

Rethinking Graph Masked Autoencoders through Alignment and Uniformity

Liang Wang^{1,2*}, Xiang Tao^{1,2*}, Qiang Liu^{1,2}, Shu Wu^{1,2†}, Liang Wang^{1,2}

¹Center for Research on Intelligent Perception and Computing
State Key Laboratory of Multimodal Artificial Intelligence Systems
Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences
{liang.wang, xiang.tao}@cripac.ia.ac.cn, {qiang.liu, shu.wu, wangliang}@nlpr.ia.ac.cn

Abstract

Self-supervised learning on graphs can be bifurcated into contrastive and generative methods. Contrastive methods, also known as graph contrastive learning (GCL), have dominated graph self-supervised learning in the past few years, but the recent advent of graph masked autoencoder (GraphMAE) rekindles the momentum behind generative methods. Despite the empirical success of GraphMAE, there is still a dearth of theoretical understanding regarding its efficacy. Moreover, while both generative and contrastive methods have been shown to be effective, their connections and differences have yet to be thoroughly investigated. Therefore, we theoretically build a bridge between GraphMAE and GCL, and prove that the node-level reconstruction objective in GraphMAE implicitly performs context-level GCL. Based on our theoretical analysis, we further identify the limitations of the GraphMAE from the perspectives of alignment and uniformity, which have been considered as two key properties of high-quality representations in GCL. We point out that GraphMAE’s alignment performance is restricted by the masking strategy, and the uniformity is not strictly guaranteed. To remedy the aforementioned limitations, we propose an Alignment-Uniformity enhanced Graph Masked AutoEncoder, named AUG-MAE. Specifically, we propose an easy-to-hard adversarial masking strategy to provide hard-to-align samples, which improves the alignment performance. Meanwhile, we introduce an explicit uniformity regularizer to ensure the uniformity of the learned representations. Experimental results on benchmark datasets demonstrate the superiority of our model over existing state-of-the-art methods. The code is available at: <https://github.com/AzureLeon1/AUG-MAE>.

1. Introduction

Graph self-supervised learning can be categorized into two distinct types, contrastive and generative methods (Wu et al. 2021; Xie et al. 2022; Liu et al. 2022). Motivated by the InfoMax principle, contrastive methods, also known as graph contrastive learning (GCL), maximize the mutual information between positive pairs. The contrastive loss is proved to asymptotically optimize two properties, representation

alignment and *uniformity*, which are considered to lead to high-quality representations (Wang and Isola 2020). On the other hand, the basic idea behind generative methods is to reconstruct the masked portions of data with generative models, such as autoencoders and autoregressive models. The reconstruction process reveals inherent data patterns and encode them into learned representations.

In the past few years, contrastive methods have dominated graph self-supervised learning due to their superior performance, and have gained sufficient theoretical analysis and understanding. Recently, graph masked autoencoder (GraphMAE) (Hou et al. 2022) is proposed and demonstrates that generative methods can also achieve competitive, and even better, performance when appropriately designed. GraphMAE analyzes the deficiencies of early generative methods in terms of reconstruction target, decoder structure, and optimization objective. This model addresses these deficiencies in a sophisticated manner and achieves performance beyond that of the contrastive methods. Many subsequent studies take GraphMAE as a foundation to further improve the model structure and apply it to different domains (Zhang et al. 2022b; Li et al. 2022; Tan et al. 2023; Xia et al. 2023; Ye, Xia, and Huang 2023). These studies spark renewed interest and reflection on generative methods.

However, despite the recent empirical success of GraphMAE, there is still a lack of sufficient understanding regarding its efficacy. Additionally, it remains unknown whether there exists a connection between GraphMAE and GCL. Specifically, the following questions arise: *Why is GraphMAE effective? Are GraphMAE and GCL completely different methods, or do they share any commonality?*

To answer these questions, we conducted a theoretical analysis of GraphMAE. To facilitate the understanding of the relationship between GraphMAE and GCL, we do not analyze GraphMAE independently, but build a bridge between GraphMAE and GCL. Specifically, we first view the learning process of GraphMAE as using the contexts (ego-graphs) of the masked nodes to restore the original features of these nodes. Then, we theoretically prove that the node-level reconstruction loss in GraphMAE is lower bounded by the context-level alignment loss. This indicates that GraphMAE has the ability to align positive pairs defined in contrastive learning, and *GraphMAE implicitly performs*

*These authors contributed equally.

†Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

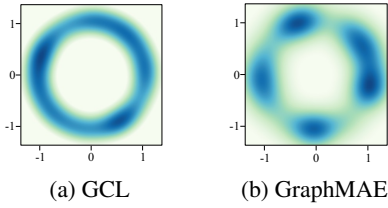


Figure 1: Distribution of nodes representations on the unit hypersphere learned by GCL (taking GRACE (Zhu et al. 2020) as an example) and GraphMAE (Hou et al. 2022). The representations learned by GCL is more uniformly distributed than GraphMAE.

context-level GCL in its learning process.

Since we have established the connection between GraphMAE and GCL through theoretical analysis, we are able to identify the limitations of GraphMAE from the perspective of representation alignment and uniformity: (1) For alignment, although GraphMAE is proven to have the ability to align representations of positive pairs, the practical alignment performance not only depends on the model’s ability, but also influenced by the masking strategy. Further, the random masking strategy adopted in GraphMAE ignores the difficulty of aligning positive samples. (2) For uniformity, the representation uniformity is not strictly guaranteed in GraphMAE. Specifically, GraphMAE can naturally avoid full feature collapse, i.e., the learned representations do not collapse to be the same. However, as shown in Fig. 1, we observe that the representations still suffer from partial dimensional collapse (Jing et al. 2022; Guo et al. 2023), i.e., the representations shrink along a certain dimension and are not uniformly distributed in the feature space. Therefore, the uniformity of the representations can also be improved.

To overcome these limitations, we propose an Alignment-Uniformity enhanced Graph Masked AutoEncoder, named AUG-MAE. Specifically, we propose an easy-to-hard adversarial masking strategy to provide richer hard-to-align samples, which improves the alignment performance. Moreover, we introduce an explicit uniformity regularizer to ensure the uniformity of the learned representations. Experimental results on benchmark datasets demonstrate the superiority of our model over existing state-of-the-art methods. Meanwhile, the learned representations are better aligned and more uniformly distributed in the feature space. The main contributions of our work are outlined as follows:

- We conduct a theoretical analysis of GraphMAE and demonstrate that it implicitly performs context-level GCL. Further, we identify the limitations of GraphMAE from the perspective of alignment and uniformity.
- We propose an AUG-MAE model. For alignment, we employ an easy-to-hard adversarial masking strategy to generate hard-to-align positive pairs. For uniformity, we introduce an explicit uniformity regularizer.
- We conduct extensive experiments on benchmark datasets, which show that AUG-MAE outperforms state-of-the-art methods on downstream tasks, and achieves better alignment and uniformity.

2. Related Work

In this section, we succinctly review existing studies for graph self-supervised learning and two measurements of representation learning, i.e., alignment and uniformity.

Graph Self-Supervised Learning. Graph self-supervised learning has been proposed as a promising paradigm for learning graph representations without labels. Existing methods can be categorized into contrastive and generative.

Contrastive methods learn meaningful representations by maximizing the mutual information between augmented views (Zhang et al. 2022a; Zhu et al. 2022). Some early efforts focus on contrastive modes (Zhu et al. 2021a,b; Qiu et al. 2020; You et al. 2020), and several recent studies focus on the negative-sample-free technique (Thakoor et al. 2022). Despite the progress made by these studies, GCL still relies on intricate designs.

Although early generative methods lagged behind contrastive methods, the recently proposed GraphMAE (Hou et al. 2022) has greatly improved the empirical performance of generative methods through sophisticated designs and triggered many subsequent studies, such as WGDN (Cheng et al. 2023), SeeGera (Li et al. 2023), and HGMAE (Tian et al. 2023). However, the theoretical support of GraphMAE is still not thoroughly investigated.

Alignment and Uniformity. Several studies analyze how the contrastive objective influences the representation learning. Wang and Isola (2020) first identify two properties induced from contrastive learning objective: alignment of positive pairs and uniformity of the representation distribution on the hypersphere. Both alignment and uniformity play vital roles in enhancing the discriminative and generalization abilities of contrastive learning models, then these two properties are utilized to measure and improve the quality of learned representations (Xia et al. 2022; Yu et al. 2022; Pu, Zhao, and Zheng 2022; Zhang, Wang, and Wang 2022).

3. Preliminary

3.1. Problem Formulation

Let $\mathcal{G} = (\mathcal{V}, \mathbf{A}, \mathbf{X})$ denote a given graph, where $\mathcal{V} = \{v_i\}_{i=1}^N$ represents the node set. The adjacency matrix and the feature matrix are denoted as $\mathbf{A} \in \{0, 1\}^{N \times N}$ and $\mathbf{X} \in \mathbb{R}^{N \times d}$ respectively, where $\mathbf{x}_i \in \mathbb{R}^d$ is the feature of v_i and $\mathbf{A}_{ij} = 1$ iff there is an edge between v_i and v_j . In the setting of graph self-supervised learning, there is no available label information during training. Our goal is to learn a GNN encoder $f(\cdot)$ receiving the graph structure and features, and producing low-dimensional node representations. We denote $\mathbf{Z} = f(\mathbf{X}, \mathbf{A}) \in \mathbb{R}^{N \times d'}$ as the learned node representations, where $\mathbf{z}_i \in \mathbb{R}^{d'}$ is the representation of node v_i . The representations are l_2 -normalized on the unit hypersphere $\mathcal{S}^{d'-1}$, which is common in machine learning.

3.2. Graph Masked Autoencoders

We choose the canonical GraphMAE (Hou et al. 2022) as the object of analysis because it serves as the foundation for various subsequent models. GraphMAE first randomly sample a subset of nodes $\tilde{\mathcal{V}} \in \mathcal{V}$ based on a uniform distribution.

Then, the node features of these selected nodes are masked:

$$\tilde{\mathbf{x}}_i = \begin{cases} \mathbf{x}_{[\text{MASK}]} & v_i \in \tilde{\mathcal{V}}, \\ \mathbf{x}_i & v_i \notin \tilde{\mathcal{V}}, \end{cases} \quad (1)$$

where $\mathbf{x}_{[\text{MASK}]} \in \mathbb{R}^d$ is the learnable mask token, and $\tilde{\mathbf{x}}_i \in \mathbb{R}^d$ is the feature of node v_i after masking.

The GraphMAE model $h = g \circ f$ is an encoder-decoder architecture, where a GNN-based encoder f maps the contexts (ego-graphs) of masked nodes to latent features, and a GNN-based decoder g reconstructs the features of masked nodes from the latent contexts. The task performed by GraphMAE can be interpreted as the reconstruction of the original features of masked nodes from their $(l_e + l_d)$ -hop contexts, where l_e and l_d denote the numbers of encoder layers and decoder layers, respectively. We use c_i to denote the $(l_e + l_d)$ -hop context of node v_i after masking. The reconstructed feature of node v_i is $\hat{\mathbf{x}}_i = h(c_i) = g(f(c_i))$.

Finally, GraphMAE adopts the scaled cosine error (SCE) on masked features as the reconstruction loss:

$$\mathcal{L}_{\text{SCE}} = \mathbb{E}_{v_i \in \tilde{\mathcal{V}}} (1 - \mathbf{x}_i^\top h(c_i))^\gamma, \quad (2)$$

where the cosine similarity is simply represented as the dot product since the original feature and reconstructed feature are both l_2 -normalized. The scaling factor $\gamma \geq 1$ is a hyperparameter that adjusts the weight of each sample with the reconstruction error.

3.3. Alignment and Uniformity Loss

The alignment and uniformity properties are necessary for high-quality representations, and highly related to the contrastive learning (Wang and Isola 2020).

Alignment loss aims to make the representations of semantically similar samples as close as possible, and thus the representations can be invariant to unneeded noise factors. Alignment loss is consistent with the contrastive objective of maximizing the agreement of positive pairs. Formally, the alignment loss is defined as:

$$\mathcal{L}_{\text{Align}} = \mathbb{E}_{(v, v^+) \sim p_{\text{pos}}} \|\mathbf{z} - \mathbf{z}^+\|^2. \quad (3)$$

where p_{pos} is the distribution of positive pairs, and \mathbf{z} is the learned representations of data sample v .

Uniformity loss prefers the uniform distribution on the unit hypersphere, so as to preserve maximal information of data. Uniformity helps to avoid feature collapse and learn discriminable representations. The uniformity loss is defined as the logarithm of the average pairwise Gaussian potential:

$$\mathcal{L}_{\text{Uni}} = \log \mathbb{E}_{v_i, v_j \stackrel{\text{i.i.d.}}{\sim} p_{\text{data}}} e^{-t \|\mathbf{z}_i - \mathbf{z}_j\|^2}, \quad (4)$$

where p_{data} is the distribution of data, and t is a hyperparameter for Gaussian potential kernel. In the contrastive objective, uniformity is achieved by pulling away the distance between negative pairs.

4. Alignment-Uniformity Enhanced Graph Masked Autoencoders

In this section, we first conduct a theoretical analysis of GraphMAE and identify its limitations from the perspective of alignment and uniformity. Subsequently, we propose our AUG-MAE model to overcome these limitations.

4.1. Theoretical Understanding of GraphMAE

There is a viewpoint that generative and contrastive methods adhere to different philosophies, where contrastive methods deal with the inter-data information and generative methods focus on the intra-data information (Wu et al. 2021). However, we perform a deep analysis and give an insight that *generative methods, such as GraphMAE, perform implicit context-level graph contrastive learning*.

Since GraphMAE is based on the autoencoder framework, we first assume that it is capable of accomplishing the vanilla autoencoding task: reconstructing the original input.

Assumption 4.1. For any graph decoder g , we assume that there exists a pseudo-inverse graph encoder f_g such that the resulting pseudo graph autoencoder $h_g = g \circ f_g$ satisfies $\mathbb{E}_{v_i \in \tilde{\mathcal{V}}} \|h_g(\mathbf{x}) - \mathbf{x}\|^2 \leq \varepsilon$.

This assumption is valid, since the GNN-based encoder and decoder degenerate to MLPs when input contains only one node, and MLPs have been proven to be universal approximators of arbitrary continuous functions (Lu et al. 2017).

Theorem 4.2. Under Assumption 4.1, the SCE loss in Eq. (2) can be lower bounded by a pretext loss:

$$\mathcal{L}_{\text{SCE}}(h) \geq \frac{\gamma}{2} \mathcal{L}_{\text{Pretext}}(h) - \frac{\gamma}{2} \varepsilon + \text{const}, \quad (5)$$

where $\mathcal{L}_{\text{Pretext}}(h) = -\mathbb{E}_{v_i \in \tilde{\mathcal{V}}} h_g(\mathbf{x}_i)^\top h(c_i)$.

Please refer to Appendix for the detailed proof of Theorem 4.2. Then we define the context-level alignment loss, which is an objective of context-level GCL, and we prove it is a lower bound of pretext loss above.

Definition 4.3. (Context-Level Alignment Loss) The alignment loss for positive context pairs (c, c^+) is defined as:

$$\mathcal{L}_{\text{Align}}^c(h) = -\mathbb{E}_{(c, c^+) \sim p_{\text{pos}}^c} h(c)^\top h(c^+). \quad (6)$$

This loss in the form of dot product and the loss in the form of mean squared error in Eq. (3) are equivalent because the reconstructed features $\{h(c)\}_{c \sim p_{\text{data}}^c}$ are all normalized.

Theorem 4.4. The pretext loss in Eq. (5) can be lower bounded by the context-level alignment loss in Eq. (6):

$$\mathcal{L}_{\text{Pretext}}(h) \geq \frac{1}{2} \mathcal{L}_{\text{Align}}^c(h) + \text{const}. \quad (7)$$

The detailed proof of Theorem 4.4 can be found in Appendix. Combining Theorem 4.2 and Theorem 4.4, we arrive at the main theorem showing that GraphMAE's node-level reconstruction loss can be bounded by the alignment loss of the positive context pairs drawn from the masked nodes.

Theorem 4.5. Under Assumption 4.1, GraphMAE's node-level reconstruction loss in Eq. (2) can be lower bounded by the context-level alignment loss in Eq. (6):

$$\mathcal{L}_{\text{SCE}}(h) \geq \frac{\gamma}{4} \mathcal{L}_{\text{Align}}^c(h) - \frac{\gamma}{2} \varepsilon + \text{const} \quad (8)$$

Following Theorem 4.5, a small GraphMAE's reconstruction loss implies a small context-level alignment loss, which indicates that GraphMAE implicitly aligns the representations of positive context pairs.

Intuitive explanation. Here, we provide an intuitive explanation for our theoretical result. We reiterate that GraphMAE employs the autoencoder to reconstruct the masked

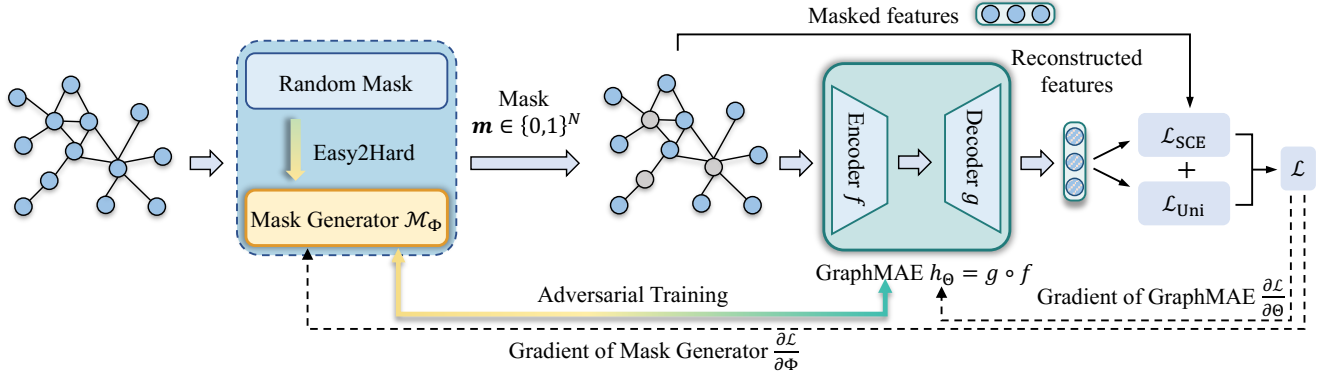


Figure 2: The overall framework of our proposed AUG-MAE model. We propose an easy-to-hard adversarial masking strategy to provide hard-to-align positive pairs, so as to improve the alignment ability of GraphMAE. Additionally, we introduce an explicit uniformity regularizer \mathcal{L}_{Uni} into the objective to enhance the uniformity of learned representations.

node features based on the contexts of these mask nodes. When two (or more) masked nodes have the same or very similar features, then their contexts have the same reconstruction goal, and thus these contexts will be encoded as similar latent representations. In the paradigm of contrastive learning, these contexts can be considered as positive pairs.

4.2. Limitations of GraphMAE

The theoretical analysis above builds a bridge between GraphMAE and GCL. Therefore, we can identify the limitations of GraphMAE with the well-established tools in GCL. Specifically, we further discuss the limitations of GraphMAE from the perspective of alignment and uniformity.

For alignment, although GraphMAE is proven to have the ability to align positive pairs, the practical alignment effect not only depends on the model’s ability, but also *influenced by the masking strategy*. Concretely, the optimization objective is computed as an expectation over a distribution, which is essentially decided by the distribution of the mask. Further, the uniform random masking strategy adopted by GraphMAE ignores the imbalanced distribution of easy and hard samples, thus cannot provide enough information about hard-to-align pairs.

For uniformity, the representation uniformity is *not strictly guaranteed in GraphMAE*. Specifically, GraphMAE avoids full feature collapse, i.e., the learned representations do not collapse to a fixed point in the feature space, as long as the masked features are not identical. However, reconstruction loss cannot lead to uniformly distributed representations. In Fig. 1, we observe that the representations still suffer from partial dimensional collapse. Therefore, the representation uniformity can also be improved.

4.3. The Proposed AUG-MAE Model

To overcome the aforementioned limitations, we propose AUG-MAE. The overall framework is illustrated in Fig. 2.

Adversarial Masking GraphMAE randomly selects nodes to mask based on the uniform distribution, which ignores the imbalanced distribution of easy and hard samples.

Random masking is able to sample a large number of easy-to-align positive pairs, but yields only a few hard-to-align positive pairs. Therefore, GraphMAE does not work well on these hard samples, limiting the quality of the learned representations. For this reason, we propose an adversarial masking strategy to mine more hard-to-align positive pairs.

To generate mask adaptively, we devise a GNN-based mask generator \mathcal{M} with parameter Φ . Given a graph \mathcal{G} , the mask generator produces a probability vector of adversarial masking $\text{prob}_{\text{adv}} \in [0, 1]^N$:

$$\text{prob}_{\text{adv}} = \mathcal{M}_{\Phi}(\mathcal{G}), \quad (9)$$

where $\text{prob}_{\text{adv},i}$ denotes the probability of node v_i being masked. Then, the Gumbel-Softmax (Jang, Gu, and Poole 2017) is applied on each $\text{prob}_{\text{adv},i}$ to generate a differentiable binary mask vector $\mathbf{m} \in \{0, 1\}^N$:

$$\mathbf{m}_i = \sigma\left(\frac{1}{\tau}(\log(\frac{\text{prob}_{\text{adv},i}}{1 - \text{prob}_{\text{adv},i}}) + (\epsilon_0 - \epsilon_1))\right), \quad (10)$$

where ϵ_0, ϵ_1 are Gumbel random noises sampled from $\text{Gumbel}(0, 1)$, τ is a temperature hyperparameter, and σ is the sigmoid function. We use the $\tilde{V}_{\text{adv}} = \{v_i | \mathbf{m}_i = 1, i = 1, 2, \dots, N\}$ to represent the set of masked nodes. Then, the node features are masked with generated mask:

$$\tilde{\mathbf{x}}_i = \begin{cases} \mathbf{x}_{[\text{MASK}]} & v_i \in \tilde{V}_{\text{adv}}, \\ \mathbf{x}_i & v_i \notin \tilde{V}_{\text{adv}}. \end{cases} \quad (11)$$

In previous studies (He et al. 2022; Hou et al. 2022), mask ratio has been empirically proven to be important for masked autoencoding. However, Eq. (9) and (10) cannot control the ratio of masked nodes. To adjust the mask ratio of the mask generator, we introduce a ratio regularizer in the form of $1/\sin(\cdot)$ to its optimization objective. Formally, the parameters of the mask generator Φ can be updated by optimizing:

$$\Phi^* = \arg \max_{\Phi} (\mathcal{L}_{\text{SCE}}(\mathcal{G}; \Theta, \Phi) - \lambda_1 \sin(\frac{\pi}{N} \sum_{i=1}^N \mathbf{m}_i)^{-1}), \quad (12)$$

where $\frac{1}{N} \sum_{i=1}^N m_i$ is mask ratio. λ_1 is the weight of the ratio regularizer, which controls the mask ratio. The large λ_1 encourages a mask ratio close to 50%. Adversarially, the parameters of GraphMAE Θ can be learned by:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}_{\text{SCE}}(\mathcal{G}; \Theta, \Phi). \quad (13)$$

During the adversarial training process, Θ and Φ are iteratively optimized so that the mask generator and GraphMAE evolve simultaneously. The mask generator gradually learns to generate hard-to-align positive pairs, while GraphMAE gradually learns how to align these pairs.

Easy-to-Hard Training To ensure the training stability, we adopt an easy-to-hard strategy for training. In the early stage of training, we employ the random masking to generate the mask, utilizing plenty of easy samples to initialize the model parameters. The advantage of this initialization is that the model can initially have the ability to handle easy samples and have a relatively accurate judgment of the difficulty of the samples. During the training process, we gradually increase the weight of the adversarial masking and decrease the weight of the random masking, so that the model can obtain further improvement from the hard samples in the later stage of training. The easy-to-hard process is controlled by:

$$prob(t) = (1 - \alpha_{\text{adv}}(t)) \cdot prob_{\text{rand}} + \alpha_{\text{adv}}(t) \cdot prob_{\text{adv}}(t), \quad (14)$$

where t denotes the current epoch, T denotes the total epochs, and $prob \in [0, 1]^N$ denotes the masking probability vector, which is the weighted sum of the masking probability vector of random masking $prob_{\text{rand}}$ and that of adversarial masking $prob_{\text{adv}}$. Then the mask m is sampled from $prob$.

During the training process (from epoch 0 to epoch T), the weight of the adversarial mask α_{adv} grows from α_0 to α_T . Correspondingly, the weight of the random mask decreases from $(1 - \alpha_0)$ to $(1 - \alpha_T)$. The change of α_{adv} is defined as:

$$\alpha_{\text{adv}}(t) = \alpha_0 + \Delta\alpha(t) = \alpha_0 + \left(\frac{t}{T}\right)^\eta \cdot (\alpha_T - \alpha_0), \quad (15)$$

where $\alpha_0, \alpha_T \in [0, 1]$, $\alpha_0 < \alpha_T$. η controls the rate of weight growth. $\eta = 1$ indicates a linear growth from easy to hard, and $\eta \neq 1$ indicates a non-linear growth.

Explicit Uniformity Regularizer As mentioned in theoretical analysis, GraphMAE naturally avoids full feature collapse, but still suffers from partial dimensional collapse. Therefore, we explicitly introduce a uniformity regularizer into the objective of GraphMAE. The objective defined by Eq. (13) is updated as:

$$\Theta^* = \arg \min_{\Theta} (\mathcal{L}_{\text{SCE}}(\mathcal{G}; \Theta, \Phi) + (1 - \alpha_{\text{adv}}) \lambda_2 \mathcal{L}_{\text{Uni}}(\mathcal{G}; \Theta)), \quad (16)$$

where \mathcal{L}_{Uni} denotes the uniformity regularizer given in Eq. (4), and λ_2 is the weight of uniformity regularizer. It is worth noting that we desire that the representations of all nodes are uniformly distributed on the hypersphere, and not just the representations of hard samples. Therefore, uniformity regularization should be paired with the random masking strategy. To achieve it, we also use $(1 - \alpha_{\text{adv}})$ to control the impact of regularizer during the training process.

To help better understand the adversarial training process, we provide the brief pseudo-code of it in Appendix.

5. Experiments

In this section, we conduct experiments to evaluate the effectiveness of AUG-MAE. We analyze it by answering the following questions: **RQ1**: How does AUG-MAE perform compared with graph self-supervised learning baselines, especially GraphMAE, in various downstream tasks? **RQ2**: How does each component influence the performance of AUG-MAE? **RQ3**: How do key hyper-parameters influence the performance of AUG-MAE? **RQ4**: How does the alignment and uniformity of the representations learned by AUG-MAE, compare with GCL and GraphMAE?

5.1. Evaluation Setups

Datasets. We select seven node classification datasets (i.e., Cora, Citeseer (Sen et al. 2008), Pubmed (Namata et al. 2012), Ogbn-arxiv (Hu et al. 2020), PPI, Reddit, and Cora-full (Bojchevski and Günnemann 2018)), and six graph classification datasets (i.e., IMDB-B, IMDB-M, PROTEINS, COLLAB, MUTAG, and REDDIT-B (Morris et al. 2020)).

Baselines. We consider both contrastive methods and generative methods as baselines. Node-level GCL baselines are compared in the node classification task, including DGI (Velickovic et al. 2019), MVGRL (Hassani and Ahmadi 2020), GRACE, BGRL (Thakoor et al. 2022), InfoGCL (Xu et al. 2021), and CCA-SSG (Zhang et al. 2021). In graph classification task, compared graph-level GCL baselines are Graph2vec (Narayanan et al. 2017), InfoGraph (Sun et al. 2020), GraphCL, JOAO (You et al. 2021), GCC, MVGRL, and InfoGCL. For generative methods, we select SeeGera, MaskGAE, and GraphMAE as baselines.

Detailed evaluation setups can be found in Appendix.

5.2. Performance Comparison (RQ1)

We compare AUG-MAE with the baselines and the results are summarized in Table 1 and Table 2. On both node classification and graph classification tasks, AUG-MAE outperforms all graph self-supervised baselines on most datasets.

Among the baseline models, GraphMAE, as a recently proposed generative method, achieves competitive performance with state-of-the-art contrastive methods. Since our work focuses on analyzing the limitations of GraphMAE and improving it, we first focus on comparing our method with GraphMAE. Our method outperforms GraphMAE on all datasets for both downstream tasks. This verifies the feasibility of improving GraphMAE from the perspective of alignment and uniformity, as well as validates the effectiveness of our proposed strategies.

On the downstream node classification task, the representations learned by our AUG-MAE are able to achieve the highest accuracy on all datasets except CiteSeer. On the graph classification task, the representations learned by our AUG-MAE also have the highest accuracy on most datasets. However, on the MUTAG and REDDIT-B datasets, although our method outperforms GraphMAE, it still does not outperform some GCL methods. We speculate that the reason may be that generative methods focus more on context-level information, while node-level and graph-level information are also important on these datasets.

| | Method | Cora | CiteSeer | PubMed | Ogbn-arxiv | PPI | Reddit | Corafull |
|-------------|----------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| Contrastive | DGI | 82.3 ± 0.6 | 71.8 ± 0.7 | 76.8 ± 0.6 | 70.3 ± 0.2 | 63.8 ± 0.2 | 94.0 ± 0.1 | 48.2 ± 0.5 |
| | MVGRL | 83.5 ± 0.4 | 73.3 ± 0.5 | 80.1 ± 0.7 | - | - | - | 52.6 ± 0.5 |
| | GRACE | 81.9 ± 0.4 | 71.2 ± 0.5 | 80.6 ± 0.4 | 71.5 ± 0.1 | 69.7 ± 0.2 | 94.7 ± 0.1 | 45.2 ± 0.1 |
| | BGRL | 82.7 ± 0.6 | 71.1 ± 0.8 | 79.6 ± 0.5 | <u>71.6 ± 0.1</u> | 73.6 ± 0.2 | 94.2 ± 0.1 | 47.4 ± 0.5 |
| | InfoGCL | 83.5 ± 0.3 | 73.5 ± 0.4 | 79.1 ± 0.2 | - | - | - | - |
| | CCA-SSG | <u>84.0 ± 0.4</u> | 73.1 ± 0.3 | 81.0 ± 0.4 | 71.2 ± 0.2 | 73.3 ± 0.2 | 95.1 ± 0.1 | <u>53.5 ± 0.4</u> |
| Generative | SeeGera | 82.8 ± 0.3 | 71.6 ± 0.2 | 79.2 ± 0.3 | 71.2 ± 0.3 | 73.4 ± 0.3 | 95.2 ± 0.2 | 52.0 ± 0.4 |
| | MaskGAE | 82.6 ± 0.3 | 73.1 ± 0.6 | <u>81.0 ± 0.3</u> | 71.2 ± 0.3 | 73.9 ± 0.3 | 95.4 ± 0.1 | 52.2 ± 0.1 |
| | GraphMAE | 84.0 ± 0.6 | 73.1 ± 0.4 | 80.9 ± 0.4 | 71.3 ± 0.6 | <u>74.1 ± 0.4</u> | <u>95.8 ± 0.4</u> | 53.3 ± 0.4 |
| | AUG-MAE | 84.3 ± 0.4 | 73.2 ± 0.4 | 81.4 ± 0.4 | 71.9 ± 0.2 | 74.3 ± 0.1 | 96.1 ± 0.1 | 57.6 ± 0.3 |

Table 1: Node classification results on benchmarks. We report Micro-F1(%) score for PPI and accuracy(%) for the other datasets. The best results are highlighted in bold and the runner ups are highlighted with underlines.

| | Method | IMDB-B | IMDB-M | PROTEINS | COLLAB | MUTAG | REDDIT-B |
|-------------|-----------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Contrastive | Graph2vec | 71.10 ± 0.54 | 50.44 ± 0.87 | 73.30 ± 2.05 | - | 83.15 ± 9.25 | 75.78 ± 1.03 |
| | InfoGraph | 73.03 ± 0.87 | 49.69 ± 0.53 | 74.44 ± 0.31 | 70.65 ± 1.13 | 89.01 ± 1.13 | 82.50 ± 1.42 |
| | GraphCL | 71.14 ± 0.44 | 48.58 ± 0.67 | 74.39 ± 0.45 | 71.36 ± 1.15 | 86.80 ± 1.34 | <u>89.53 ± 0.84</u> |
| | JOAO | 70.21 ± 3.08 | 49.20 ± 0.77 | 74.55 ± 0.41 | 69.50 ± 0.36 | 87.35 ± 1.02 | 85.29 ± 1.35 |
| | GCC | 72.0 | 49.4 | - | 78.9 | - | 89.8 |
| | MVGRL | 74.20 ± 0.70 | 51.20 ± 0.50 | - | - | <u>89.70 ± 1.10</u> | 84.50 ± 0.60 |
| Generative | InfoGCL | 75.10 ± 0.90 | <u>51.40 ± 0.80</u> | - | 80.00 ± 1.30 | 91.20 ± 1.30 | - |
| | GraphMAE | <u>75.30 ± 0.59</u> | 51.35 ± 0.78 | <u>75.30 ± 0.52</u> | <u>80.32 ± 0.42</u> | 88.19 ± 1.26 | 87.83 ± 0.25 |
| | AUG-MAE | 75.56 ± 0.61 | 51.80 ± 0.86 | 75.83 ± 0.24 | 80.48 ± 0.50 | 88.28 ± 0.98 | 87.98 ± 0.43 |

Table 2: Graph classification results on benchmarks. We report accuracy(%) for all datasets. The best results are highlighted in bold and the runner ups are highlighted with underlines.

5.3. Ablation Study (RQ2)

Table 3 summarizes the results of ablation study, from which we have the following observations.

Effect of adversarial masking. By comparing GraphMAE and Variant 1 (also Variant 2 and Variant 4), we can observe that the adversarial masking is better than random masking, and steadily boosts the performance on node classification. On the graph classification task, the straightforward introduction of adversarial masking does not seem to be helpful. But when combined with the easy-to-hard training strategy, Adv+E2H can effectively improve the performance. This can be observed by comparing Variant 2 and AUG-MAE.

Effect of easy-to-hard training. The easy-to-hard strategy is designed to assist in the adversarial masking. From the difference between the results of Variant 1 and Variant 3 (also Variant 4 and AUG-MAE) on graph classification datasets, we can find that the easy-to-hard training is very important on graph-level representation learning.

Effect of uniformity regularizer. By comparing GraphMAE and Variant 2 (also Variant 3 and AUG-MAE), we can observe that uniformity regularizer effectively improves the performance of both node classification and graph classification, which verifies the effectiveness of this regularizer.

5.4. Sensitivity Analysis (RQ3)

Fig. 3 shows the effect of varied hyper-parameter values, from which we have the following observations.

Effect of weight of ratio regularizer λ_1 . This weight affects the result of representation learning by affecting the mask ratio. As shown in Fig. 3(a), when $\lambda_1 \geq 1$, the mask ratio is around 0.5 and the best performance is achieved on Cora. When $\lambda_1 = 1e-3$, the mask ratio is around 0.75 and the best performance is achieved on PubMed and COLLAB. They are consistent with the optimal mask ratios for random masking provided by GraphMAE (Hou et al. 2022).

Effect of weight of uniformity regularizer λ_2 . The optimal choice for this weight on most datasets is $5e-4$ or $5e-5$. We find that uniformity regularizer impairs accuracy when the weight is larger than $1e-1$. This is because the excessive pursuit of uniformity can damage the distinguishability of learned representations.

Effect of parameters controlling easy-to-hard α_0, α_T, η . From Fig. 3(c) and Fig. 3(d) we observe that the impact of α_0 and α_T on graph classification datasets is obvious. Taking COLLAB as an example, appropriate values of α_0 and α_T can effectively improve the performance. We tune η in the range of $[0.5, 1.5]$, and find that our model is not sensitive to the change of η . Relatively speaking, the best performance is achieved with $\eta = 1$ on most of the datasets.

5.5. Alignment and Uniformity Analysis (RQ4)

On the Cora dataset, we take nodes with the same label as positive samples, and compute l_2 distance between them (also called supervised alignment loss (Wang et al. 2022)).

| Model | Component | | | Node Classification Dataset | | | | | |
|-----------|-----------|-----|-----|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| | Adv | E2H | Uni | Cora | CiteSeer | PubMed | Ogbn-arxiv | PPI | Reddit |
| GraphMAE | - | - | - | 84.04 \pm 0.58 | 73.11 \pm 0.38 | 80.94 \pm 0.47 | 71.32 \pm 0.55 | 74.09 \pm 0.37 | 95.79 \pm 0.36 |
| Variant 1 | ✓ | - | - | 84.30 \pm 0.75 | 73.18 \pm 0.59 | 81.25 \pm 0.53 | 71.43 \pm 0.07 | 74.12 \pm 0.36 | 95.97 \pm 0.36 |
| Variant 2 | - | - | ✓ | 84.24 \pm 0.58 | 73.16 \pm 0.54 | 81.18 \pm 0.44 | 71.42 \pm 0.30 | 74.13 \pm 0.20 | 95.90 \pm 0.15 |
| Variant 3 | ✓ | ✓ | - | 84.20 \pm 0.54 | 73.14 \pm 0.58 | 81.28 \pm 0.43 | 71.50 \pm 0.30 | 74.24 \pm 0.05 | 95.85 \pm 0.25 |
| Variant 4 | ✓ | - | ✓ | 84.32 \pm 0.45 | 73.18 \pm 0.48 | 81.30 \pm 0.54 | 71.50 \pm 0.15 | 74.13 \pm 0.36 | 96.00 \pm 0.10 |
| AUG-MAE | ✓ | ✓ | ✓ | 84.30 \pm 0.38 | 73.20 \pm 0.44 | 81.35 \pm 0.44 | 71.86 \pm 0.22 | 74.30 \pm 0.11 | 96.07 \pm 0.03 |

| Model | Component | | | Graph Classification Dataset | | | | | |
|-----------|-----------|-----|-----|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| | Adv | E2H | Uni | IMDB-B | IMDB-M | PROTEINS | COLLAB | MUTAG | REDDIT-B |
| GraphMAE | - | - | - | 75.30 \pm 0.59 | 51.35 \pm 0.78 | 75.30 \pm 0.52 | 80.32 \pm 0.42 | 88.19 \pm 1.26 | 87.83 \pm 0.25 |
| Variant 1 | ✓ | - | - | 74.32 \pm 0.75 | 49.92 \pm 0.92 | 75.72 \pm 1.04 | 79.88 \pm 0.69 | 87.00 \pm 1.40 | 87.35 \pm 0.35 |
| Variant 2 | - | - | ✓ | 75.40 \pm 0.60 | 51.50 \pm 0.49 | 75.50 \pm 0.51 | 80.37 \pm 0.44 | 88.20 \pm 1.56 | 87.90 \pm 0.42 |
| Variant 3 | ✓ | ✓ | - | 75.20 \pm 0.88 | 51.59 \pm 1.36 | 75.65 \pm 0.53 | 80.26 \pm 0.44 | 88.01 \pm 1.11 | 87.90 \pm 0.27 |
| Variant 4 | ✓ | - | ✓ | 74.56 \pm 0.58 | 50.22 \pm 0.69 | 75.75 \pm 0.43 | 80.07 \pm 0.45 | 87.41 \pm 1.27 | 87.47 \pm 0.45 |
| AUG-MAE | ✓ | ✓ | ✓ | 75.56 \pm 0.61 | 51.80 \pm 0.86 | 75.83 \pm 0.24 | 80.48 \pm 0.50 | 88.28 \pm 0.98 | 87.98 \pm 0.43 |

* Adv: adversarial masking, E2H: easy-to-hard training, Uni: uniformity regularizer.

Table 3: Ablation analysis, in which we keep different components in our AUG-MAE to form variants. We report accuracy(%) of these variants for node and graph classification datasets. The best performance is highlighted in bold.

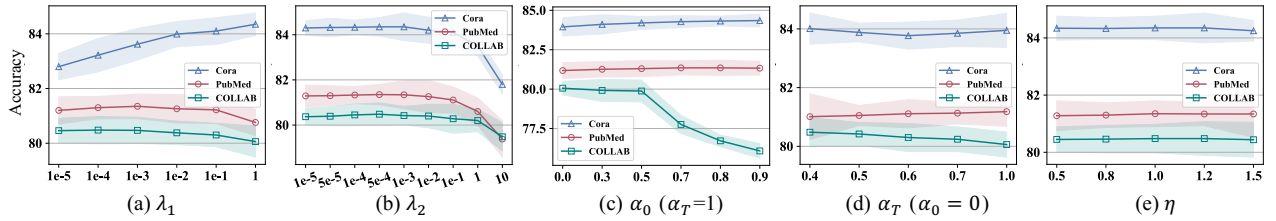


Figure 3: Effect of different hyper-parameters. The y-axis represents accuracy(%).

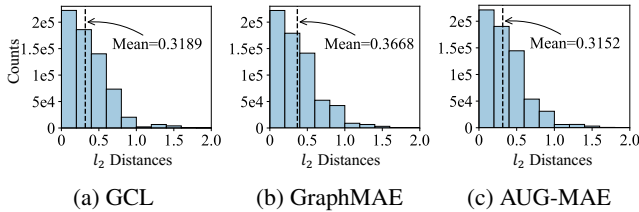


Figure 4: l_2 distances between positive representations of Cora learned by GCL, GraphMAE, and AUG-MAE. The smaller mean distance indicates the better alignment.

The statistical results are shown in Fig. 4, in which we also plot the mean distances with dashed lines. Since smaller mean distance indicates better alignment, we can observe that AUG-MAE can align similar samples better than GraphMAE, and even slightly better than GCL.

We also visualize the representation distributions learned by GCL, GraphMAE, and our AUG-MAE in Fig. 5. Compared with GraphMAE, representations learned by AUG-MAE achieve better uniformity, i.e., representations are more uniformly distributed on \mathcal{S}^1 .

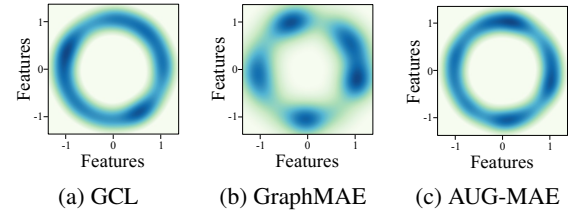


Figure 5: Representation distributions of Cora on \mathcal{S}^1 learned by GCL, GraphMAE, and AUG-MAE. We plot distributions with Gaussian kernel density estimation in \mathbb{R}^2 .

6. Conclusion

We theoretically prove that the node-level reconstruction in GraphMAE implicitly performs context-level GCL. Based on this, we identify the limitations of GraphMAE from the perspective of alignment and uniformity. To overcome them, we propose AUG-MAE equipped with an easy-to-hard adversarial masking strategy and an explicit uniformity regularizer. Experimental results show that AUG-MAE produces representations with better alignment and uniformity, and surpasses baselines on downstream tasks.

Acknowledgments

This work is jointly sponsored by National Natural Science Foundation of China (62206291, 62141608).

References

- Bojchevski, A.; and Günnemann, S. 2018. Deep Gaussian Embedding of Graphs: Unsupervised Inductive Learning via Ranking. In *ICLR*.
- Cheng, J.; Li, M.; Li, J.; and Tsung, F. 2023. Wiener Graph Deconvolutional Network Improves Graph Self-Supervised Learning. In *AAAI*.
- Guo, X.; Wang, Y.; Du, T.; and Wang, Y. 2023. ContraNorm: A Contrastive Learning Perspective on Oversmoothing and Beyond. In *ICLR*.
- Hassani, K.; and Ahmadi, A. H. K. 2020. Contrastive Multi-View Representation Learning on Graphs. In *ICML*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked Autoencoders Are Scalable Vision Learners. In *CVPR*.
- Hou, Z.; Liu, X.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. GraphMAE: Self-Supervised Masked Graph Autoencoders. In *KDD*.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In *NeurIPS*.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *ICLR*.
- Jing, L.; Vincent, P.; LeCun, Y.; and Tian, Y. 2022. Understanding Dimensional Collapse in Contrastive Self-supervised Learning. In *ICLR*.
- Li, J.; Wu, R.; Sun, W.; Chen, L.; Tian, S.; Zhu, L.; Meng, C.; Zheng, Z.; and Wang, W. 2022. MaskGAE: Masked Graph Modeling Meets Graph Autoencoders. *arXiv*.
- Li, X.; Ye, T.; Shan, C.; Li, D.; and Gao, M. 2023. SeeGera: Self-Supervised Semi-Implicit Graph Variational Auto-Encoders with Masking. In *WWW*.
- Liu, Y.; Jin, M.; Pan, S.; Zhou, C.; Zheng, Y.; Xia, F.; and Yu, P. 2022. Graph Self-Supervised Learning: A Survey. *TKDE*.
- Lu, Z.; Pu, H.; Wang, F.; Hu, Z.; and Wang, L. 2017. The Expressive Power of Neural Networks: A View from the Width. In *NeurIPS*.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. In *Graph Representation Learning and Beyond@ICML*.
- Namata, G.; London, B.; Getoor, L.; and Huang, B. 2012. Query-driven Active Surveying for Collective Classification. In *International Workshop on Mining and Learning with Graphs*.
- Narayanan, A.; Chandramohan, M.; Venkatesan, R.; Chen, L.; Liu, Y.; and Jaiswal, S. 2017. graph2vec: Learning Distributed Representations of Graphs. *arXiv*.
- Pu, S.; Zhao, K.; and Zheng, M. 2022. Alignment-Uniformity aware Representation Learning for Zero-shot Video Classification. In *CVPR*.
- Qiu, J.; Chen, Q.; Dong, Y.; Zhang, J.; Yang, H.; Ding, M.; Wang, K.; and Tang, J. 2020. GCC: Graph Contrastive Coding for Graph Neural Network Pre-Training. In *KDD*.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Gallagher, B.; and Eliassi-Rad, T. 2008. Collective Classification in Network Data. *AI Magazine*.
- Sun, F.; Hoffmann, J.; Verma, V.; and Tang, J. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *ICLR*.
- Tan, Q.; Liu, N.; Huang, X.; Choi, S.-H.; Li, L.; Chen, R.; and Hu, X. 2023. S2GAE: Self-Supervised Graph Autoencoders Are Generalizable Learners with Graph Masking. In *WSDM*.
- Thakoor, S.; Tallec, C.; Azar, M. G.; Azabou, M.; Dyer, E. L.; Munos, R.; Velickovic, P.; and Valko, M. 2022. Large-Scale Representation Learning on Graphs via Bootstrapping. In *ICLR*.
- Tian, Y.; Dong, K.; Zhang, C.; Zhang, C.; and Chawla, N. V. 2023. Heterogeneous Graph Masked Autoencoders. In *AAAI*.
- Velickovic, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. In *ICLR*.
- Wang, C.; Yu, Y.; Ma, W.; Zhang, M.; Chen, C.; Liu, Y.; and Ma, S. 2022. Towards Representation Alignment and Uniformity in Collaborative Filtering. In *KDD*.
- Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *ICML*.
- Wu, L.; Lin, H.; Gao, Z.; Tan, C.; and Li, S. Z. 2021. Self-Supervised Learning on Graphs: Contrastive, Generative, or Predictive. *TKDE*.
- Xia, J.; Wu, L.; Chen, J.; Hu, B.; and Li, S. Z. 2022. SimGRACE: A Simple Framework for Graph Contrastive Learning without Data Augmentation. In *WWW*.
- Xia, L.; Huang, C.; Huang, C.; Lin, K.; Yu, T.; and Kao, B. 2023. Automated Self-Supervised Learning for Recommendation. In *WWW*.
- Xie, Y.; Xu, Z.; Zhang, J.; Wang, Z.; and Ji, S. 2022. Self-Supervised Learning of Graph Neural Networks: A Unified Review. *TPAMI*.
- Xu, D.; Cheng, W.; Luo, D.; Chen, H.; and Zhang, X. 2021. InfoGCL: Information-Aware Graph Contrastive Learning. In *NeurIPS*.
- Ye, Y.; Xia, L.; and Huang, C. 2023. Graph Masked Autoencoder for Sequential Recommendation. In *SIGIR*.
- You, Y.; Chen, T.; Shen, Y.; and Wang, Z. 2021. Graph Contrastive Learning Automated. In *ICML*.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In *NeurIPS*.
- Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. Are Graph Augmentations Necessary?: Simple Graph Contrastive Learning for Recommendation. In *SIGIR*.

- Zhang, H.; Wu, Q.; Yan, J.; Wipf, D.; and Yu, P. S. 2021. From Canonical Correlation Analysis to Self-supervised Graph Neural Networks. In *NeurIPS*.
- Zhang, M.; Zhu, Y.; Liu, Q.; Wu, S.; and Wang, L. 2022a. Deep Contrastive Multiview Network Embedding. In *CIKM*.
- Zhang, Q.; Wang, Y.; and Wang, Y. 2022. How Mask Matters: Towards Theoretical Understandings of Masked Autoencoders. In *NeurIPS*.
- Zhang, S.; Chen, H.; Yang, H.; Sun, X.; Yu, P. S.; and Xu, G. 2022b. Graph Masked Autoencoders with Transformers. *arXiv*.
- Zhu, Y.; Xu, Y.; Cui, H.; Yang, C. J.; Liu, Q.; and Wu, S. 2022. Structure-Enhanced Heterogeneous Graph Contrastive Learning. In *SDM*.
- Zhu, Y.; Xu, Y.; Liu, Q.; and Wu, S. 2021a. An Empirical Study of Graph Contrastive Learning. In *NeurIPS Datasets and Benchmarks*.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2020. Deep Graph Contrastive Representation Learning. In *Graph Representation Learning and Beyond@ICML*.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2021b. Graph Contrastive Learning with Adaptive Augmentation. In *WWW*.