

Node Representation Learning in Graph via Node-to-Neighbourhood Mutual Information Maximization

Wei Dong¹, Junsheng Wu^{2*}, Yi Luo², Zongyuan Ge³, Peng Wang^{4*}

¹School of Computer Science and Engineering, ²School of Software, Northwestern Polytechnical University
³Monash University
⁴University of Wollongong

Abstract

The key towards learning informative node representations in graphs lies in how to gain contextual information from the neighbourhood. In this work, we present a simple-yet-effective self-supervised node representation learning strategy via directly maximizing the mutual information between the hidden representations of nodes and their neighbourhood, which can be theoretically justified by its link to graph smoothing. Following InfoNCE, our framework is optimized via a surrogate contrastive loss, where the positive selection underpins the quality and efficiency of representation learning. To this end, we propose a topology-aware positive sampling strategy, which samples positives from the neighbourhood by considering the structural dependencies between nodes and thus enables positive selection upfront. In the extreme case when only one positive is sampled, we fully avoid expensive neighbourhood aggregation. Our methods achieve promising performance on various node classification datasets. It is also worth mentioning by applying our loss function to MLP based node encoders, our methods can be orders of faster than existing solutions. Our codes and supplementary materials are available at <https://github.com/dongwei156/n2n>.

1. Introduction

Graph-structured data is ubiquitous because almost nothing in the world exists in isolation. As an effective graph modeling tool, Graph Neural Networks (GNNs) have gained increasing popularity in a wide range of domains such as computer vision [6, 26, 29], natural language processing [35], knowledge representation [9], social networks [15], and molecular property prediction [33], just name a few.

In this work, we focus on the node classification task in graphs where the key is to learn informative structure-aware node representations by gaining contextual information from the neighbourhood. This motivates a massive proliferation of message passing techniques in GNNs. Among these methods, a dominant idea follows an AGGREGATION-COMBINE-PREDICTION pipeline, where AGGREGATION step aggregates the neighbouring information into vectorized representations via various neighbourhood aggregators such as *mean* [15], *max* [15], *attention* [31], and *ensemble* [8], which are COMBINED with the node representations via sum or concatenation to realize neighbourhood information fusion. To model multi-hop message passing, the AGGREGATION and COMBINE operations tend to be repeated before the ultimate node representations are obtained for label prediction. In other words, the information exchange in this pipeline is driven by a node classification loss in the PREDICTION phase. Under the umbrella of supervised learning, this line of methods unifies node representation learning and classification. A potential problem is that they may suffer from the scalability issue due to the expensive labeling cost for large-scale graph data.

As a remedy to expensive human annotation, Self-Supervised Learning (SSL) has shown proven success in computer vision [4, 5] and natural language processing [2, 12], which, however, is less sufficiently explored in graph modeling. The key challenge thereof lies in how to design suitable *pretext* task from non-Euclidean graph-structured data to learn informative node representations. Inheriting the idea from computer vision, many recent attempts approach graph-based SSL by designing topological graph augmentations to generate multi-view graphs for contrastive loss. Off-the-shelf GNNs [20, 37] tend to be used as default options for node/graph encoding.

In this work, we propose a simple-yet-effective SSL alternative to learn node representations by adopting aggregation-free Multi-layer Perceptron (MLP) as node encoder and directly maximizing the mutual information be-

*Wei Dong's email is dw156@mail.nwpu.edu.cn. Corresponding authors: Peng Wang (pengw@uow.edu.au) and Junsheng Wu (wujunsheng@nwpu.edu.cn)

tween the hidden representations of nodes and their neighbourhood. Intuitively, by aligning the node representation to the neighbourhood representation, we can distill useful contextual information from the surrounding, analogous to knowledge distillation [7, 23]. Theoretically, the proposed Node-to-Neighbourhood (N2N) mutual information maximization essentially encourages graph smoothing based on a quantifiable graph smoothness metric. Following InfoNCE [22], the mutual information can be optimized by a surrogate contrastive loss, where the key boils down to positive sample definition and selection.

To further improve the efficiency and scalability of our N2N network as well as ensuring the quality of selected positives, we propose a Topology-Aware Positive Sampling (TAPS) strategy, which samples positives for a node from the neighbourhood by considering the structural dependencies between nodes. This enables us to select the positives upfront. In the extreme case when only one positive is used for contrastive learning, we can avoid the time-consuming neighbourhood aggregation step, but still achieve promising node classification performance.

We conduct experiments on six graph-based node classification datasets and the results show the promising potential of the proposed node representation strategy. The contributions of this work can be summarized as follows:

- We propose a simple-yet-effective MLP-based self-supervised node representation learning strategy with the idea of maximizing the mutual information between nodes and surrounding neighbourhood. We reveal our N2N mutual information maximization strategy essentially encourages graph smoothing by resorting to a quantifiable smoothness metric.
- The mutual information maximization problem is optimized by a surrogate contrastive loss. A scalable TAPS strategy is proposed which enables the positives to be selected upfront. In the case when only one positive is considered, we avoid expensive neighbourhood aggregation but still obtain satisfactory node classification performance.
- Experiments on six graph-based node classification datasets show our methods not only achieve competitive performance but have other appealing properties as well. For example, our methods can be orders of faster than existing solutions.

2. Related Work

In this section, we briefly review existing work on (1) Graph Neural Networks (GNNs) and (2) Graph Contrastive Learning (GCL).

Graph Neural Networks. GNNs target to learn structure-aware node/graph representations based on the

topology structure of the graph data [36]. A main focus thereof is to design effective message passing strategies to encourage information propagation in the graph. Early attempts [27] learn node representations by following a recurrent fashion, where the node states are updated by applying a propagation function iteratively until equilibrium is reached. Inspired by convolutional neural networks that were proposed originally for grid-like topology, such as images, convolution-like propagation was introduced into graph data [10]. As a prevailing variant of this line of work, Graph Convolutional Network (GCN) [20] stacks a set of 1-hop spectral filters [10] and nonlinear activation function to learn node representations. GCN ignites a wave of following-up work with the aim of improving the efficiency or effectiveness of information exchange in graphs. Simplifying Graph Convolutional Network (SGC) [34] reduces the excessive complexity of GCN by removing the nonlinear activation function to obtain collapsed aggregation matrix. L2-GCN [39] simplifies the training complexity of GCN by using an efficient layer-wise training strategy. GraphSAGE-Mean [15] applies *mean* aggregator to fixed number of randomly sampled neighbours to reduce the computational cost and adopts concatenation to merge the node and neighbourhood information. FastGCN [3] also adopts sampling strategy to reduce computational footprint of GCN by interpreting graph convolutions as integral transforms of embedding functions. Another line of work aims to design sophisticated neighbourhood aggregation strategies. Graph Attention Network (GAT) [31] stacks masked self-attention layers in order that the nodes can adaptively attend their neighbours. The use of multiple aggregators (*mean, sum, max, normalized mean, scaling*, etc) is also proven to be able to benefit the expressive power of GNNs either empirically [11] or theoretically [8]. Parallely, efforts are also dedicated towards theoretical understanding about the properties of various GNN models. For example, in [37], the authors present a theoretical framework for analyzing the representational power of GNNs based on Weisfeiler-Lehman (WL) graph isomorphism test. In [18], CS-GNN is designed to understand and improve the use of graph information in GNNs by leveraging feature and label based smoothness metrics.

The aforementioned GNNs are under the umbrella of supervised learning, where a predominant training pipeline is AGGREGATION-COMBINE-PREDICTION, which unifies node representation learning and node classification in an end-to-end fashion. The supervised GNNs show proven success in various applications by virtue of sufficient labeled data.

Graph Contrastive Learning. With the aim of learning high-quality data representations without human annotation, SSL has shown very promising performance in computer vision and natural language processing. Motivated

by such success, marrying SSL and GNNs is also gaining increasing attention, where the key boils down to designing effective *pretext* tasks for graph-structured data. In this work, we focus mainly on Graph Contrastive Learning (GCL) as it is directly related to the proposed method.

One key in contrastive learning is how to define positive samples as it directly determines contrasting mode. The work in [38] directly performs graph-level contrastive learning, where two views of the same graph are obtained by performing augmentations such as dropping nodes, edge perturbation and attribute masking. The motivation behind is that the semantics in graphs are invariant to minor topological change. Another line of work is inspired by Deep InfoMax in computer vision [17], where the image representation is learned by maximizing the mutual information between the learned representation and the input image. By applying this mutual information maximization idea in graphs, several variants are derived. Deep Graph InfoMax (DGI) [32] aligns the global representation of a graph to its hidden node representation by contrasting against the node representations derived from a corrupted graph. To avoid the readout function and corrupted operation in [32], Graphical Mutual Information (GMI) [24] aligns the output node representation to the input sub-graph. The work in [16] learns node and graph representation by maximizing mutual information between node representations of one view and graph representations of another view obtained by graph diffusion. InfoGraph [30] works by taking graph representation and patch representation as pairs and determine whether they are from the same graph. Although our method also adopts mutual information maximization, our method differs from the aforementioned methods from several perspectives. Firstly, we avoid the operations such as graph augmentation, iterative aggregation and readout function by aligning nodes to their neighbourhood in the output layer of a MLP encoder. Secondly, instead of designing the *pretext* task heuristically, our framework can be theoretically justified by its connection to graph smoothing. Finally, a topology-aware sampling strategy is proposed in our work that converts our framework into highly efficient node-to-node contrastive learning but still maintains promising node classification performance.

3. Methodology

In this section, we firstly introduce notations, symbols, and necessary background about GNN models. We then present our idea of N2N mutual information maximization and its link to graph smoothing. Finally, we elaborate the proposed TAPS strategy.

3.1. GNN Framework

We denote a graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ with a set of nodes \mathcal{V} , a set of edges \mathcal{E} , and an adjacency matrix \mathbf{A} . Each

node $v \in \mathcal{V}$ has a feature vector $\vec{v} \in \mathcal{X}$ with node feature space \mathcal{X} of dimensionality D . GNNs utilize a neighbourhood aggregation scheme to learn latent node representation $\vec{h}^{(l)} \in \mathbb{R}^{D^{(l)}}$ in l -th layer for each node v , and a prediction function is applied to the node representations of final hidden layer to predict the class label y_v of each node v . Based on such notations, a commonly adopted pipeline for supervised GNNs can be defined as:

$$\begin{aligned} \vec{s}_i^{(l-1)} &= \text{AGGREGATION}(\{\vec{h}_j^{(l-1)} : v_j \in \mathcal{N}_i\}), \\ \vec{h}_i^{(l)} &= \text{COMBINE}(\{\vec{s}_i^{(l-1)}, \vec{h}_i^{(l-1)}\}), \\ \mathcal{L}_{\text{CE}} &= \text{PREDICTION}(\{\vec{h}_i^{(L)}, y_{v_i}\}), \end{aligned} \quad (1)$$

where the AGGREGATION function can be any form of aggregators such as *mean*, *max*, *sum*, *attention*, and *ensemble* that learn the neighbourhood representation $\vec{s}_i^{(l-1)}$ from the set $\{\vec{h}_j^{(l-1)} : v_j \in \mathcal{N}_i\}$ based on the neighbourhood \mathcal{N}_i , and the COMBINE function updates $\vec{h}_i^{(l-1)}$ to a new representation $\vec{h}_i^{(l)}$ in l -th layer by combining $\vec{s}_i^{(l-1)}$ with $\vec{h}_i^{(l-1)}$. A L -layer GNN iterates the above two operations L times and the PREDICTION function is applied in the output layer for node classification. A de-facto loss function for the PREDICTION layer is Cross-Entropy (CE) loss.

3.2. Node-to-Neighbourhood (N2N) Mutual Information Maximization

We learn topology-aware node representation by maximizing the mutual information between the hidden representations of nodes and their neighbourhood, which is partially motivated by knowledge distillation [23] in computer vision. In this section, we firstly present the definition of the N2N mutual information, which is followed by the optimization of the mutual information and its link to graph smoothing.

We denote the Probability Density Function (PDF) of the node representation $\vec{h}_i^{(l)}$ over the feature space $\mathcal{X}^{D^{(l)}}$ in $[0, 1]^{D^{(l)}}$ as $p(H(\mathbf{x})^{(l)})$, where $\mathbf{x} \in \mathcal{X}^{D^{(l)}}$ and $H(\cdot)^{(l)}$ is a mapping function from \mathbf{x} to $\vec{h}_i^{(l)}$; the PDF of the neighbourhood representation $\vec{s}_i^{(l)}$ as $p(S(\mathbf{x})^{(l)})$ with the mapping function $S(\cdot)^{(l)}$ from \mathbf{x} to $\vec{s}_i^{(l)}$; and the joint PDF between node and neighbourhood is $p(S(\mathbf{x})^{(l)}, H(\mathbf{x})^{(l)})$. We define the mutual information between the node representations and their corresponding neighbourhood representation as:

$$\begin{aligned} I(S(\mathbf{x})^{(l)}; H(\mathbf{x})^{(l)}) &= \\ & \int_{\mathcal{X}^{D^{(l)}}} p(S(\mathbf{x})^{(l)}, H(\mathbf{x})^{(l)}) \cdot \log \frac{p(S(\mathbf{x})^{(l)}, H(\mathbf{x})^{(l)})}{p(S(\mathbf{x})^{(l)}) \cdot p(H(\mathbf{x})^{(l)})} d\mathbf{x}. \end{aligned} \quad (2)$$

This operation encourages each node representation to distill the contextual information presented in its neighbour-

hood representation. However, mutual information is notoriously difficult to compute, particularly in continuous and high-dimensional space. Fortunately, scalable estimation enabling efficient computation of mutual information is made possible through Mutual Information Neural Estimation (MINE) [1], which converts mutual information maximization into minimizing the InfoNCE loss [22]. The surrogate InfoNCE loss function of the N2N mutual information in Eq. (2) is defined as:

$$\mathcal{L}_{\text{InfoNCE}} = -\mathbb{E}_{v_i \in \mathcal{V}} \left[\log \frac{\exp(\text{sim}(\vec{s}_i^{(l)}, \vec{h}_i^{(l)})/\tau)}{\sum_{v_k \in \mathcal{V}} \exp(\text{sim}(\vec{h}_k^{(l)}, \vec{h}_i^{(l)})/\tau)} \right], \quad (3)$$

which estimates the mutual information via node sampling, where the $\text{sim}(\cdot, \cdot)$ function denotes the cosine similarity, the $\exp(\cdot)$ function implies the exponential function, and τ is the temperature parameter. The positive pair is $(\vec{s}_i^{(l)}, \vec{h}_i^{(l)})$ and the negative pair is $(\vec{h}_k^{(l)}, \vec{h}_i^{(l)})_{i \neq k}$.

In essence, maximizing $I(S(\mathbf{x})^{(l)}; H(\mathbf{x})^{(l)})$ can play the role of graph smoothing, which has proven to be able to benefit node/graph prediction. To elaborately prove this point, we resort to a feature smoothness metric in [18]:

$$\delta_f^{(l)} = \frac{\|\sum_{v_i \in \mathcal{V}} (\sum_{v_j \in \mathcal{N}_i} (\vec{h}_i^{(l)} - \vec{h}_j^{(l)}))^2\|_1}{|\mathcal{E}| \cdot D^{(l)}}, \quad (4)$$

where $\|\cdot\|_1$ is the Manhattan norm. The work [18] further proposes that the information gain from the neighbourhood representation $\vec{s}_i^{(l)}$ is defined as the Kullback-Leibler divergence:

$$D_{KL}(S(\mathbf{x})^{(l)} \| H(\mathbf{x})^{(l)}) = \int_{\mathcal{X}^{D^{(l)}}} p(S(\mathbf{x})^{(l)}) \cdot \log \frac{p(S(\mathbf{x})^{(l)})}{p(H(\mathbf{x})^{(l)})} d\mathbf{x}, \quad (5)$$

which is positively correlated to the feature smoothness metric $\delta_f^{(l)}$, i.e., $D_{KL}(S(\mathbf{x})^{(l)} \| H(\mathbf{x})^{(l)}) \sim \delta_f^{(l)}$. This standpoint implies that a large feature smoothness value $\delta_f^{(l)}$ means significant disagreement between node representations $\{\vec{h}_i^{(l)}\}$ and their corresponding neighbourhood representations $\{\vec{s}_i^{(l)}\}$. This inspires the following theorem (Ref. Appendix A for proof):

Theorem 1 For a graph \mathcal{G} with the set of features $\mathcal{X}^{D^{(l)}}$ in space $[0, 1]^{D^{(l)}}$, the information gain represented by $D_{KL}(S(\mathbf{x})^{(l)} \| H(\mathbf{x})^{(l)})$ is negatively correlated to the mutual information $I(S(\mathbf{x})^{(l)}; H(\mathbf{x})^{(l)})$ and thus maximizing $I(S(\mathbf{x})^{(l)}; H(\mathbf{x})^{(l)})$ essentially minimizes $D_{KL}(S(\mathbf{x})^{(l)} \| H(\mathbf{x})^{(l)})$ and $\delta_f^{(l)}$, which attains the goal of graph smoothing:

$$I(S(\mathbf{x})^{(l)}; H(\mathbf{x})^{(l)}) \sim \frac{1}{D_{KL}(S(\mathbf{x})^{(l)} \| H(\mathbf{x})^{(l)})} \sim \frac{1}{\delta_f^{(l)}}. \quad (6)$$

3.3. Topology-Aware Positive Sampling (TAPS)

Up to now, we obtain the neighbourhood representation $\vec{s}_i^{(l)}$ of a node by applying the AGGREGATION function to all neighbours of the node. This solution may suffer from two issues. Firstly, the whole neighbourhood can include redundant or even noisy information. Secondly, the aggregation operation is computationally expensive. To address these two problems, we propose a TAPS strategy for self-supervised node representation learning. The basic idea is that we measure the topological dependencies between a node and its neighbours and sample positives of the node based on the ranked dependency values.

For a node v_i , we use a variable X_i to represent its topological information. X_i can take the value of either \mathcal{N}_i or $\bar{\mathcal{N}}_i = \mathcal{V} - \mathcal{N}_i$, where the former corresponds to the neighbourhood information and the latter is the contextual information complementary to the neighbourhood. Based on the definition of X_i , we define $p(X_i = \mathcal{N}_i) = \frac{|\mathcal{N}_i|}{|\mathcal{V}|}$ and probability $p(X_i = \bar{\mathcal{N}}_i) = \frac{|\mathcal{V} - \mathcal{N}_i|}{|\mathcal{V}|}$, where $|\cdot|$ is the cardinality function. Basically $p(X_i = \mathcal{N}_i)$ indicates when we sample a node randomly on the graph, the probability that the node will fall into the neighbourhood of v_i . Furthermore, for two neighbouring nodes v_i and v_j , we can define the following joint probabilities:

$$\begin{aligned} p(X_i = \mathcal{N}_i, X_j = \mathcal{N}_j) &= \frac{|\mathcal{N}_i \cap \mathcal{N}_j|}{|\mathcal{V}|}, \\ p(X_i = \mathcal{N}_i, X_j = \bar{\mathcal{N}}_j) &= \frac{|\mathcal{N}_i \cap (\mathcal{V} - \mathcal{N}_j)|}{|\mathcal{V}|}, \\ p(X_i = \bar{\mathcal{N}}_i, X_j = \mathcal{N}_j) &= \frac{|(\mathcal{V} - \mathcal{N}_i) \cap \mathcal{N}_j|}{|\mathcal{V}|}, \\ p(X_i = \bar{\mathcal{N}}_i, X_j = \bar{\mathcal{N}}_j) &= \frac{|(\mathcal{V} - \mathcal{N}_i) \cap (\mathcal{V} - \mathcal{N}_j)|}{|\mathcal{V}|}, \end{aligned} \quad (7)$$

where $p(X_i = \mathcal{N}_i, X_j = \mathcal{N}_j)$ is the probability that the randomly selected node will fall into the intersected neighbours of v_i and v_j . Motivated by mutual information, we define the graph-structural dependency between v_i and v_j as:

Definition 1 Graph-structural dependency between neighbouring node v_i and v_j is defined as:

$$I(X_i; X_j) = \sum_{X_i} \sum_{X_j} p(X_i, X_j) \cdot \log \frac{p(X_i, X_j)}{p(X_i) \cdot p(X_j)}, \quad (8)$$

s.t. $v_j \in \mathcal{N}_i$.

The graph-structural dependency value above basically measures the topological similarity of two nodes. A large value suggests the strong dependency between two nodes.

In TAPS strategy, we select positives of v_i by ranking the dependency values between the neighbouring nodes and v_i and then obtain the neighbourhood representation $\vec{s}_i^{(l)}$ by

applying aggregator, *e.g.*, *mean*, to the selected positives. In particular, when only one positive is selected, we directly select the node v_j with maximum dependency value to v_i and thus avoid the expensive aggregation operation. Meanwhile, because the topology structure of a graph relies on the adjacency matrix only, TAPS allows us to perform positive sampling upfront, which can avoid the positive sampling overhead during training.

4. Training Frameworks

According to the relationship among GNN encoders, *pretext* tasks, and *downstream* tasks, there are three types of graph-based self-supervised training schemes [21]. The first type is Pre-training and Fine-tuning (PT&FT). The pre-training stage firstly initializes the parameters of the GNN encoder with *pretext* tasks. After this, this pre-trained GNN encoder is fine-tuned under the supervision of specific *downstream* tasks. The second is Joint Learning (JL) scheme, where the GNN encoder, *pretext* and *downstream* tasks are trained jointly. The last type is Unsupervised Representation Learning (URL). Akin to PT&FT, URL also follows a two-stage training scheme, where the first stage trains the GNN encoder based on the *pretext* task but in the second *downstream* task stage, the GNN encoder is frozen. In our work, we adopt both JL and URL pipelines to train and evaluate our network.

JL Training Framework. Fig. 1 (a) illustrates our JL training pipeline. As can be seen, unlike most existing graph-based SSL work that uses GNN as node/graph encoder, we simply use a shallow MLP as encoder, which is more efficient. In JL scheme, we apply InfoNCE loss and Cross-Entropy loss jointly on top of the node representations obtained as output of the MLP encoder:

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_{CE} + \alpha\mathcal{L}_{\text{InfoNCE}}, \quad (9)$$

where α is a trade-off parameter used to balance the two loss functions.

URL Training Framework. The URL framework, as shown in Fig. 1 (b), involves two training stages: the pre-training *pretext* task trains the MLP encoder using InfoNCE loss $\mathcal{L}_{\text{InfoNCE}}$, and the *downstream* task learns linear node classifiers using Cross-Entropy loss \mathcal{L}_{CE} .

5. Experiments

In this section, we start by introducing the experimental setups including datasets, comparing methods, and the implementation details. Then we present the performance comparison to existing methods. Finally, ablation studies are presented to reveal other appealing properties of the proposed methods.

5.1. Experimental Setup

Datasets. We conduct experiments on six real-world node classification datasets: Cora [39], Pubmed [39], Citeseer [39], Amazon Photo [28], Coauthor CS [28], and Coauthor Physics [28]. The first three datasets are constructed as citation networks¹, Amazon Photo is for the Amazon co-purchase graph, and Coauthor CS & Physics are for co-authorship graphs. All of the above six graphs are connected and undirected.

Our Models. We denote N2N-TAPS- x as our model sampling top- x positive neighbours based on TAPS, *e.g.*, N2N-TAPS-1 samples the one neighbour with maximum dependency value as positive. We evaluate our methods with 1–5 positives to inspect how the positive size influences the node classification performance. We use N2N-random-1 as our baseline where one positive is sampled randomly from the neighbourhood of a node. By default, we aggregate all the neighbours of a node using *mean* aggregator as positive, which is denoted as N2N.

Existing Methods. We employ three types of methods for comparison: GNN methods, conventional SSL methods, and GCL methods. For each type, some representatives are chosen. The GNN solutions include GCN [20], SGC [34], L2-GCN [39], GraphSAGE-Mean [15], FastGCN [3], GAT [31], SplineCNN [13], PNA [8], and CS-GCN [18]. For conventional SSL methods, we incorporate DeepWalk [25] and node2vec [14] for comparison. The state-of-the-art GCL methods used for comparison include DGI [32], GMI [24], MVGRL [16], InfoGraph [30], and Graph-MLP [19]. The key properties of such SSL/GCL methods and our methods are summarized in Table 1 from three perspectives: learning type, training framework, and encoder type.

Table 1. Summaries of typical SSL/GCL methods and our method.

Model	Learning Type	Framework	Encoder
DeepWalk [25]	SSL	URL	Shallow MLP
node2vec [14]		URL	Shallow MLP
DGI [32]	GCL	URL	GCN
GMI [24]		URL	GCN
MVGRL [16]		URL	GCN
InfoGraph [30]		URL	GIN
Graph-MLP [19]		JL	Shallow MLP
N2N-TAPS- x		JL & URL	Shallow MLP

Implementation Details. For fair comparison, we follow the common practice to fix the number of hidden layers in our method and the compared GNNs and GCL encoders to be 2. For all datasets, we set the dimensionality of the hidden layer to be 512. Some other important hyper-parameters include: dropout ratio is 0.6 for Cora, Citeseer, and Coauthor CS, 0.2 for Pubmed, 0.4 for Amazon Photo, and 0.5 for Coauthor Physics; L2-regularization

¹For these three datasets, we use the train/val/test splits in [18, 39].

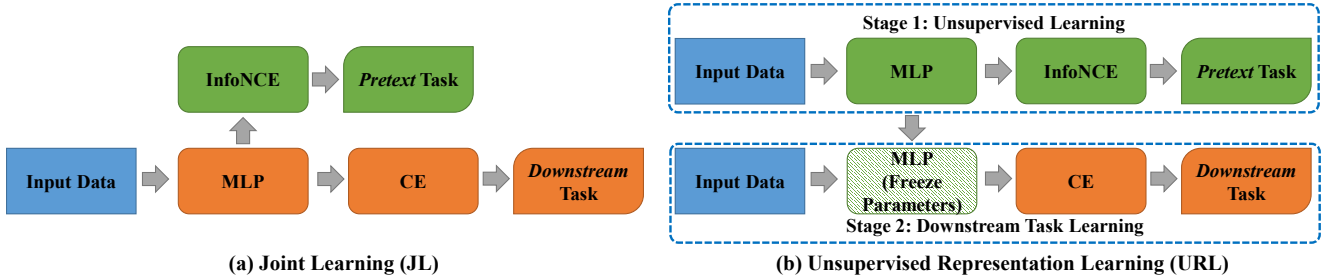


Figure 1. Illustration of two training pipelines adopted for the proposed model.

is 0.01 for Cora, Citeseer, and Coauthor CS, 0.001 for Pubmed and Amazon Photo, and 0.05 for Coauthor Physics; training epochs are 1000 for all datasets; learning rate is 0.01 for Pubmed and 0.001 for the other datasets; nonlinear activation is ReLU function. For N2N-TAPS- x (JL), α is set to 0.9 for Cora, Pubmed, and Citeseer, 0.99 for Amazon Photo and Coauthor CS, and 0.999 for Coauthor Physics; temperature τ is 5 for Cora, Pubmed, Citeseer, and Amazon Photo, 100 for Coauthor CS, and 30 for Coauthor Physics. The temperature τ of N2N-TAPS- x (URL) is 5 for all datasets. These hyper-parameters are determined via cross-validation. We implement our models by Tensorflow 2.4. All of the experiments are performed on a machine with Intel 8-core i7-10870H CPU (2.20GHz), 32GB CPU memory, and one GeForce RTX 3080 Laptop card (16GB GPU memory). For each dataset, we run all models five times, and the mean and standard deviation of the micro-f1 score are used as the evaluation metric.

5.2. Overall Results

Table 2 shows the performance comparison between our method and other selected methods. From the results we have the following observations: (1). Our N2N models, either N2N(JL) or N2N(URL), consistently outperform the comparing methods on all the six datasets. The margins can go up to 3% on datasets like Cora, Pubmed, and Coauthor CS. This shows the competitiveness of our N2N mutual information maximization strategy over GNNs and other GCL based solutions for node representation learning. Another issue worth mentioning is that since our N2N based methods avoid topological augmentations and simply use MLP as node encoders, our methods are more efficient in terms of training and inference. (2). Within the N2N family, we generally observe improvement when sampling more positives based on TAPS but the improvements are marginal. This demonstrates the potential of N2N-TAPS-1 as it avoids neighbourhood aggregation operation which is known to be expensive. However, when the single positive is sampled randomly from the neighbourhood, the performance drops significantly. This result shows the proposed TAPS strategy indeed can sample topologically meaningful positives. (3). Within the existing methods, the GCL solu-

tions have comparable performance or even slightly better performance comparing to the fully supervised GNN variants. This observation shows SSL can be a promising alternative in graph-based representation learning.

5.3. Ablation Studies

In this section, we conduct additional ablation studies to reveal other appealing properties of the proposed methods.

N2N (JL) based on Random Positive Sampling. To further justify the necessity and advantage of our TAPS strategy, we run experiments on random positive sampling by varying the sampling size from 1 to 5. We choose two datasets, *i.e.*, Amazon Photo and Coauthor Physics, for this experiment because their average node degree > 5 . For each sampling size, we run the experiments three times with different random seeds. The results are shown in Table 3. From the table we can clearly observe that random positive sampling results in large performance variance which means random sampling fails to identify consistent and informative neighbours.

Time Consumption. Our methods are expected to be more efficient comparing to existing work. On one hand, our work adopts MLP as node encoder and thus avoids the expensive node aggregation in the encoding phase. On the other hand, TAPS enables us to sample limited high-quality positives upfront. Especially, when one positive is selected, we fully get rid of the aggregation operation.

Table 4 shows the time consumption comparison. From the results we can see our methods can be orders of faster than the typical GNN and GCL based methods. GraphMLP [19] also adopts MLP as encoder but it aligns a node to all the nodes that can be reached from this node. This explains its slowness on large datasets such as CS and Physics.

Evaluation of TAPS Strategy. TAPS is an important component in our framework to ensure the quality and efficiency of positive sampling. In Table 2 we have shown the advantage of TAPS over random sampling on our N2N-TAPS-1 model. In this section, we apply TAPS sampling to another sampling based GNN baseline, GraphSAGE-Mean, to verify if TAPS can be used as a general neighbourhood sampling strategy to identify informative neighbours. Fig. 2 shows the results. By default, GraphSAGE-

Table 2. Performance comparison to existing methods on six datasets. The results for the comparing methods are obtained either by running the publicly released code or through our own implementation. Mean and standard deviation of 5-fold Micro-f1 scores are reported as evaluation metric. The best results for each dataset are highlighted in **bold**.

Model	Cora	Pubmed	Citeseer	Amazon Photo	Coauthor CS	Coauthor Physics
GCN [20]	84.72±0.08	87.02±0.06	78.20±0.15	90.02±0.10	90.52±0.21	91.04±0.06
SGC [34]	84.25±0.10	86.68±0.06	77.65±0.12	89.36±0.16	90.03±0.15	90.12±0.08
L2-GCN [39]	84.56±0.03	86.80±0.06	77.06±0.08	89.16±0.26	90.52±0.04	91.15±0.10
GraphSAGE-Mean [15]	85.04±0.12	87.15±0.14	77.82±0.15	90.05±0.04	90.40±0.08	90.89±0.12
FastGCN [3]	84.08±0.04	86.92±0.08	77.65±0.05	88.65±0.12	90.00±0.05	89.60±0.25
GAT [31]	85.23±0.15	87.85±0.14	78.05±0.26	86.78±0.26	91.10±0.10	91.17±0.15
SplineCNN [13]	85.45±0.16	87.82±0.08	78.83±0.13	89.08±0.18	91.13±0.20	90.82±0.16
PNA [8]	85.40±0.12	87.20±0.07	78.28±0.05	90.23±0.14	91.35±0.16	90.68±0.14
CS-GCN [18]	85.14±0.04	87.75±0.08	78.85±0.16	90.12±0.14	91.14±0.08	90.23±0.08
DeepWalk [25]	77.84±0.12	86.52±0.12	60.24±0.28	83.95±0.28	84.75±0.28	86.25±0.22
node2vec [14]	75.15±0.06	85.20±0.02	65.52±0.18	84.16±0.18	85.28±0.18	85.58±0.18
DGI [32]	85.08±0.05	87.03±0.08	78.82±0.15	90.06±0.17	90.85±0.08	89.88±0.20
GMI [24]	85.26±0.08	87.26±0.16	78.69±0.16	89.25±0.10	90.80±0.20	90.05±0.08
MVGRL [16]	85.38±0.06	87.25±0.12	78.08±0.06	88.23±0.10	90.62±0.08	89.68±0.18
InfoGraph [30]	84.32±0.08	87.56±0.12	78.85±0.12	90.10±0.20	90.42±0.08	90.18±0.12
Graph-MLP [19]	82.50±0.10	87.25±0.13	78.86±0.08	89.25±0.14	90.25±0.20	89.45±0.10
N2N-Random-1 (JL)	83.46±0.18	86.20±0.08	76.85±0.30	86.25±0.15	89.65±0.06	88.75±0.12
N2N-TAPS-1 (JL)	85.46±0.08	91.08±0.16	80.24±0.13	89.90±0.08	92.07±0.06	90.70±0.20
N2N-TAPS-2 (JL)	86.36±0.16	90.76±0.14	80.34±0.06	89.77±0.07	92.11±0.08	90.28±0.24
N2N-TAPS-3 (JL)	86.74±0.14	90.74±0.05	80.64±0.14	89.71±0.10	92.21±0.07	90.81±0.17
N2N-TAPS-4 (JL)	86.52±0.15	90.78±0.07	81.06±0.11	89.89±0.08	92.27±0.07	90.42±0.22
N2N-TAPS-5 (JL)	87.10±0.08	90.78±0.20	80.84±0.11	89.99±0.06	92.37±0.05	91.38±0.12
N2N (JL)	87.52±0.20	90.92±0.09	80.90±0.21	90.12±0.26	93.06±0.07	91.41±0.09
N2N-Random-1 (URL)	82.50±0.10	84.08±0.12	75.06±0.12	86.07±0.12	88.23±0.08	86.56±0.04
N2N-TAPS-1 (URL)	84.66±0.11	88.42±0.07	77.44±0.07	89.40±0.05	91.72±0.03	90.35±0.06
N2N-TAPS-2 (URL)	85.60±0.08	89.26±0.08	78.56±0.08	90.49±0.04	91.56±0.03	90.68±0.02
N2N-TAPS-3 (URL)	87.96±0.08	89.24±0.05	78.36±0.04	90.61±0.05	91.53±0.08	91.20±0.08
N2N-TAPS-4 (URL)	88.04±0.11	89.32±0.09	78.54±0.07	90.35±0.07	92.03±0.06	90.89±0.05
N2N-TAPS-5 (URL)	87.84±0.09	89.88±0.05	79.08±0.07	90.65±0.07	91.99±0.04	91.52±0.02
N2N (URL)	88.20±0.05	89.30±0.04	79.54±0.02	91.08±0.11	92.31±0.05	91.77±0.08

Table 3. Performance of N2N (JL) model with random positive sampling on Amazon Photo and Coauthor Physics. We vary the positive sampling size from 1 to 5 and for each sampling size we run the experiments three times with different random seeds.

Model	Amazon Photo			Coauthor Physics		
	1	2	3	1	2	3
Random Seeds						
N2N-Random-1 (JL)	86.25±0.15	85.52±0.14	84.08±0.20	88.75±0.12	87.25±0.10	86.20±0.28
N2N-Random-2 (JL)	87.68±0.28	85.74±0.18	85.62±0.38	88.62±0.24	87.28±0.14	86.02±0.10
N2N-Random-3 (JL)	87.95±0.10	86.28±0.08	85.20±0.16	89.24±0.16	87.64±0.10	86.52±0.10
N2N-Random-4 (JL)	88.25±0.26	86.06±0.20	86.21±0.42	88.82±0.10	86.52±0.20	87.30±0.10
N2N-Random-5 (JL)	88.30±0.12	85.52±0.14	86.56±0.30	89.65±0.20	87.82±0.15	87.22±0.14

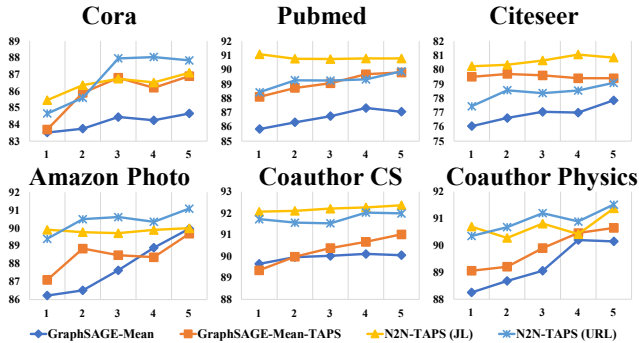


Figure 2. The performance comparison among GraphSAGE-Mean, GraphSAGE-Mean-TAPS, and N2N-TAPS- x (JL and URL) based on varying number of sampled neighbours. Random neighbour sampling is used for GraphSAGE-Mean and TAPS is used for GraphSAGE-Mean-TAPS.

Mean [15] uses random sampling to select neighbours for aggregation, which has the risk of absorbing noisy information. We replace random sampling in GraphSAGE-Mean with TAPS and leave all the other implementation intact. Its performance is obviously boosted and generally using more neighbours can benefit the performance more. This

Table 4. The time cost comparison among typical GNN/GCL methods and our N2N-TAPS- x models. The number indicates the time consumption for training a method on a dataset for one epoch. The symbol † behind N2N-TAPS- x models indicates inference time. The unit of time is milliseconds (ms).

Model	Cora	Pubmed	Citeseer	Amazon Photo	Coauthor CS	Coauthor Physics
GCN [20]	1.68	5.88	6.64	6.08	10.06	11.31
GraphSAGE-Mean [15]	1.56	6.52	6.46	6.28	9.65	12.06
FastGCN [3]	1.21	5.73	6.89	5.24	9.06	10.04
DGI [32]	1.82	7.20	8.56	8.25	12.00	15.68
GMI [24]	2.08	7.93	9.62	10.18	15.54	21.16
MVGRL [16]	1.95	7.64	9.06	9.30	13.11	16.92
InfoGraph [30]	1.86	8.65	11.16	11.42	12.82	15.20
Graph-MLP [19]	0.45	1.08	0.56	0.50	1.82	3.64
N2N-TAPS-1 (JL)	0.09	0.12	0.11	0.12	0.22	0.24
N2N-TAPS-5 (JL)	0.15	0.14	0.18	0.19	0.46	0.63
N2N (JL)	0.30	0.19	0.19	0.20	1.19	1.68
N2N-TAPS-1 (JL)†	0.05	0.07	0.13	0.08	0.06	0.07
N2N-TAPS-5 (JL)†	0.06	0.09	0.16	0.09	0.08	0.10
N2N (JL)†	0.08	0.11	0.19	0.09	0.12	0.13
N2N-TAPS-1 (URL)	0.10	0.15	0.13	0.16	0.30	0.28
N2N-TAPS-5 (URL)	0.18	0.18	0.16	0.25	0.42	0.54
N2N (URL)	0.22	0.25	0.18	0.38	1.25	1.80
N2N-TAPS-1 (URL)†	0.06	0.04	0.15	0.09	0.08	0.10
N2N-TAPS-5 (URL)†	0.05	0.06	0.16	0.10	0.10	0.13
N2N (URL)†	0.06	0.06	0.16	0.14	0.13	0.15

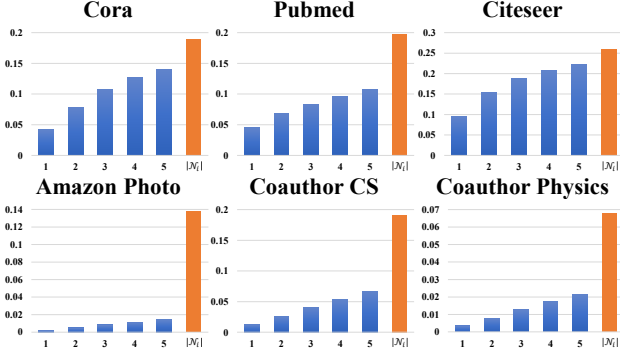


Figure 3. Label smoothness values obtained by using TAPS strategy to sample 1 to 5 neighbours (blue bars) and the label smoothness value obtained by considering all neighbours without any sampling strategy (orange bar).




Figure 4. Subgraph statistics on Cora by using TAPS strategy to perform the subgraph partition. The horizontal axis indicates the number of nodes in the subgraph; the vertical axis implies the number of such subgraphs.

observation shows us again that it is important to consider the structural dependencies to select useful neighbours to enrich node representations.

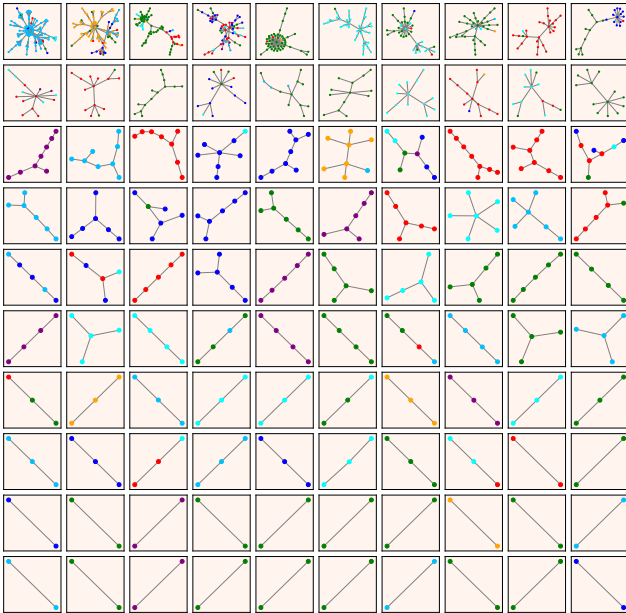


Figure 5. Visualization of part of the subgraphs derived from TAPS on Cora. We rank the sizes of the subgraphs and divide them into 10 intervals. For each row, we visualize the top-10 subgraphs within the corresponding interval from top (top 10% interval) to bottom (80% ~ 90% interval). Different node colors represent different labels.

Label Smoothness Analysis. To verify the quality of the neighbourhood sampling by using TAPS strategy, we introduce the label smoothness metric proposed in CS-GNN [18]: $\delta_l = \sum_{(v_i, v_j) \in \mathcal{E}} (1 - \mathbb{I}(v_i \simeq v_j)) / |\mathcal{E}|$, where $\mathbb{I}(\cdot)$ is an indicator function that has $\mathbb{I}(v_i \simeq v_j) = 1$ if $y_{v_i} = y_{v_j}$. Otherwise, $\mathbb{I}(v_i \simeq v_j) = 0$ when $y_{v_i} \neq y_{v_j}$. By virtue of the label smoothness, a large δ_l suggests that the nodes with different labels are regarded to be connected neighbours, while a small δ_l indicates a graph \mathcal{G} possessing the higher-quality neighbourhood structure, *i.e.*, most neighbours of a node have the same label as the node. Such high-quality neighbourhood can contribute homogeneous information gain to their corresponding central nodes [18].

Fig. 3 shows that the label smoothness values gradually increases by expanding the sampling size from 1 to 5 with our TAPS strategy. Without any sampling strategy, the label smoothness value of the whole graph is the highest. This phenomenon suggests that our TAPS strategy can promote the neighbourhood sampling quality as the sampling size decreases, explaining why the proposed N2N-TAPS-1 model has competitive performance on some datasets.

TAPS strategy is essentially a subgraph partition scheme. A good partition is expected to lead to subgraphs faithful to the node labels. Fig. 4 shows the statistical distribution in terms of subgraph size (nodes in a subgraph) and the number of such subgraphs derived from TAPS. The details of the subgraph partition on Cora are visualized in Fig. 5, where different node colors represent different labels. In each of the subgraph, most of the nodes have the same color (same label), even in some large subgraphs, implying TPAS generates high-quality neighbourhood. This visualization also reveals that our TAPS strategy is able to model multi-hop contextual information in graph although we do not explicitly do so. The details of the statistical distribution and the subgraph partition for other datasets can be found in Appendix B and C.

6. Conclusions

This work presented a simple-yet-effective self-supervised node representation learning strategy by directly optimizing the alignment between hidden representations of nodes and their neighbourhood through mutual information maximization. Theoretically, our formulation encourages graph smoothing. We also proposed a TAPS strategy to identify informative neighbours and improve the efficiency of our framework. It is worth mentioning when only one positive is selected, our model can fully avoid neighbourhood aggregation but still maintain promising node classification performance. An interesting future work will be extending the proposed self-supervised node representation learning and neighbourhood sampling strategy to heterogeneous graph data.

References

- [1] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual information neural estimation. In *International Conference on Machine Learning*, pages 531–540. PMLR, 2018. 4
- [2] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 1
- [3] Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: Fast learning with graph convolutional networks via importance sampling. In *International Conference on Learning Representations*, 2018. 2, 5, 7
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1
- [5] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 1
- [6] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-label image recognition with graph convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1
- [7] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4794–4802, 2019. 2
- [8] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 5, 7
- [9] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering*, 31(5):833–852, 2018. 1
- [10] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 3844–3852, 2016. 2
- [11] Nima Dehmamy, Albert-Laszlo Barabasi, and Rose Yu. Understanding the representation power of graph neural networks in learning graph topology. *Advances in Neural Information Processing Systems*, 32:15413–15423, 2019. 2
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019. 1
- [13] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *CVPR*, pages 869–877, 2018. 5, 7
- [14] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016. 5, 7
- [15] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017. 1, 2, 5, 7
- [16] Kaveh Hassani and Amir Hoseini Khasahmadi. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*, pages 4116–4126. PMLR, 2020. 3, 5, 7
- [17] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2018. 3
- [18] Yifan Hou, Jian Zhang, James Cheng, Kaili Ma, Richard TB Ma, Hongzhi Chen, and Ming-Chang Yang. Measuring and improving the use of graph information in graph neural networks. In *International Conference on Learning Representations*, 2019. 2, 4, 5, 7, 8, 11
- [19] Yang Hu, Haoxuan You, Zhecan Wang, Zhicheng Wang, Erjin Zhou, and Yue Gao. Graph-mlp: Node classification without message passing in graph. *arXiv preprint arXiv:2106.04051*, 2021. 5, 6, 7
- [20] Thomas Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. 1, 2, 5, 7
- [21] Yixin Liu, Shirui Pan, Ming Jin, Chuan Zhou, Feng Xia, and Philip S Yu. Graph self-supervised learning: A survey. *arXiv preprint arXiv:2103.00111*, 2021. 5
- [22] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 4
- [23] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019. 2, 3
- [24] Zhen Peng, Wenbing Huang, Minnan Luo, Qinghua Zheng, Yu Rong, Tingyang Xu, and Junzhou Huang. Graph representation learning via graphical mutual information maximization. In *Proceedings of The Web Conference 2020*, pages 259–270, 2020. 3, 5, 7
- [25] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014. 5, 7
- [26] Xiaojuan Qi, Renjie Liao, Jiaya Jia, Sanja Fidler, and Raquel Urtasun. 3d graph neural networks for rgb-d semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5199–5208, 2017. 1
- [27] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural

- network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008. [2](#)
- [28] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018. [5](#)
- [29] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with directed graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7912–7921, 2019. [1](#)
- [30] Fan-Yun Sun, Jordon Hoffman, Vikas Verma, and Jian Tang. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. In *International Conference on Learning Representations*, 2020. [3](#), [5](#), [7](#)
- [31] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2019. [1](#), [2](#), [5](#), [7](#)
- [32] Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph info-max. In *International Conference on Learning Representations*, 2018. [3](#), [5](#), [7](#)
- [33] Oliver Wieder, Stefan Kohlbacher, Méline Kuenemann, Arthur Garon, Pierre Ducrot, Thomas Seidel, and Thierry Langer. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, 2020. [1](#)
- [34] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International Conference on Machine Learning*, pages 6861–6871, 2019. [2](#), [5](#), [7](#)
- [35] Lingfei Wu, Yu Chen, Kai Shen, Xiaojie Guo, Hanning Gao, Shucheng Li, Jian Pei, and Bo Long. Graph neural networks for natural language processing: A survey. *arXiv preprint arXiv:2106.06090*, 2021. [1](#)
- [36] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020. [2](#)
- [37] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks. In *International Conference on Learning Representations*, 2019. [1](#), [2](#)
- [38] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33:5812–5823, 2020. [3](#)
- [39] Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. L2-gcn: Layer-wise and learned efficient training of graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2127–2135, 2020. [2](#), [5](#), [7](#)