf{D}^{-1/2} \mathbf{P} \mathbf{D}^{-1/2}$  where  $\mathbf{P} \in \mathbb{R}^{N \times N}$  is a symmetric probability matrix with  $\mathbf{P}_{xx^0} = w_{xx^0}$  and  $\mathbf{D}_P \in \mathbb{R}^{N \times N}$  is a diagonal matrix with  $(\mathbf{D}_P)_{xx} = w_x$ .

We also introduce the following Lemma in (HaoChen et al., 2021) which asserts that multiplying the embedding matrix on the right by an invertible matrix does not affect the linear probing error.

**Lemma B.2.** (Lemma 3.1 (page 8) in (HaoChen et al., 2021)). Consider an embedding matrix  $\mathbf{F} \in \mathbb{R}^{N \times k}$  and a linear classifier  $\mathbf{B} \in \mathbb{R}^{k \times r}$ . Let  $\mathbf{D} \in \mathbb{R}^{N \times N}$  be a diagonal matrix with positive diagonal entries and  $\mathbf{Q} \in \mathbb{R}^{k \times k}$  be an invertible matrix. Then, for any matrix  $\mathbf{F}' = \mathbf{D}^{-1} \mathbf{F} \mathbf{Q}$ , the linear classifier  $\mathbf{B}' = \mathbf{Q}^{-1} \mathbf{B}$  on  $\mathbf{F}'$  has the same prediction as  $\mathbf{B}$  on  $\mathbf{F}$ . Thus, we have  $\mathbb{E}(F) = \mathbb{E}(F')$ .

Intuitively, this Lemma suggests that although there might not be a single unique optimal solution, when we employ the representation within the context of linear probing, the linear classifier can efficiently handle variations caused by affine transformations. Thus, it produces identical classification errors across different variants when operating in optimal settings.

**Theorem 4.2.** Let  $\mathbf{f}_M$  be the global minimum of generalized contrastive loss (1) in Equation (4) and  $y(v)$  denote the label of  $v$ .  $\lambda_1, \dots, \lambda_N$  are the eigenvalues with descending order of the normalized adjacency matrix. Then, the linear probing error  $\text{off}_M$  is upper-bounded by:

$$\mathbb{E}(f_M), \min_W \frac{1}{j|v|} \sum_{v \in V} \mathbb{1}[g_{k+1}(w(v)) \neq y(v)] \leq \frac{1}{1 - \lambda_{k+1}}; \quad (13)$$

where  $\lambda_k = \frac{1}{j|v|} \sum_{v \in V} \frac{1}{j|N(v)|} \sum_{u \in N(v)} \mathbb{1}[y(v) = y(u)]$  and  $K$  is the dimension of the representation.

Proof. Based on Equation (8), we further have the following:

$$\begin{aligned}
 L_{ECL} &= \frac{1}{jVj} \sum_{v \in V} \left( \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) \right) + \log \sum_{v \in 2V} \exp(f_M(v) \cdot f_G(v)) : \\
 &= \frac{1}{jVj} \sum_{v \in 2V} \left( \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) \right) + \log \sum_{v \in 2V} \frac{\exp(f_M(v) \cdot f_G(v))}{jVj} \\
 &= \frac{1}{jVj} \sum_{v \in 2V} \left( \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) \right) + \log \sum_{v \in 2V} \frac{\exp(f_M(v) \cdot f_G(v))}{jVj} + \log jVj \\
 &\stackrel{c}{=} \frac{1}{jVj} \sum_{v \in 2V} \left( \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) \right) + \log \sum_{v \in 2V} \frac{\exp(f_M(v) \cdot f_G(v))}{jVj} \\
 &= \frac{1}{jVj} \sum_{v \in 2V} \left( \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) \right) + \sum_{v \in 2V} \frac{1}{jVj} (f_M(v) \cdot f_G(v)) \tag{14} \\
 &= \frac{1}{jVj} \sum_{v \in 2V} \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) + \frac{1}{jVj} \frac{1}{jVj} \sum_{v \in 2V} \sum_{v \in 2V} f_M(v) \cdot f_G(v) \\
 &= \frac{1}{jVj} \sum_{v \in 2V} \frac{1}{jN(v)j} \sum_{u \in 2N(v)} f_M(v) \cdot f_G(u) + \frac{1}{jVj} \frac{1}{jVj} \sum_{v \in 2V} \sum_{v \in 2V} (f_M(v) \cdot f_G(v))^2 \tag{15} \\
 &\stackrel{c}{=} \frac{1}{jVj} \sum_{v \in 2V} \frac{1}{jN(v)j} \sum_{u \in 2N(v)} 2f_M(v) \cdot f_G(u) + \frac{1}{jVj} \frac{1}{jVj} \sum_{v \in 2V} \sum_{v \in 2V} (f_M(v) \cdot f_G(v))^2, L_{cross} ; \tag{16}
 \end{aligned}$$

where the symbol  $\stackrel{c}{=}$  indicates equality up to a multiplicative and/or additive constant. Here, we utilize Jensen's inequality in (14). Inequality (15) holds because  $f_M(v)$  and  $f_G(u)$  are  $\ell_2$  normalized, and we assume the embedding heads consisting of last-layer ReLU neural networks. We define the two metrics  $M$  and  $G$ . Equation (16) holds if we set the temperature of positive pairs is twice to it of negative pairs. Here  $(M)_v = jVj^{1=2} f_M(v)$  and  $(G)_u = jVj^{1=2} f_G(u)$ . Then, the loss in Equation (16) is equivalent to the low-rank asymmetric matrix factorization loss up to a constant:

$$L_{AMF} = kD^{-1}A - M \cdot G^T k = L_{cross} + \text{const} \tag{17}$$

According to Eckart–Young–Mirsky theorem (Eckart & Young, 1936), the optimal solution  $M$  and  $G$  of  $L_{AMF}$  can be respectively represented as follows:

$$M \cdot (G)^T = U^K \text{diag}(\lambda_1; \dots; \lambda_K) (V^K)^T \tag{18}$$

where we denote  $D^{-1}AU^T$  as the spectral decomposition of  $D^{-1}A$ .  $(\lambda_1; \dots; \lambda_K)$  are the  $K$ -largest eigenvalue of  $D^{-1}A$ . The  $k$ -th column of  $U^K \in \mathbb{R}^{jVj \times K}$  is the corresponding eigenvector of the  $k$ -th largest eigenvalue and  $V^K \in \mathbb{R}^{jVj \times K}$  is a unitary matrix. Then the optimal solution  $M$  and  $G$  can be represented as follows:

$$M = U^K B R; \quad G = V^K \text{diag}(\lambda_1; \dots; \lambda_K) B^{-1} R; \tag{19}$$

where  $R \in \mathbb{R}^{K \times K}$  is a unitary matrix and  $B$  is an invertible diagonal matrix. Since  $(M)_v = jVj^{1=2} f_M(v)$  and  $(G)_u = jVj^{1=2} f_G(u)$ , we have:

$$f_M(v) = jVj^{1=2} ((U^K)_v B R)^T; \quad f_G(u) = jVj^{1=2} ((V^K)_u \text{diag}(\lambda_1; \dots; \lambda_K) B^{-1} R)^T : \tag{20}$$

Similar, if we consider optimizing following uni-model spectral contrastive loss:

$$L_{uni} = \frac{1}{jVj} \sum_{v \in V} \frac{1}{jN(v)j} \sum_{u \in 2N(v)} 2f_M(v) \cdot f_M(u) + \frac{1}{jVj} \frac{1}{jVj} \sum_{v \in 2V} \sum_{v \in 2V} (f_M(v) \cdot f_M(v))^2 : \tag{21}$$

The optimal solution  $\hat{f}_M$  of this uni-model spectral contrastive loss can be represented as follow:

$$\hat{f}_M(v) = jVj^{1=2} ((U_{uni}^K)_v B_{uni} R_{uni})^T : \tag{22}$$

Since the uni-model spectral contrastive loss in Equation (21) also decomposes the matrix  $\mathbf{D}^{-1}\mathbf{A}$ , the  $\mathbf{U}_{uni}^K = \mathbf{U}^K$ . As  $\mathbf{B}$ ,  $\mathbf{R}$ ,  $\mathbf{B}_{uni}$ ,  $\mathbf{R}_{uni}$  are invertible matrices, and the product of the invertible matrices is still invertible, we have the following:

$$f_M(v) = \hat{f}_M(v)\mathbf{T}, \quad (23)$$

where  $\mathbf{T} = (\mathbf{B}_{uni})^{-1}(\mathbf{R}_{uni})^{-1}\mathbf{B}\mathbf{R}$ . With Lemma B.2, we establish that  $E(f_M) = E(\hat{f}_M)$ . Additionally, we observe that the loss in Equation (21) shares the same form as the spectral contrastive loss when we define  $\frac{1}{jV_j}\mathbf{D}^{-1}\mathbf{A} = \hat{\mathbf{A}}$  i.e.,  $w_x = \frac{1}{jV_j}$  and  $w_{x^o_jx} = (\mathbf{D}^{-1}\mathbf{A})_{x;x^o}$ . It's worth noting that  $\mathbf{D}^{-1}\mathbf{A} = \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$  forms a symmetric matrix due to our random sampling process, which ensures that the same neighbors are sampled for each central node, approximately resulting in equal node degrees. Thus, with Lemma B.1, we can obtain the following:

$$E(f_M) = E(\hat{f}_M), \quad \min_w \frac{1}{jV_j} \times_{v \in V} [g_f, w(v) \notin y(v)] \quad \frac{1}{1 - \sigma_{K+1}} \quad (24)$$

where  $\frac{1}{jV_j} = \frac{1}{jV_j} \prod_{v \in V} \frac{1}{jN(v)}$  and  $\sigma_{K+1}$  is the  $(K+1)$ -th largest singular value of the normalized adjacency matrix  $\mathbf{D}^{-1}\mathbf{A}$ . Given the above, the proof is finished.  $\square$

## C. Experimental Details

Table 7. Statistics of Datasets.

	Cora	Citeseer	Pubmed	Photo	Flickr	WikiCS	Actor	Wisconsin	Cornell	Texas	snap-patents	ogbn-arxiv	ogbn-papers100M
#Nodes	2,708	3,327	19,717	7,650	89,250	11,701	7,600	251	183	183	2,923,922	169,343	111,059,956
#Edges	5,278	4,552	44,324	119,081	899,756	216,123	33,544	466	295	309	13,975,788	1,166,243	1,615,685,872
#Classes	7	6	3	8	7	10	5	5	5	5	5	40	172
$H(G)$	0.83	0.71	0.79	0.85	0.32	0.65	0.22	0.16	0.11	0.06	0.22	0.66	0.54
$S(G)$	0.89	0.81	0.87	0.91	0.33	0.75	0.68	0.42	0.40	0.79	0.29	0.79	0.71

### C.1. One-hop Node Homophily Level

We use the node homophily ratio to measure the one-hop neighbor homophily of the graph (Pei et al., 2019). Specifically, the node homophily ratio  $H(G)$  can be computed as follows:

$$H(G) = \frac{1}{jV_j} \times_{v \in V} \frac{1}{N(v)} \times_{u \in N(v)} (y(v) = y(u)) \quad (25)$$

### C.2. One-hop Neighborhood Context Similarity

To validate the assumption that nodes belonging to an identical semantic category are likely to exhibit similar patterns in their one-hop neighborhoods, even in heterophilic graphs, we examine whether nodes with the same label demonstrate similar distributions of labels in their neighborhoods regardless of homophily. We evaluate this characteristic by computing the class neighborhood similarity (Ma et al., 2021), which is defined as:

$$s(m; m^o) = \frac{1}{jV_m jV_{m^o}} \times_{u \in V_m; v \in V_{m^o}} \cos(d(u); d(v)); \quad (26)$$

where  $M$  denotes the total number of classes,  $V_m$  represents the set of nodes classified as  $m$ , and  $d(u)$  is the empirical histogram of the labels of node  $u$ 's neighbors across  $M$  classes. The cosine similarity function is represented by  $\cos(\cdot)$ . This metric for cross-class neighborhood similarity quantifies the differences in neighborhood distributions between varying classes. When  $m = m^o$ ;  $s(m; m^o)$  determines the intra-class similarity. To quantify the neighborhood similarity, we take the average of the intra-class similarities across all classes:

$$S(G) = \frac{1}{M} \sum_{m=1}^M s(m; m) \quad (27)$$

If nodes with identical labels exhibit similar neighborhood distributions, then the class neighborhood similarity  $S(G)$  will be high. Table 7 shows that heterophilic graphs exhibit stronger neighborhood similarity, even when the homophily ratio is low.

### C.3. Datasets Details

The statistics of the benchmark datasets, including homophily levels and 1-hop neighborhood similarities, are given in Table 7. All datasets and public splits can be found in PyTorch Geometric: <https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html>.

**Cora, Citeseer, and Pubmed.** (Yang et al., 2016) These datasets serve as some of the most prevalent benchmarks for node classification. Each one constitutes a graph representing citations, with nodes symbolizing documents and edges depicting citation relationships between them. The classification of each node is determined by the respective research field. Features of the nodes are derived from a bag-of-words model applied to their abstracts. We utilize the public split: a fixed 20 nodes from each class for training and another distinct set of 500 and 1,000 nodes for validation and testing, respectively.

**WikiCS.** (Mernyei & Cangea, 2020) This graph consists of nodes corresponding to Computer Science articles, with edges based on hyperlinks and 10 classes representing different branches of the field. We adopt a 10/10/80% training/validation/testing public split provided by PyTorch Geometric.

**Photo.** (McAuley et al., 2015) This graph originates from the Amazon co-purchase graph (McAuley et al., 2015), where nodes denote products and edges connect pairs of items often bought together. In the Photo dataset, products are categorized into eight classes based on their category, and the node features are represented by a bag-of-words model of the product’s reviews. We employ a public split of the nodes into training, validation, and testing sets, following a 10/10/80% ratio as described in (Thakoor et al., 2021).

**Flickr.** (Zeng et al., 2020) In this graph, each node symbolizes an individual image uploaded to Flickr. An edge is established between the nodes of two images if they share certain attributes, such as geographic location, gallery, or user comments. The node features are represented by a 500-dimensional bag-of-words model provided by NUS-WIDE. Regarding labels, we examined the 81 tags assigned to each image and manually consolidated them into 7 distinct classes, with each image falling into one of these categories. We use a random node division method, adhering to a 50/25/25% split for training, validation, and testing sets, following (Zeng et al., 2020).

**Cornell, Wisconsin and Texas.** (Pei et al., 2019) These are networks of webpages gathered from the computer science departments of various universities by Carnegie Mellon University. In each network, the nodes represent individual webpages, while the edges signify hyperlinks between them. The features of the nodes are depicted using bag-of-words representations of the webpages. The objective is to categorize each node into one of five classes.

**Actor.** (Pei et al., 2019) This is a subgraph induced solely by actors, derived from the broader film-director-actor-writer network. In this subgraph, nodes represent actors, while edges denote the co-occurrence of two nodes on the same Wikipedia page. The features of the nodes are constituted by keywords found on Wikipedia pages. Labels are categorized into five groups based on the content of the actor’s corresponding Wikipedia page.

For **Texas, Wisconsin, Cornell, and Actor**, we use the raw data provided by Geom-GCN (Pei et al., 2019) with the standard fixed 10-fold split for our experiment. In addition to the above graphs, we also conduct experiments on the following three large-scale graphs: snap-patents, ogbn-arxiv and ogbn-papers100M,

**Snap-patents.** (Lim et al., 2021) The Snap-patents dataset encompasses a collection of utility patents from the United States, where each node represents a patent, and edges are formed between patents that cite one another. The features of the nodes are extracted from the metadata of the patents. In this work, we introduce a task aiming to predict the time at which a patent was granted, which is categorized into five classes. We used the unprocessed data from (Lim et al., 2021), employing the standard 10-fold split for our experimental setup.

**Ogbn-arxiv and Ogbn-papers100M.** (Hu et al., 2020) These two large-scale datasets are collected by Hu et al. (2020). Ogbn-arxiv and Ogbn-papers100M are citation networks where each node represents a paper. The corresponding features consist of titles and abstracts, and node labels are the primary categories of the papers (Chien et al., 2021). We used the public split ratio provided in the OGB benchmark.

### C.4. Transductive and Inductive Settings for Unsupervised Representation Learning

**Transductive Setting.** To fully evaluate the model, we consider two settings: transductive (tran) and inductive (ind). In the transductive setting, our evaluation consists of two phases. Initially, we pre-train models on graph  $G$ , followed by the generation of representations for all nodes within the graph, denoted as  $z_v$  for  $v \in V$ . Subsequently, we employ a linear



Table 8. Graph classification results (%) on MUTAG and PROTEINS

Method	Graph-MLP	VGAE	CCA-SSG	BGRL	GraphECL
MUTAG	75.8±2.0	84.4±0.6	85.8±0.4	86.8±1.3	<b>88.5±1.2</b>
PROTEINS	71.1±1.5	74.0±0.5	73.1±0.6	73.5±0.7	<b>75.2±0.3</b>

classifier trained in fixed learned representations using labeled data  $Z^L$  and  $Y^L$ . Finally, we assess the remaining inferred representations  $Z^U$  with the corresponding labels  $Y^U$ .

**Inductive Setting.** In the unsupervised inductive setting, we randomly select 20% of the nodes as a test set for inductive evaluation. Specifically, we partition the unlabeled nodes  $V^U$  into two separate subsets: observed and inductive (i.e.,  $V^U = V_{obs}^U \sqcup V_{ind}^U$ ). This leads to the creation of three distinct graphs:  $G = G^L \sqcup G_{obs}^U \sqcup G_{ind}^U$ , where no nodes are shared between  $G^L \sqcup G_{obs}^U$  and  $G_{ind}^U$ . Importantly, during training, we remove the edges that connect  $G^L \sqcup G_{obs}^U$  and  $G_{ind}^U$ . Upon completing the self-supervised pretraining in  $G^L \sqcup G_{obs}^U$ , we generate representations for all nodes. Consequently, the learned representations and associated labels are partitioned into three separate sets:  $Z = Z^L \sqcup Z_{obs}^U \sqcup Z_{ind}^U$  and  $Y = Y^L \sqcup Y_{obs}^U \sqcup Y_{ind}^U$ . A downstream classifier is then trained on the learned representations  $Z^L$  and labels  $Y^L$ . Finally, we evaluate the remaining representations  $Z^U$  and  $Z_{ind}^U$  in the downstream classifier with labels  $Y_{obs}^U$  and  $Y_{ind}^U$ , respectively.

### C.5. Setup and Hyper-parameter Settings

We utilized the official implementations publicly released by the authors for the baselines. To ensure a fair comparison, we conducted a grid search to determine the optimal hyperparameters. Our experiments were conducted on a machine equipped with NVIDIA RTX A100 GPUs with 80GB memory. For all experiments, we employed the Adam optimizer (Kingma & Ba, 2014). A small-scale grid search was used to select the best hyperparameters for all methods. Specifically, for our approach, we explored the following hyperparameter ranges:  $\eta$  from  $\{0.001, 0.01, 0.1, 0.5, 1g\}$ ,  $K$  from  $\{256, 512, 1024, 2048, 4096g\}$ ,  $\beta$  from  $\{0.5, 0.75, 0.99, 1g\}$ , and the number of negative pairs  $M$  from  $\{1, 5, 10g\}$  when negative sampling was used. Furthermore, we tuned the learning rate from the set  $\{1e-3, 5e-3, 1e-4g\}$  and the weight decay from the set  $\{0, 1e-4, 3e-4, 1e-6g\}$ . The selection of the optimal hyperparameter configuration was based on the accuracy on the validation set.

## D. Additional Experimental Results

### D.1. Graph Classification Performance

For the graph classification task, we can use a non-parameterized graph pooling (readout) function, such as MeanPooling, to obtain the graph-level representation. In our experiments, we focus on graph classification using two benchmarks: PROTEINS and MUTAG. We follow the same experimental setup as for GraphCL (You et al., 2020). The results are presented in Table 8. From the table, we observe that our GraphECL performs well on the graph classification task and achieves better performance compared to the baselines. This observation, coupled with the node classification results, underscores the effectiveness of GraphECL in acquiring more expressive and resilient node representations for a variety of downstream tasks. These findings further validate that modeling one-hop neighborhood patterns confers advantages on downstream tasks on real-world graphs with varying degrees of homophily.

### D.2. Performance on Long Range Graph Benchmark

To further evaluate the effectiveness of GraphECL in capturing inter-neighborhood information, we also evaluate it on Long Range Graph Benchmark (Dwivedi et al., 2022). Specifically, we compare GraphECL with two graph contrastive learning methods, BGRL and CCA-SSG, on PascalVOC-SP (Dwivedi et al., 2022) in Table 9. For all methods, we employ the GCN as the backbone. From the table below, we can observe that GraphECL performs well in PascalVOC-SP, achieving better performance compared to the baselines. This further strengthens GraphECL on capturing inter-neighborhood information.

### D.3. The Effect of Size of Negative Pairs

We conducted a sweep over the size of negative samples, denoted as  $M$ , to study its impact on performance. We varied  $M$  across the values  $\{1;5;10\}$ . For each value  $M$ , we first learned node representations and subsequently applied these learned representations to node classification. The results of this experiment are shown in Figure 9. From the figure, we observe




Figure 7. The effect of the dimensions of representations.




Figure 8. The effect of the hyperparameter  $\lambda$ .




Figure 9. The effect of the size of negative pairs.

that even a small number of negative samples, such as  $M = 5$ , is sufficient to achieve good performance across all graphs, demonstrating that GraphECL is particularly robust to reduced negative pairs.

#### D.4. The Effect of Representation Dimension

We investigate the impact of different dimensions of representations. Figures 7 show the results of node classification with varying dimensions on homophilic and heterophilic graphs. From the figure, we can observe that larger dimensions often yield better results for both homophilic and heterophilic graphs. This observation is consistent with Theorem 4.2, which shows that a larger dimension can effectively reduce the upper bound of downstream errors. Training with extremely large dimensions for some graphs may lead to a slight drop of performance, as GraphECL may suffer from the over-fitting issue.

Table 9. The performance on the long-range graph benchmark PascalVOC-SP.

Method	BGRL	CCA-SSG	GraphECL
PascalVOC-SP	0.1356±0.0087	0.1437±0.0095	<b>0.1588±0.0091</b>

