

Single-View Graph Contrastive Learning with Soft Neighborhood Awareness

Qingqiang Sun¹, Chaoqi Chen², Ziyue Qiao^{1*}, Xubin Zheng¹, Kai Wang^{3†},

¹Great Bay University

²Shenzhen University

³Central South University

{qqsun, zyqiao, xbzheng}@gbu.edu.cn, cqchen1994@gmail.com, kaiwang@csu.edu.cn

Abstract

Most graph contrastive learning (GCL) methods heavily rely on cross-view contrast, thus facing several concomitant challenges, such as the complexity of designing effective augmentations, the potential for information loss between views, and increased computational costs. To mitigate reliance on cross-view contrasts, we propose SIGNA, a novel single-view graph contrastive learning framework. Regarding the inconsistency between structural connection and semantic similarity of neighborhoods, we resort to soft neighborhood awareness for GCL. Specifically, we leverage dropout to obtain structurally-related yet randomly-noised embedding pairs for neighbors, which serve as potential positive samples. At each epoch, the role of partial neighbors is switched from positive to negative, leading to probabilistic neighborhood contrastive learning effect. Moreover, we propose a normalized Jensen-Shannon divergence estimator for a better effect of contrastive learning. Experiments on diverse node-level tasks demonstrate that our simple single-view GCL framework consistently outperforms existing methods by margins of up to 21.74% (*PPI*). In particular, with soft neighborhood awareness, SIGNA can adopt MLPs instead of complicated GCNs as the encoder in transductive learning tasks, thus speeding up its inference process by $109\times$ to $331\times$.

Code — <https://github.com/sunisfighting/SIGNA>

Introduction

Over the past few years, Self-Supervised Learning (SSL) of representations has emerged as a promising and popular research topic since it does not rely on the quantity and quality of labels (Wu et al. 2021; Liu et al. 2022; Xie et al. 2022). As one of the most competitive SSL paradigms, Contrastive Learning (CL) has made its mark in computer vision (He et al. 2020; Chen et al. 2020), natural language processing (Tian, Krishnan, and Isola 2020), and graph domains (Velickovic et al. 2019; Zhu et al. 2020b; Peng et al. 2020b; Sun et al. 2024).

The core spirit of CL is to pull together representations of related instances (positive samples) and push apart representations of unrelated pairs (negative samples). Similar

*Corresponding Author

†Corresponding Author

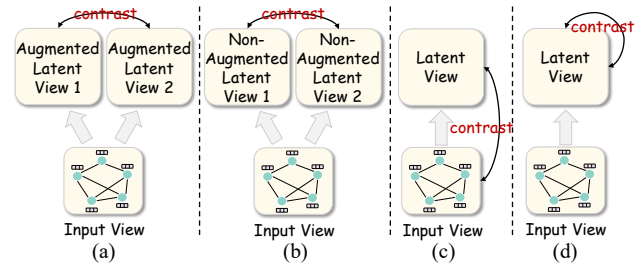


Figure 1: (a) Augmentation based cross-view contrast; (b) Non-augmentation based cross-view contrast; (c) Input-latent cross-view contrast; (d) Single-view contrast (ours).

to CL in computer vision and natural language processing domains, Graph Contrastive Learning (GCL) methods typically rely on cross-view contrasts, *i.e.*, each positive or negative pair consists of two latent embeddings from different views. These strategies include augmentation-based dual-branch methods (Zhu et al. 2020b; You et al. 2020; Shen et al. 2023), non-augmentation-based dual-branch methods (Thakoor et al. 2021; Mo et al. 2022), and input-latent single-branch methods (Peng et al. 2020b), which are summarized in Fig. 1 and discussed in detail in the related work.

Despite the significant progress achieved by cross-view contrastive learning methods on graph data, these approaches face several challenges. Firstly, designing effective augmentation techniques often requires substantial manual effort and fine-tuning, leading to increased implementation complexity. Secondly, cross-view contrast may result in information loss or inconsistencies between views, which can negatively impact model performance. Lastly, the computational cost of cross-view methods is typically high, especially for those dual-branch ones, making them less scalable.

To circumvent these challenges, this paper seeks to explore a simpler yet little-studied alternative: single-view contrastive learning. The most pivotal point is how to obtain intra-view positive and negative samples. Note that nodes in the graph are intricately correlated with each other rather than completely independent (Dang et al. 2021; Wang et al. 2022; Yu, Shi, and Wang 2024; Wang et al. 2023, 2024). Therefore, a natural potential solution within this framework is to conduct contrast according to the topological rela-

relationship between nodes, *i.e.*, pull embeddings of neighbors together and push those of non-neighbors away. However, such a trivial solution may be prone to overfitting topological information and thus results in suboptimal performance. A typical example is GAE (Kipf and Welling 2016b), which aims to reconstruct the graph structure as much as possible but can only achieve less competitive performance than other self-supervised methods that ignore topological information, such as GRACE (Zhu et al. 2020b). In such a context, two questions are naturally raised: (i) *What role should neighborhoods play in node-level GCL?* (ii) *Can neighborhood awareness help single-view contrastive learning surpass cross-view contrastive learning?*

Regarding these two questions, we first investigate the homophily nature across diverse real-world datasets from both global and local perspectives. Statistical results indicate that there is an obvious inconsistency between structural connection and semantic similarity. Inspired by this, we propose **Single-view Graph** contrastive learning with soft Neighborhood Awareness (SIGNA). Instead of relying on augmentations or other cross-view contrastive techniques to generate contrastive pairs, SIGNA adopts soft neighborhood awareness to realize single-view contrast. Specifically, we first use dropout to obtain randomly-perturbed embeddings, which implicitly provides more diverse embedding combinations for contrast. At each epoch, the role of neighbors is allowed to switch from positive to negative rather than being consistently designated as positive, thus leading to probabilistic neighborhood contrastive learning. Furthermore, we propose a normalized Jensen-Shannon divergence (Norm-JSD) estimator, which combines the advantages of both JSD and InfoNCE, resulting in a better contrastive effect.

We evaluate SIGNA on three kinds of node-level tasks, including transductive node classification, inductive node classification, and node clustering. SIGNA consistently outperforms existing methods across all tasks, achieving performance gains of up to 21.74% (PPI). In particular, thanks to reasonable neighborhood awareness, SIGNA is enabled to use a simple Multilayer Perceptron (MLP) as the encoder in transductive learning tasks, which is $109\times$ to $331\times$ faster than a GCN-based encoder with the same settings in terms of the inference time. The distribution analysis and visualization of learned representations showcase the superiority of SIGNA in striking a better balance between intra-class aggregation and inter-class separation, compared with representative neighborhood-aware GCL methods. In summary, our contributions can be highlighted as follows:

- We study the role of neighborhoods by comprehensively investigating their homophily nature on real-world datasets from both global and local perspectives, thus offering nuanced insights on soft neighborhood awareness.
- We propose a simple yet effective single-view GCL framework, SIGNA, which alleviates the heavy reliance on traditional cross-view contrastive learning with the help of our soft neighborhood awareness strategy.
- Extensive experiments demonstrate that SIGNA outperforms existing methods across various node-level tasks, showcasing the feasibility of single-view GCL and the rationality of soft neighborhood awareness.

Related Works

To position our contributions in the literature, we briefly review related works here.

GCL Paradigms. Existing GCL methods mainly rely on cross-view contrast, which can be further grouped into three classes, *i.e.*, augmentation based cross-view contrast, non-augmentation based cross-view contrast, and input-latent cross-view contrast, as illustrated in Fig. 1. (a) *Augmentation based cross-view contrast* is represented by DGI (Velickovic et al. 2019), MVGRL (Hasani and Khasahmadi 2020), GRACE (Zhu et al. 2020b), GraphCL (You et al. 2020), GCC (Qiu et al. 2020), and MERIT (Jin et al. 2021), NCLA (Shen et al. 2023), etc. Augmented views from the same instance are pulled closer while those views from distinct instances are pushed away. (b) Instead of augmenting input data, *non-augmentation based cross-view contrast* resorts to using two discrepant encoders or directly perturbing latent embeddings to obtain contrastive pairs, such as SUGRL (Mo et al. 2022), AFGRL (Lee, Lee, and Park 2022), SimGRACE (Xia et al. 2022), SimGCL (Yu et al. 2022), and COSTA (Zhang et al. 2022b). (c) *Input-latent cross-view contrast* is represented by GMI (Peng et al. 2020b), which aims to maximize mutual information between the input graph and latent embeddings.

In this paper, we aim to study the simpler single-view GCL, which has seldom been studied in the literature.

Neighborhood-aware GCL. Regarding the limitations of augmentation-invariant GCL, concurrent studies turn to neighborhood-aware techniques to retain semantics. For example, GMI (Peng et al. 2020b) maximizes the mutual information across both feature and edge representations between the input and output spaces. GraphCL-NS (Hafidi et al. 2022) makes use of graph structure to sample negatives, *i.e.*, from l -th order neighbors of the anchor node. Based on NT-Xent (Zhu et al. 2020b), NCLA (Shen et al. 2023) adopts a neighbor contrastive loss that regards both intra-view and inter-view neighbors as positives. Without generating augmented views, Local-GCL (Zhang et al. 2022a) fabricates positive samples for each node using first-order neighbors. Instead of binary contrastive justification, GSCL (Ning et al. 2022) uses fine-grained contrastive justification according to the hop distance of neighborhoods. AFGRL (Lee, Lee, and Park 2022) obtains local positives by jointly considering the adjacency matrix and k -NNs in the embedding space. SUGRL (Mo et al. 2022) employs both GCN and MLP encoders to obtain contrastive pairs. For each anchor node, its positive counterparts are constructed with its GCN output and the MLP outputs of its neighbors.

Unlike existing works that fully trust neighborhoods to capture semantic information for downstream tasks, we resort to soft neighborhood awareness to prevent the encoder from overfitting uncertain signals.

Rethinking the Role of Neighborhoods

To implement effective single-view contrastive learning, it is of great significance to understand *what role neighbours should play in graph contrastive learning*. Thus, we first empirically investigate the homophily nature of neighborhoods.




Figure 2: Global and local homophily statistics. *Top*: global homophily ratios on different datasets. *Bottom Left*: the distribution of local homophily counts on Photo. *Bottom Right*: the distribution of local homophily ratios on Photo.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an unweighted graph with a node set \mathcal{V} and an edge set \mathcal{E} . We denote the feature matrix and adjacency matrix by $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{V}|} \in \mathbb{R}^{|\mathcal{V}| \times F}$ and $\mathbf{A} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$, respectively. The one-hot label matrix is denoted by $\mathbf{Y} = \{\mathbf{y}_i\}_{i=1}^{|\mathcal{V}|} \in \mathbb{R}^{|\mathcal{V}| \times c}$. Following (Zhu et al. 2020a), the definition of global homophily ratio is given by:

Definition 0.1 (Global Homophily Ratio). Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, its global homophily ratio is defined as the probability that two connected nodes share the same label:

$$\mathcal{H}_{global} = \frac{1}{|\mathcal{E}|} \sum_{u,v \in \mathcal{V}} \mathbb{I}[(u, v) \in \mathcal{E}] \cdot \mathbb{I}[\mathbf{y}_u = \mathbf{y}_v]. \quad (1)$$

Furthermore, we define local homophily count and local homophily ratio as well:

Definition 0.2 (Local Homophily Count). Given a node $u \in \mathcal{V}$ and its one-hop neighbors \mathcal{N}_u , its local homophily count is defined as the number of neighbors with the same label:

$$\mathcal{H}_{local}^{\#}(u) = \sum_{v \in \mathcal{N}_u} \mathbb{I}[\mathbf{y}_u = \mathbf{y}_v]. \quad (2)$$

Definition 0.3 (Local Homophily Ratio). Given a node $u \in \mathcal{V}$ and its one-hop neighbors \mathcal{N}_u , its local homophily ratio is defined as the probability that its neighbors share the same label with it:

$$\mathcal{H}_{local}(u) = \frac{1}{|\mathcal{N}_u|} \sum_{v \in \mathcal{N}_u} \mathbb{I}[\mathbf{y}_u = \mathbf{y}_v]. \quad (3)$$

We compute global homophily ratios \mathcal{H}_{global} on six real-world datasets, and explore the distributions of local homophily counts $\{\mathcal{H}_{local}^{\#}(u) : u \in \mathcal{V}\}$ and ratios $\{\mathcal{H}_{local}(u) : u \in \mathcal{V}\}$ on Amazon Photo. The statistics are illustrated in Figure 2. As can be observed: (a) the global homophily ratio varies considerably across datasets, with values ranging from 0.3195 (*Flickr*) to 0.9314 (*Physics*), indicating that it is ubiquitous that structural connections between neighbors do not coincide with their semantic relations; and (b) although local homophily is relatively more concentrated in

smaller counts and greater ratios, the overall distribution is hard to estimate, let alone identify exact semantically positive neighbors in the context of unsupervised learning (Sun, Zhang, and Lin 2023). In a nutshell, overemphasizing neighborhood affinity in contrastive learning is risky since the model would be provided with noisy and even detrimental learning signals, which result in suboptimal performance. Motivated by the above findings, we consider *soft neighborhood awareness* to realize single-view GCL.

Methodology

The overall framework of our proposed SIGNA is depicted in Figure 3. SIGNA employs a single-branch paradigm for single-view contrast, eliminating the need for augmented inputs, disparate encoders, or perturbed embeddings to generate multiple views. Simply pulling embeddings of neighbors and pushing away those of non-neighbors has shown to be less competitive (Kipf and Welling 2016b). To pursue soft neighborhood awareness, SIGNA is equipped with three main components, *i.e.*, the encoder with dropout, the stochastic neighbor masking, and the normalized Jensen-Shannon divergence (Norm-JSD) estimator.

Encoder with Dropout. Dropout has been adopted as *minimal data augmentation* in contrastive learning (Gao, Yao, and Chen 2021; Xu et al. 2023). To obtain positive pairs, they pass the same instance to the dropout-contained encoder twice so as to generate two *related yet discrepant* embeddings. By contrast, we only pass the input graph to the encoder with dropouts *ONCE*. In this way, the embedding pairs of neighbors are *structurally related yet randomly noised*, making them competent for positive samples in single-view contrastive learning. Note that we do not rely on augmentations to generate contrastive pairs. The introduction of random noise just implicitly creates more diverse embedding combinations for contrast. Thus, the encoder is less likely to overfit the relationship between node pairs when conducting contrastive learning.

Specifically, the encoder is composed of a stack of L identical layers (in this paper, we fix $L = 2$ to match the common depth of existing works). Each layer has a BaseEncoder which can either be a linear layer (on transductive tasks) or a graph convolutional layer (Kipf and Welling 2016a) (on inductive tasks). Before each BaseEncoder, we employ Dropout (with rescaling) (Srivastava et al. 2014) to inject noise. The output of BaseEncoder is further processed by a nonlinear Activation function and a LayerNorm. Let $\mathbf{H}^{(l-1)}$ be the input embeddings of l -th encoder layer (particularly, $\mathbf{H}^{(0)} = \mathbf{X}$), the forward encoding process within an encoder layer can be formally described as:

$$\mathbf{H}_d^{(l)} = \text{Dropout}(\mathbf{H}^{(l-1)}; p), \quad (4)$$

$$\mathbf{H}_e^{(l)} = \text{Linear}(\mathbf{H}_d^{(l)}) \text{ or } \mathbf{H}_e^{(l)} = \text{GConv}(\mathbf{H}_d^{(l)}, \mathbf{A}), \quad (5)$$

$$\mathbf{H}_a^{(l)} = \text{Activation}(\mathbf{H}_e^{(l)}), \quad (6)$$

$$\mathbf{H}^{(l)} = \mathbf{H}_n^{(l)} = \text{LayerNorm}(\mathbf{H}_a^{(l)}), \quad (7)$$

where p denotes the dropout rate, $\text{Linear}(\mathbf{H}_d^{(l)}) = \mathbf{H}_d^{(l)} \mathbf{W}^{(l)}$ denotes a linear layer, and $\text{GConv}(\mathbf{H}_d^{(l)}, \mathbf{A}) =$

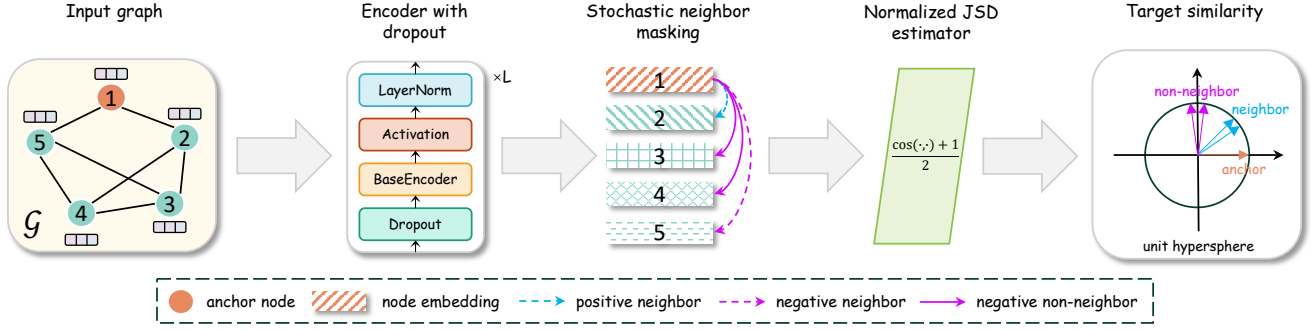


Figure 3: The framework of our proposed SIGNA. The contrast is conducted within a single graph view, and thus SIGNA has only one branch. The encoder with dropout implicitly provides more embedding combinations for robust contrast. The role of neighbors is variable, while non-neighbors are fixed as negative samples. The normalized JSD estimator facilitates better contrastive effect. The goal of SIGNA is to realize soft neighborhood awareness.

$\hat{\mathbf{D}}^{-\frac{1}{2}} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}_d^{(l)} \mathbf{W}^{(l)}$ denotes a graph convolutional layer with $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$ and $\hat{D}_{ii} = \sum_{j=1}^N \hat{A}_{ij}$. We denote the parameter set of the to-be-trained encoder by $\theta = \{\mathbf{W}^{(l)}\}_{l=1, \dots, L}$. The output of the last encoder layer is our learned representations, *i.e.*, $\mathbf{H} = \mathbf{H}^{(L)}$.

Stochastic Neighbor Masking. According to previous homophily analyses, neighbors actually have a non-negligible probability of owning different labels. In light of such uncertainty, we seek to mask a fraction of neighbors and consider the remaining neighbors as positive samples. To do that, we independently draw a masking indicator m_v for a neighboring node v of u from a Bernoulli distribution with probability $\alpha \in [0, 1]$, *i.e.*, $m_v \sim \mathcal{B}(\alpha), \forall v \in \mathcal{N}_u$. Specifically, the remaining neighbor set at each epoch is:

$$\mathcal{N}'_u = \{v \in \mathcal{N}_u \mid m_v = 0\}. \quad (8)$$

Just in case all neighbors are masked, we add the anchor node itself to its positive set for implementation convenience (no extra training signal is provided). As for negative set, it consists of all non-neighbors as well as those masked neighbors. Then, the positive and negative sets are:

$$\mathcal{P}_u = \mathcal{N}'_u \cup \{u\}, \quad \mathcal{Q}_u = \mathcal{V} \setminus \mathcal{P}_u. \quad (9)$$

Theorem 0.4 (Probabilistic Neighborhood Contrastive Learning). *Let S_{uv} be the target similarity between embeddings of the anchor node u and any other node $v \neq u$ within the graph, and assume that $S_{uv} = \delta$ if $v \in \mathcal{P}_u$ otherwise $S_{uv} = \lambda$ ($v \in \mathcal{Q}_u$), where δ, λ are determined by the objective function. Then, we have: (a) $\mathbb{E}_{v \in \mathcal{N}_u}(S_{uv}) = \delta(1 - \alpha) + \lambda\alpha$; (b) $\mathbb{E}_{v \notin \mathcal{N}_u}(S_{uv}) = \lambda$.*

Proof. (a) Since neighbors are randomly masked by a probability of α , we have $p(v \in \mathcal{P}_u \mid v \in \mathcal{N}_u) = 1 - \alpha$ and $p(v \in \mathcal{Q}_u \mid v \in \mathcal{N}_u) = \alpha$. Thus, $\mathbb{E}_{v \in \mathcal{N}_u}(S_{uv}) = \delta(1 - \alpha) + \lambda\alpha$. (b) Since $p(v \in \mathcal{Q}_u \mid v \notin \mathcal{N}_u) = 1$, $\mathbb{E}_{v \notin \mathcal{N}_u}(S_{uv}) = \lambda$ holds. \square

Remark. With stochastic masking, neighbors are flipped back and forth as positive and negative samples along the training process, while non-neighbors consistently act as negative samples. As a result, the expectation of target similarity scores between neighbors lies between δ and λ .

Normalized JSD Estimator. According to the empirical evidence provided by (Hjelm et al. 2018), InfoNCE (Noise Contrastive Estimation (Oord, Li, and Vinyals 2018)) and DV (Donsker-Varadhan representation of the KL-divergence (Donsker and Varadhan 1983)) require a large number of negative samples to be competitive, while JSD (Jensen-Shannon Divergence estimator (Nowozin, Cseke, and Tomioka 2016)) is less sensitive to the number of negative samples. Since we sample positive samples from neighbors, the number of negative samples is inherently reduced. Therefore, we are committed to optimizing an objective in the form of JSD estimator:

$$\mathcal{J}_{\text{JSD}}(u) = \mathbb{E}_{v^+ \sim \mathcal{P}_u} [\log \mathcal{D}_\phi(u, v^+)] + \mathbb{E}_{v^- \sim \mathcal{Q}_u} [\log(1 - \mathcal{D}_\phi(u, v^-))], \quad (10)$$

where $\mathcal{D}_\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a discriminator function modeled by a neural network with parameters ϕ to measure the similarity between two instances and scale it into the range of $[0, 1]$. Typically, the discriminator of JSD estimator is implemented using the inner product plus sigmoid function (Nowozin, Cseke, and Tomioka 2016; Hjelm et al. 2018; Peng et al. 2020b). Yet, it has been empirically verified that ℓ_2 normalization plays an important role in contrastive learning (Chen et al. 2020), which projects embeddings to the unit hypersphere before computing similarity (Wang and Isola 2020). Hence, we introduce a normalized discriminator for the JSD estimator:

$$\mathcal{D}_\phi^{\text{norm}}(u, v) = \frac{1}{2} \left(\frac{\mathbf{z}_u^\top \mathbf{z}_v}{\|\mathbf{z}_u\|_2 \|\mathbf{z}_v\|_2} + 1 \right) = \frac{\cos(\mathbf{z}_u, \mathbf{z}_v) + 1}{2}, \quad (11)$$

where $\mathbf{z} = g_\phi(\mathbf{h}) = \mathbf{W}_g^{(2)} \sigma(\mathbf{W}_g^{(1)} \mathbf{h})$ is a MLP projector parameterized by $\phi = \{\mathbf{W}_g^{(1)}, \mathbf{W}_g^{(2)}\}$. With ℓ_2 normalization, the similarity metric between embedding vectors is calculated by cosine similarity instead of inner product. And we adopt a simple linear scaler to restrict the range of similarity scores. The instantiated estimator with the normalized discriminator is called *Norm-JSD*. In Table 1, we compare Norm-JSD against two typical estimators, JSD and InfoNCE, which implies that Norm-JSD succeeds in combining the advantages of both JSD and InfoNCE while remaining concise.

Perspective	InfoNCE	JSD	Norm-JSD
Robust to #neg.	no	yes	yes
ℓ_2 normalization	yes	no	yes
Scaler	temperature	sigmoid	linear

Table 1: Comparing Norm-JSD against JSD and InfoNCE.

The superiority of Norm-JSD over JSD and InfoNCE is also empirically demonstrated in our ablation studies.

Combining Equation (9), Equation (10), and Equation (11) together, we arrive at the loss function of SIGNA:

$$\ell(u) = -\frac{1}{|\mathcal{P}_u|} \sum_{v^+ \in \mathcal{P}_u} \log(\mathcal{D}_\phi^{\text{norm}}(u, v^+)) - \frac{1}{|\mathcal{Q}_u|} \sum_{v^- \in \mathcal{Q}_u} \log(1 - \mathcal{D}_\phi^{\text{norm}}(u, v^-)), \quad (12)$$

$$\mathcal{L}_{\text{SIGNA}} = \frac{1}{|\mathcal{V}|} \sum_{u \in \mathcal{V}} \ell(u). \quad (13)$$

Note that the target similarity scores for positive and negative samples are $\mathcal{D}_\phi^{\text{norm}}(u, v^+) = 1$ and $\mathcal{D}_\phi^{\text{norm}}(u, v^-) = 0$, respectively. Namely, $\delta = 1$ and $\lambda = 0$. Recalling Theorem 0.4, we have the following corollary:

Corollary 0.5. *With Norm-JSD, the expectation of target similarity for neighbors equals $1 - \alpha$ and the expectation of target similarity for non-neighbors equals 0, i.e., $\mathbb{E}_{v \in \mathcal{N}_u}(S_{uv}) = 1 - \alpha \in [0, 1]$ and $\mathbb{E}_{v \notin \mathcal{N}_u}(S_{uv}) = 0$.*

Remark. The above corollary indicates that, with soft neighborhood awareness, neighbors are expected to maintain a *moderate* level of local tolerance with anchors, which leaves some leeway for future tuning in downstream tasks. By contrast, a large number of non-neighbors will be separated as far as possible. In other words, SIGNA is allowed to learn a desirable globally-uniform yet locally-tolerant embedding space (Wang and Isola 2020; Wang and Liu 2021). Such an effect is observed in our experimental analysis.

Experiments

Experimental Setup

Datasets. We comprehensively evaluate SIGNA on three kinds of node-level tasks across 7 datasets with various scales and properties (Velickovic et al. 2019; Jiao et al. 2020; Thakoor et al. 2021; Lee, Lee, and Park 2022). *Wiki CS*, *Amazon Photo*, *Amazon Computers*, *Coauthor CS*, and *Coauthor Physics* are used for transductive node classification and node clustering tasks. Two larger-scale datasets, *Flickr* and *PPI*, are used for inductive node classification on a single graph and multiple graphs, respectively. Statistics of these datasets are presented in the Appendix.

Baselines. We primarily compare SIGNA against representative and state-of-the-art unsupervised methods for node representation learning, including GAE (Kipf and Welling 2016b), DGI (Velickovic et al. 2019), GMI (Peng et al. 2020b), MVGRL (Hassani and Khasahmadi 2020), GRACE (Zhu et al. 2020b), Subg-Con (Jiao et al. 2020),

S²GRL (Peng et al. 2020a), CCA-SSG (Zhang et al. 2021), BGRL (Thakoor et al. 2021), GraphCL-NS (Hafidi et al. 2022), SUGRL (Mo et al. 2022), and AFGRL (Lee, Lee, and Park 2022).

Evaluation Protocols. We first train all models in a fully unsupervised manner, and then the trained encoder is frozen and used for testing in downstream tasks (see Appendix).

Implementation Details. Details about encoder implementation, hyperparameter selection, and computing infrastructure are provided in appendix due to the space limitation.

Main Results and Analysis

Transductive Node Classification. The empirical performance of various methods across five datasets are summarized in Table 2. As demonstrated by the results, SIGNA consistently and substantially outperforms baseline methods across all five benchmark datasets. Specifically, SIGNA improves the accuracy by absolute margins of 1.55% (*Wiki CS*), 2.04% (*Amazon Photo*), 0.58% (*Amazon Computers*), 1.71% (*Coauthor CS*), and 0.66% (*Coauthor Physics*) over those competitive runner-ups. Besides, a notable finding is that the performances of those competitors are less stable across different benchmarks. For example, AFGRL (Lee, Lee, and Park 2022) outperforms BGRL (Thakoor et al. 2021) on four out of five datasets while being surpassed by BGRL (Thakoor et al. 2021) with a margin of 1.74% on *Wiki CS*. By contrast, our method achieves more robust leading performances, indicating the efficacy and universality of SIGNA. Furthermore, unlike existing methods that rely on graph convolutional layers to generate representations, simple linear layers are used as our BaseEncoders, which discard the complicated message aggregation process on numerous edges. To show the benefit, we investigate the time cost of the inference phase with GCN and MLP being basic encoders (all other settings are kept consistent), which are reported in Table 3. As illustrated, our MLP-based encoder is about $109\times$ to $331\times$ faster than the GCN-based encoder with the same settings in terms of the inference time. In particular, the gap widens rapidly as the size of the edge set increases, since the inference of MLP-based encoders is independent of edge-oriented aggregation.

Inductive Node Classification. The empirical performances on a single graph (*Flickr*) and multiple graphs (*PPI*) are reported in Table 4 and Table 5, respectively. According to the results in terms of micro-averaged F1-score, SIGNA surpasses the previous best method by 3.13% (*Flickr*) and 21.74% (*PPI*), showcasing its remarkable superiority in inductive learning tasks. It is noteworthy that SIGNA even outperforms the best supervised baseline GraphSAGE (Hamilton, Ying, and Leskovec 2017) by 1.83% on *Flickr*. On the *PPI* dataset, while the performance of supervised baselines remains unsurpassed, our method significantly reduces the disparity between unsupervised approaches and those supervised benchmarks.

Node Clustering. The comparison of clustering performance against GRACE (Zhu et al. 2020b), GCA (Zhu et al. 2021), BGRL (Thakoor et al. 2021), and AFGRL (Lee, Lee, and Park 2022) is shown in Table 6. The closer the value of both NMI and Homogeneity is to 1, the better the cluster-

Method	Training Data	Wiki CS	Amazon Photo	Amazon Computers	Coauthor CS	Coauthor Physics
GCN	X, A, Y	77.19±0.12	92.42±0.22	86.51±0.54	93.03±0.31	95.65±0.16
GAT	X, A, Y	77.65±0.11	92.56±0.35	86.93±0.29	92.31±0.24	95.47±0.15
GAE	X, A	75.25±0.28	91.62±0.13	85.27±0.19	90.01±0.17	94.92±0.07
VGAE	X, A	75.63±0.19	92.20±0.11	86.37±0.21	92.11±0.09	94.52±0.00
DGI	X, A	75.35±0.14	91.61±0.22	83.95±0.47	92.15±0.63	94.51±0.52
GMI	X, A	74.85±0.08	90.68±0.17	82.21±0.31	OOM	OOM
MVGRL	X, A	77.52±0.08	92.08±0.01	87.52±0.21	92.18±0.15	95.33±0.03
GRACE	X, A	78.19±0.41	92.24±0.45	86.35±0.44	92.93±0.22	95.26±0.10
CCA-SSG	X, A	78.64±0.72	93.14±0.14	88.74±0.28	92.91±0.20	95.38±0.06
SUGRL	X, A	79.12±0.67	93.07±0.15	88.93±0.21	92.83±0.23	95.38±0.11
BGRL	X, A	79.36±0.53	92.87±0.27	89.68±0.31	93.21±0.18	95.56±0.12
AFGRL	X, A	77.62±0.74	93.22±0.28	89.88±0.33	93.27±0.17	95.69±0.08
SIGNA	X, A	80.91±0.46	95.32±0.19	90.46±0.25	94.98±0.20	96.35±0.09

Table 2: Performance on transductive node classification tasks. The best and second-best performances among unsupervised methods are highlighted in bold and underlined, respectively.

Encoder	Wiki CS	Photo	Comp.	CS	Phys.
GCN	37.67	30.62	60.75	62.85	145.70
MLP	0.25	0.28	0.31	0.35	0.44
Ratio	151×	109×	196×	180×	331×

Table 3: Inference time (millisecond) on transductive node classification tasks with different BaseEncoders.

Method	Training Data	F1-score
FastGCN	X, A, Y	48.1±0.5
GCN	X, A, Y	48.7±0.3
GraphSAGE	X, A, Y	50.1±1.3
Unsup-GraphSAGE	X, A	36.5±1.0
DGI	X, A	42.9±0.1
GMI	X, A	44.5±0.2
Subg-Con	X, A	48.8±0.1
SIGNA	X, A	51.93±0.04

Table 4: Performance on single-graph inductive node classification task (Flickr) in terms of micro-averaged F1-score.

ing effect. As revealed by Table 6, SIGNA exhibits dominating performance over all competitors across five datasets. In particular, SIGNA consistently outperforms AFGRL (Lee, Lee, and Park 2022), a clustering-oriented method, demonstrating the superior capability of SIGNA to learn underlying cluster structure in the unsupervised context.

Why Does SIGNA Work

To facilitate direct understanding of what SIGNA has done, we first analyze the relationship between neighborhood awareness and model performance, which is illustrated in Figure 4. Compared with representative neighborhood-aware (SUGRL) and non-neighborhood-aware (GRACE) GCL methods, SIGNA learns a moderate of structural information due to our soft neighborhood awareness strategies, thus yielding better generalization performances on downstream tasks.

Method	Training Data	F1-score
GaAN-mean	X, A, Y	96.9±0.2
GAT	X, A, Y	97.3±0.2
FastGCN	X, A, Y	63.7±0.6
Unsup-GraphSAGE	X, A	46.5
Random-Init	X, A	62.6±0.2
DGI	X, A	63.8±0.2
GMI	X, A	64.6±0.0
S ² GRL	X, A	66.0±0.0
GRACE	X, A	66.2±0.1
Subg-Con	X, A	66.9±0.2
GraphCL-NS	X, A	65.9±0.0
BGRL-GAT-Encoder	X, A	70.49±0.05
SIGNA	X, A	92.25±0.03

Table 5: Performance on multi-graph inductive node classification task (PPI) in terms of micro-averaged F1-score.

Besides, our histogram analysis reveals that pulling together neighbors is less rational than pushing away non-neighbors, and thus we should avoid overly pulling together neighbors, which exactly matches the spirit of SIGNA. Please refer to our appendix for details.

In Figure 5, we use t-SNE (Van der Maaten and Hinton 2008) to visualize representations learned by SIGNA, as well as those learned by GRACE (Zhu et al. 2020b) and SUGRL (Mo et al. 2022) (representatives of augmentation-invariant and neighbourhood-aware methods). For fairness comparison, their random seeds are kept consistent. As can be seen, GRACE learns a uniformly-distributed embedding space, which impairs intra-class aggregation (e.g., purple samples). On the contrary, SUGRL improves intra-class similarities yet fails to separate different classes (e.g., purple samples and their surroundings). By contrast, SIGNA strikes a better balance between intra-class aggregation and inter-class separation (e.g., purple samples and their surroundings). Clearly, this ability contributes to its superiority in classification and clustering tasks.

Method	Wiki CS		Am. Photo		Am. Computers		Co. CS		Co. Physics	
	NMI	Homo.	NMI	Homo.	NMI	Homo.	NMI	Homo.	NMI	Homo.
Raw features	0.2633	0.2738	0.3273	0.3376	0.2389	0.2617	0.7103	0.7446	0.5207	0.5576
GRACE	<u>0.4282</u>	<u>0.4423</u>	0.6513	0.6657	0.4793	0.5222	0.7562	0.7909	0.5128	0.5546
GCA	0.3373	0.3525	0.6443	0.6575	0.5278	0.5816	0.7620	0.7965	0.5202	0.5654
BGRL	0.3969	0.4156	<u>0.6841</u>	<u>0.7004</u>	0.5364	0.5869	0.7732	0.8041	0.5568	0.6018
AFGRL	0.4132	0.4307	0.6563	0.6743	<u>0.5520</u>	<u>0.6040</u>	<u>0.7859</u>	<u>0.8161</u>	<u>0.5782</u>	<u>0.6174</u>
SIGNA	0.4593	0.4763	0.7635	0.7823	0.5608	0.6057	0.8047	0.8408	0.5907	0.6393

Table 6: Performance on node clustering tasks in terms of NMI and Homogeneity.

Variant	Wiki CS	Am. Photo	Am. Computers	Co. CS	Co. Physics
SIGNA	80.91±0.46	95.32±0.19	90.46±0.25	94.98±0.20	96.35±0.09
×Dropout	80.04±0.55	93.85±0.26	89.72±0.36	93.95±0.15	95.93±0.12
Dropout⇒NFM	75.21±0.45	91.16±0.53	87.43±0.36	93.92±0.20	95.95±0.09
×StochMask	80.58±0.65	94.87±0.28	90.09±0.34	93.74±0.15	96.22±0.08
StochMask⇒AllMask	68.14±0.69	80.59±0.99	76.64±0.60	89.22±0.25	94.30±0.14
Norm-JSD⇒JSD	78.11±0.59	76.90±0.54	70.02±0.77	91.84±0.18	92.42±0.13
Norm-JSD⇒InfoNCE	79.03±0.71	93.64±0.39	88.10±0.35	94.88±0.15	95.96±0.11
×All	77.92±0.48	74.61±1.22	69.16±0.83	86.74±0.34	50.57±0.10

Table 7: Ablation study of SIGNA. ×: remove components; ⇒: replace components with alternatives.




Figure 4: The soft neighborhood awareness of SIGNA yields better performance.




Figure 5: Representation visualization via t-SNE.

Ablation Study

In Table 7, we evaluate the contribution of different components of SIGNA. Overall, soft neighborhood awareness proves of great significance. Our three components are all reasonably designed, as removing any one of them results in a noticeable performance degradation, especially when

all three are removed, the performance drops dramatically. Besides, our specific findings include: (i) when we replace dropout with node feature masking (NFM), a widely-used input augmentation technique (Zhu et al. 2020b, 2021; Jin et al. 2021), the performance decreases dramatically; (ii) our Norm-JSD surpasses two other estimators by large margins, while the performance of JSD lags far behind that of InfoNCE, indicating the importance of our normalized discriminator; (iii) masking all neighbors significantly reduces the performance due to the missing of structural information (recall that we use MLP as encoder for transductive tasks); Yet, overemphasizing neighborhood affinities without stochastic masking also yields an inferior performance. In short, our soft neighborhood awareness is more suitable for GCL.

Conclusion

In this work, we resort to soft neighborhood awareness for single-view GCL. To begin with, we explore homophily statistics on real-world datasets, which reveal that overconfidence in neighborhoods would be risky. Motivated by this, we propose a simple yet effective GCL framework (SIGNA) to pursue a moderate level of neighborhood affinities. Specifically, SIGNA considers neighboring nodes as potential positive samples but is equipped with three nuanced designs to relax restrictions on the retention of structural information. Experimental results demonstrate that SIGNA consistently outperforms existing GCL methods in a wide variety of node-level tasks. This work is expected to serve as a pioneer in exploring reasonable neighborhood awareness for contrastive learning of node representation.

Acknowledgments

The work of Ziyue Qiao was supported by National Natural Science Foundation of China under Grant No. 62406056. The work of Kai Wang was supported in part by National Natural Science Foundation of China under Grant No. U24A20270 and Grant No. 62373378.

References

- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Dang, Z.; Deng, C.; Yang, X.; Wei, K.; and Huang, H. 2021. Nearest neighbor matching for deep clustering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13693–13702.
- Donsker, M. D.; and Varadhan, S. S. 1983. Asymptotic evaluation of certain Markov process expectations for large time. IV. *Communications on Pure and Applied Mathematics*, 36(2): 183–212.
- Gao, T.; Yao, X.; and Chen, D. 2021. Simcse: Simple contrastive learning of sentence embeddings.
- Hafidi, H.; Ghogho, M.; Ciblat, P.; and Swami, A. 2022. Negative sampling strategies for contrastive self-supervised learning of graph representations. *Signal Processing*, 190: 108310.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Hassani, K.; and Khasahmadi, A. H. 2020. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*, 4116–4126. PMLR.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2018. Learning deep representations by mutual information estimation and maximization.
- Jiao, Y.; Xiong, Y.; Zhang, J.; Zhang, Y.; Zhang, T.; and Zhu, Y. 2020. Sub-graph contrast for scalable self-supervised graph representation learning. In *2020 IEEE international conference on data mining (ICDM)*, 222–231. IEEE.
- Jin, M.; Zheng, Y.; Li, Y.-F.; Gong, C.; Zhou, C.; and Pan, S. 2021. Multi-scale contrastive siamese networks for self-supervised graph representation learning.
- Kipf, T. N.; and Welling, M. 2016a. Semi-supervised classification with graph convolutional networks.
- Kipf, T. N.; and Welling, M. 2016b. Variational graph auto-encoders.
- Lee, N.; Lee, J.; and Park, C. 2022. Augmentation-free self-supervised learning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 7372–7380.
- Liu, Y.; Jin, M.; Pan, S.; Zhou, C.; Zheng, Y.; Xia, F.; and Yu, P. 2022. Graph self-supervised learning: A survey. *IEEE Transactions on Knowledge and Data Engineering*.
- Mo, Y.; Peng, L.; Xu, J.; Shi, X.; and Zhu, X. 2022. Simple unsupervised graph representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 7797–7805.
- Ning, Z.; Wang, P.; Wang, P.; Qiao, Z.; Fan, W.; Zhang, D.; Du, Y.; and Zhou, Y. 2022. Graph soft-contrastive learning via neighborhood ranking.
- Nowozin, S.; Cseke, B.; and Tomioka, R. 2016. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 29.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding.
- Peng, Z.; Dong, Y.; Luo, M.; Wu, X.-M.; and Zheng, Q. 2020a. Self-supervised graph representation learning via global context prediction.
- Peng, Z.; Huang, W.; Luo, M.; Zheng, Q.; Rong, Y.; Xu, T.; and Huang, J. 2020b. Graph representation learning via graphical mutual information maximization. In *Proceedings of The Web Conference 2020*, 259–270.
- 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 *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 1150–1160.
- Shen, X.; Sun, D.; Pan, S.; Zhou, X.; and Yang, L. T. 2023. Neighbor Contrastive Learning on Learnable Graph Augmentation.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958.
- Sun, Q.; Wang, K.; Zhang, W.; Cheng, P.; and Lin, X. 2024. Interdependence-Adaptive Mutual Information Maximization for Graph Contrastive Learning. *IEEE Transactions on Knowledge and Data Engineering*.
- Sun, Q.; Zhang, W.; and Lin, X. 2023. Progressive hard negative masking: From global uniformity to local tolerance. *IEEE Transactions on Knowledge and Data Engineering*, 35(12): 12932–12943.
- Thakoor, S.; Tallec, C.; Azar, M. G.; Munos, R.; Veličković, P.; and Valko, M. 2021. Bootstrapped representation learning on graphs. In *ICLR 2021 Workshop on Geometrical and Topological Representation Learning*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive multiview coding. In *European conference on computer vision*, 776–794. Springer.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Veličković, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. *ICLR (Poster)*, 2(3): 4.

- Wang, F.; and Liu, H. 2021. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2495–2504.
- Wang, T.; and Isola, P. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 9929–9939. PMLR.
- Wang, Y.; Yuan, L.; Chen, Z.; Zhang, W.; Lin, X.; and Liu, Q. 2023. Towards efficient shortest path counting on billion-scale graphs. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 2579–2592. IEEE.
- Wang, Y.; Yuan, L.; Zhang, W.; Li, X.; Chen, Z.; and Liu, Q. 2024. Simpler is More: Efficient Top-K Nearest Neighbors Search on Large Road Networks. *Proc. VLDB Endow.*, 17(13): 4683–4695.
- Wang, Y.; Zhang, Y.; Zhang, F.; Lin, M.; Zhang, Y.; Wang, S.; and Sun, X. 2022. Ada-nets: Face clustering via adaptive neighbour discovery in the structure space. *arXiv preprint arXiv:2202.03800*.
- Wu, L.; Lin, H.; Tan, C.; Gao, Z.; and Li, S. Z. 2021. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE Transactions on Knowledge and Data Engineering*.
- 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 *Proceedings of the ACM Web Conference 2022*, 1070–1079.
- Xie, Y.; Xu, Z.; Zhang, J.; Wang, Z.; and Ji, S. 2022. Self-supervised learning of graph neural networks: A unified review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xu, Y.; Wang, Z.; Wang, Z.; Guo, Y.; Fan, R.; Tian, H.; and Wang, X. 2023. SimDCL: dropout-based simple graph contrastive learning for recommendation. *Complex & Intelligent Systems*, 1–13.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33: 5812–5823.
- Yu, C.; Shi, Y.; and Wang, J. 2024. Contextually affinitive neighborhood refinery for deep clustering. *Advances in Neural Information Processing Systems*, 36.
- 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 *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 1294–1303.
- Zhang, H.; Wu, Q.; Wang, Y.; Zhang, S.; Yan, J.; and Yu, P. S. 2022a. Localized Contrastive Learning on Graphs.
- Zhang, H.; Wu, Q.; Yan, J.; Wipf, D.; and Yu, P. S. 2021. From canonical correlation analysis to self-supervised graph neural networks. *Advances in Neural Information Processing Systems*, 34: 76–89.
- Zhang, Y.; Zhu, H.; Song, Z.; Koniusz, P.; and King, I. 2022b. COSTA: covariance-preserving feature augmentation for graph contrastive learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2524–2534.
- Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020a. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems*, 33: 7793–7804.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2020b. Deep graph contrastive representation learning.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2021. Graph contrastive learning with adaptive augmentation. In *Proceedings of the Web Conference 2021*, 2069–2080.