

Conditional Distribution Learning for Graph Classification

Jie Chen^{1, 5}, Hua Mao², Chuanbin Liu^{3*}, Zhu Wang^{4*}, Xi Peng^{1, 5}

¹College of Computer Science, Sichuan University, Chengdu, China

²Department of Computer and Information Sciences, Northumbria University, Newcastle, U.K.

³School of Economics and Management, China University of Petroleum (Beijing), Beijing, China

⁴Law School, Sichuan University, Chengdu, China

⁵National Key Laboratory of Fundamental Algorithms and Models for Engineering Numerical Simulation, Sichuan University, Chengdu, China

{chenjie2010, wangzhu}@scu.edu.cn, hua.mao@northumbria.ac.uk, liuchuanbingl@163.com, pengx.gm@gmail.com

Abstract

Leveraging the diversity and quantity of data provided by various graph-structured data augmentations while preserving intrinsic semantic information is challenging. Additionally, successive layers in graph neural network (GNN) tend to produce more similar node embeddings, while graph contrastive learning aims to increase the dissimilarity between negative pairs of node embeddings. This inevitably results in a conflict between the message-passing mechanism (MPM) of GNNs and the contrastive learning (CL) of negative pairs via intraviews. In this paper, we propose a conditional distribution learning (CDL) method that learns graph representations from graph-structured data for semisupervised graph classification. Specifically, we present an end-to-end graph representation learning model to align the conditional distributions of weakly and strongly augmented features over the original features. This alignment enables the CDL model to effectively preserve intrinsic semantic information when both weak and strong augmentations are applied to graph-structured data. To avoid the conflict between the MPM and the CL of negative pairs, positive pairs of node representations are retained for measuring the similarity between the original features and the corresponding weakly augmented features. Extensive experiments with several benchmark graph datasets demonstrate the effectiveness of the proposed CDL method.

Introduction

Graph-structured data typically represent non-Euclidean data structures, which are often visualized as graphs with labels (Kazemi et al. 2020). The analysis of such data has attracted considerable attention in various applications, e.g., social network recommendation (Fan et al. 2022), molecule classification (Wang et al. 2022), and traffic flow analysis (Li and Zhu 2021). In graph classification tasks, graph-structured data consists of multiple graphs, each with its own independent topology. In practice, obtaining large-scale labelled annotations for graph-structured data is often prohibitively costly or even impossible. Semisupervised graph classification aims to predict the labels of graphs via a small number of labeled graphs and many unlabeled graphs.

Numerous graph neural networks (GNNs) have been proposed for the analysis of graph-structured data, e.g., the

graph convolutional network (GCN) (Kipf and Welling 2017), graph attention network (GAT) (Veličković et al. 2018) and their variants. Through the message-passing mechanism, GNNs aggregate semantic information from neighboring nodes in the graph topology to central nodes. In graph classification tasks, GNNs aim to transform individual graphs into corresponding low-dimensional vector representations at the graph level (Yue et al. 2022). Consequently, GNNs are capable of integrating graph-structured data into comprehensive graph representations, preserving the intrinsic structure of graph-structured data (Vignac, Loukas, and Frossard 2020). Furthermore, data augmentation for graph-structured data has been extensively investigated to enhance the generalizability of GNN models.

Inspired by the previous success of visual data augmentation (Wang and Qi 2023; Chen et al. 2023, 2020b; Oord, Li, and Vinyals 2018), graph contrastive learning (GCL) leverages the advantages of graph data augmentation to maximize the similarity between positive pairs and the dissimilarity between negative pairs (Wu et al. 2023). For example, You et al. (You et al. 2020) proposed a GCL framework that employs four types of augmentations to learn invariant graph representations across different augmented views, including node dropping, edge perturbation, attribute masking and subgraph extraction. Zhu et al. (Zhu et al. 2021) proposed a graph contrastive representation learning method with adaptive augmentation that incorporates various priors for topological and semantic information of the graph. These GCL-based approaches effectively enhance the generalizability of graph representation learning models.

GCL for graph classification has raised significant interests recently (Ju et al. 2024; Han et al. 2022; Yue et al. 2022). For example, Han et al. (Han et al. 2022) adopted \mathcal{G} -Mixup to augment graphs for graph classification. Yue et al. (Yue et al. 2022) proposed a graph label-invariant augmentation (GLIA) method that adds perturbations to the node embeddings of graphs for creating augmented graphs. However, these approaches have overlooked two limitations in GCL. First, different augmentations for graph-structured data often result in varying performances for GNNs (Yue et al. 2022). For example, heavy edge perturbation often disrupts the intrinsic semantic information contained in the graph-structured data. Therefore, a significant challenge is

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

how to leverage the diversity and quantity provided by various graph-structured data augmentations while preserving the intrinsic semantic information. Second, successive GNN layers tend to make node representations more similar due to the message-passing mechanism (Chen et al. 2020a). In contrast, GCL aims to enhance the dissimilarity of negative pairs of node embeddings through mutual information maximization (Yue et al. 2022). This inevitably results in a potential conflict between message-passing mechanism of GNNs and the contrastive learning of negative pairs in intraviews.

In this paper, we propose a conditional distribution learning (CDL) method that learns graph representations from graph-structured data for semisupervised graph classification. We define slight perturbations to node attributes or graph topology as weak augmentations, whereas significant perturbations are referred to as strong augmentations. We first present an end-to-end graph representation learning model to align the conditional distributions of weakly and strongly augmented features over the original features. Aligning the conditional distributions effectively reduces the risk of disrupting intrinsic semantic information when strong augmentation is applied to the graph-structured data. Moreover, only positive pairs of node representations are retained for measuring the similarity between the original features and their corresponding weakly augmented features. From the perspective of the mutual information loss, we explain how the conflict between the message-passing mechanism of GNNs and the contrastive learning of negative pairs of node representations can be avoided. Additionally, we introduce a semisupervised learning scheme that includes a pretraining stage and a fine-tuning stage for semisupervised graph classification. Extensive experimental results with benchmark graph datasets demonstrate that the proposed CDL method is highly competitive with state-of-the-art graph classification methods.

Our major contributions are summarized as follows.

- We present an end-to-end graph representation learning model that takes advantage of both weak and strong augmentations of graph-structured data to learn graph representations for semisupervised graph classification.
- Conditional distribution learning is introduced to characterize the consistent conditional distributions of weakly and strongly augmented node embeddings over the original node embeddings.
- The similarity loss function is introduced to alleviate the potential conflict between the message-passing mechanism of the GNN and the contrastive learning of negative pairs in GCL.

Preliminaries

Problem Formulation

Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, \mathcal{V} and \mathcal{E} represent the sets of nodes and edges, respectively. Let $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ denote the feature matrix of nodes, where d is the feature dimension, and n is the number of nodes. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be the adjacency matrix of \mathcal{G} . If there exists an edge from node i to node j , then $A_{ij} = 1$; otherwise, $A_{ij} = 0$. The

degree matrix \mathbf{D} is defined as $\mathbf{D} = \text{diag}[d_1, d_2, \dots, d_n] \in \mathbb{R}^{n \times n}$, and its diagonal elements are $d_i = \sum_{v_j \in \mathcal{V}} A_{ij}$.

Given an undirected graph \mathcal{G} , let $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ denote the adjacency matrix of \mathcal{G} with added self-loops. The diagonal degree matrix $\tilde{\mathbf{D}}$ is defined as $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$. Given a set of graphs $\mathcal{G}_s = \{(\mathcal{G}_i, y_i)\}_{i=1}^{n_g}$ with $\mathbf{X}_i \in \mathbb{R}^{d \times n_i}$ and c categories, the purpose of graph classification is to learn a function f that can map each graph to its corresponding label, where n_g represents the number of graphs, n_i denotes the number of nodes in \mathcal{G}_i , and y_i is the label of \mathcal{G}_i .

Graph Classification Techniques

GNNs have emerged as effective methods for graph classification (Yue et al. 2022; Han et al. 2022). Without loss of generalizability, the hidden representation of node j of a graph \mathcal{G}_i is iteratively updated by aggregating information from its neighbors $\mathcal{N}(j)$ at the l -th layer, i.e.,

$$\begin{aligned} \mathbf{z}_j^{(l)} &= f_a(\mathbf{h}_k^{(l-1)} : k \in \mathcal{N}(j)) \\ \mathbf{h}_j^{(l)} &= f_c(\mathbf{h}_j^{(l-1)}, \mathbf{z}_j^{(l)}) \end{aligned} \quad (1)$$

where $f_a(\cdot)$ is an aggregation function, $\mathcal{N}(i)$ represents the neighbors of node i , and $f_c(\cdot)$ combines the hidden representations of the neighbors and itself at the previous layer. For graph classification, GNNs ultimately apply a readout function to aggregate the hidden representations of nodes into a graph-level representation at the final layer, i.e.,

$$\mathbf{h}_{\mathcal{G}_i} = f_{\text{readout}}(\mathbf{h}_j^{(L)} : j \in \mathcal{V}(\mathcal{G}_i)) \quad (2)$$

where L denotes the index of the final layer (Feng et al. 2022). The graph-level representation $\mathbf{h}_{\mathcal{G}_i}$ can be employed to predict the label of \mathcal{G}_i , i.e.,

$$\mathbf{q}_i = \text{softmax}(\mathbf{h}_{\mathcal{G}_i}) \quad (3)$$

where $\text{softmax}(\cdot)$ is applied to transform an input vector into a probability distribution for graph classification tasks.

Related Work

GCL has recently attracted significant interest. For example, the GCL framework exhibits a pretraining strategy of GNNs through maximizing the agreement between two augmented views of graphs (You et al. 2020). Given n graphs, $2n$ augmented graphs can be generated by applying a data augmentation technique, such as edge perturbation or attribute masking. The normalized temperature-scaled cross entropy loss (NT-Xent) (Oord, Li, and Vinyals 2018; Sohn 2016) is adopted to maximize the consistency between positive pairs $\mathbf{u}_i, \mathbf{v}_i$ compared with negative pairs, i.e.,

$$\mathcal{L}_i = -\log \frac{f(\mathbf{u}_i, \mathbf{v}_i)}{f(\mathbf{u}_i, \mathbf{v}_i) + \sum_{k \neq i} f(\mathbf{u}_i, \mathbf{v}_k) + \sum_{k \neq i} f(\mathbf{u}_i, \mathbf{u}_k)} \quad (4)$$

where $f(\cdot) = \exp(\text{sim}(\cdot)/\tau)$, $\text{sim}(\mathbf{u}_i, \mathbf{v}_j) = \frac{\mathbf{u}_i \mathbf{v}_j}{\|\mathbf{u}_i\| \|\mathbf{v}_j\|}$ is the cosine similarity and τ is a temperature parameter. In

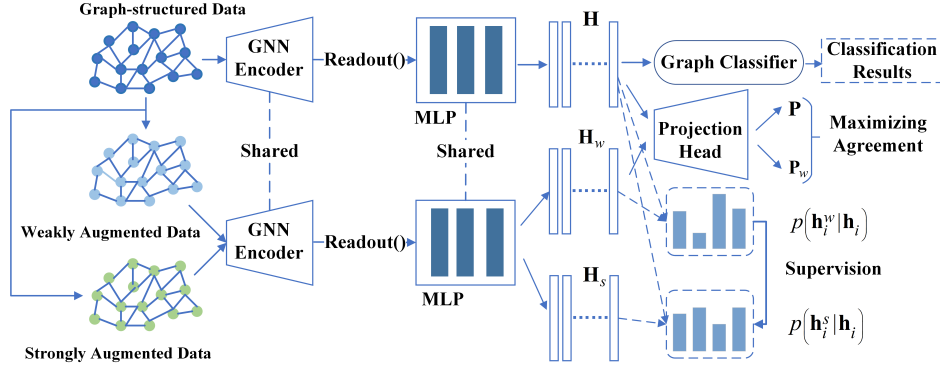


Figure 1: Framework of the CDL model. The graph-level representations, \mathbf{H} , \mathbf{H}_s and \mathbf{H}_w , are produced by a shared GNN encoder module using the graph-structured data, a strongly augmented view and a weakly augmented view, respectively. $p(\mathbf{h}_i^s | \mathbf{h}_i)$ and $p(\mathbf{h}_i^w | \mathbf{h}_i)$ represent the conditional distributions given the raw graph-level representation of the i th node. \mathbf{P} and \mathbf{P}_w denote the projected representations of the raw graph-structured data and a weakly augmented view, respectively.

GCL, a conflict arises in the message-passing mechanism of GNNs when combined with contrastive learning of negative pairs via intraviews. This is because certain samples simultaneously contribute to the gradients of both positive and negative pairs during message passing (Ji et al. 2024).

Recent studies have been proposed to investigate semisupervised graph classification (Ju et al. 2024; Yue et al. 2022; Zeng and Xie 2021; Chen, Bian, and Sun 2019). These methods typically leverage a small number of labeled graphs and massive numbers of available unlabeled graphs to explore the semantic information of graph-structured data. Most existing GCL-based methods for graph classification tasks inherently employ contrastive learning of both positive and negative pairs (Wang et al. 2024; Yue et al. 2022). This inevitably leads to a conflict between the message-passing mechanism of GNNs and the contrastive learning of negative pairs via intraviews. Additionally, data augmentation often disrupt the intrinsic semantic information contained in the graph-structured data (Yue et al. 2022), which makes enhancing the generalizability and robustness of the graph classification models challenging.

Conditional Distribution Learning on Multiple Graphs

Framework Overview

Fig. 1 shows the framework of the CDL model, which consists of three modules: a shared GNN encoder module, a projection head module and a conditional distribution construction module. The shared GNN encoder module learns graph-level representations, \mathbf{H} , \mathbf{H}_s and \mathbf{H}_w , from the raw graph-structured data, a strongly augmented view and a weakly augmented view, respectively. These graph-level representations are obtained from the node embeddings of the graphs through a graph pooling layer. In particular, \mathbf{H} is used for semisupervised graph classification tasks. The projection head module is designed to produce the projected representations of the raw graph-structured data augmentation and a weakly augmented view for GCL. The conditional distri-

bution construction module constructs the conditional distributions of strongly and weakly augmented node embeddings given the original node embeddings.

Conditional Distribution Learning

We formally introduce the details of the proposed CDL method, including a semisupervised learning scheme for graph classification and the CDL model. In particular, the semisupervised training scheme for graph classification consists of two stages: a pretraining stage and a fine-tuning stage.

Pretraining Stage When many unlabeled graphs are available, an intuitively designed pretraining framework is an effective approach for semisupervised graph classification. The pretraining framework of the CDL model includes a shared GNN encoder module of graph-structured data and a projection head module. The shared GNN encoder module consists of three components, including a GNN encoder component, a pooling layer component and a multilayer perceptron (MLP) component. In the shared GNN encoder module, we perform graph representation learning according to the graph classification techniques.

For simplicity, a GCN (You et al. 2020; Kipf and Welling 2017) is used as an example to describe the backbone of the GNN encoder module. The formula for the l th GCN layer in the GNN encoder is defined as follows:

$$\mathbf{H}^{(l)} = \sigma \left(\tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2} \mathbf{H}^{(l-1)} \mathbf{W}^{(l)} \right) \quad (5)$$

where $\mathbf{W}^{(l)}$ is a weight matrix of the l th layer, $\mathbf{H}^{(l)}$ represents a node-level embedding matrix with $\mathbf{H}^{(0)} = \mathbf{X}$, and σ is a nonlinear activation function, e.g., $\text{ReLU}(\cdot) = \max(0, \cdot)$. Then, we employ a global sum pooling layer as a READOUT function to obtain graph-level representations from the node-level embeddings, i.e., $\mathbf{H}_G = \text{pooling}(\mathbf{H}^{(L)})$. Finally, the MLP component is composed of several fully connected layers, with computations defined as follows:

$$\mathbf{H} = \text{MLP} \left(\mathbf{H}_G, \mathbf{W}^{(l)} \right) \quad (6)$$

where a nonlinear activation function, $\text{ReLU}(\cdot)$, is applied for each linear layer. To improve the generalizability of the CDL model, we apply a batch normalization function in individual layers (Ioffe and Szegedy 2015).

During the pretraining stage, the graph-level representations, \mathbf{H} and \mathbf{H}_w , are produced by the shared GNN encoder module via the raw graph-structured data and a weakly augmented view, respectively. We further employ the projection head module (Chen et al. 2020b), which consists of a two-layer MLP, to obtain the projections of \mathbf{H} and \mathbf{H}_w for contrastive learning, i.e.,

$$\begin{aligned} \mathbf{P} &= \mathbf{W}^{(2)} \sigma \left(\mathbf{W}^{(1)} \mathbf{H} \right), \\ \mathbf{P}_w &= \mathbf{W}^{(2)} \sigma \left(\mathbf{W}^{(1)} \mathbf{H}_w \right). \end{aligned} \quad (7)$$

The pretraining framework of the CDL model aims to ensure the consistency of the intrinsic semantic information between a weakly augmented view and the original graph-structured data. Therefore, we further investigate the loss function employed during the pretraining stage.

Conditional Distribution Learning When dealing with relatively small amounts of graph-structured data, GNN-based models often learn redundant features due to the message-passing and aggregation mechanisms. Graph-structured data augmentation aims to increase the diversity and quantity of graph-structured data, which can improve the generalizability and robustness of GCL-based models. To effectively learn essential features from graph-structured data, GCL-based models apply strong augmentation to graph-structured data. However, strong augmentations, such as edge perturbation or attribute masking, often significantly perturb the original graph-structured data, potentially altering the intrinsic semantic information contained in the graph-structured data and thus resulting in limited generalizability of GCL-based models.

Inspired by recent advances of contrastive learning techniques in computer vision (Wang and Qi 2023), we utilize a conditional distribution to enforce the consistency of intrinsic semantic information between an augmented view and the original graph-structured data. To minimize the risk of altering intrinsic semantic information with strong augmentation, we introduce a conditional distribution learning strategy to guide graph representation learning. This learning strategy explores the supervision information embedded in the original graph-structured data for weak and strong augmentations. We consider the distribution of relative similarity between a weakly augmented node embedding and the original node embedding. Specifically, the conditional distribution of \mathbf{h}_i^w given \mathbf{h}_i is defined as follows:

$$\begin{aligned} p(\mathbf{h}_i^w | \mathbf{h}_i) &= \frac{p(\mathbf{h}_i^w, \mathbf{h}_i)}{p(\mathbf{h}_i)} \\ &= \frac{\exp(\text{sim}(\mathbf{h}_i^w, \mathbf{h}_i)/\tau)}{\exp(\text{sim}(\mathbf{h}_i^w, \mathbf{h}_i)/\tau) + \sum_{k=1}^K \exp(\text{sim}(\mathbf{h}_k^w, \mathbf{h}_i)/\tau)} \end{aligned} \quad (8)$$

where K denotes the number of negative samples chosen from a weakly augmented view, and \mathbf{h}_i and \mathbf{h}_i^w represent the corresponding original node embedding and weak node embedding, respectively. Additionally, the other conditional distributions of \mathbf{h}_i^s given \mathbf{h}_i , i.e., $p(\mathbf{h}_i^s | \mathbf{h}_i)$, can be constructed similarly to $p(\mathbf{h}_i^w | \mathbf{h}_i)$. Thus, we construct the conditional distributions of weakly and strongly augmented node embeddings given the original node embeddings.

To maintain the consistency of intrinsic semantic information contained in the graph-structured data between weak and strong augmentations, we align the conditional distributions of weakly and strongly augmented features over the original features, represented as $p(\mathbf{h}_i^w | \mathbf{h}_i)$ and $p(\mathbf{h}_i^s | \mathbf{h}_i)$. Specifically, we present a metric named the distribution divergence to quantify the divergence between these two conditional distributions as follows:

$$\mathcal{L}_d = -\frac{1}{n_g} \sum_{i=1}^{n_g} p(\mathbf{h}_i^w | \mathbf{h}_i) \log(p(\mathbf{h}_i^s | \mathbf{h}_i)). \quad (9)$$

By minimizing \mathcal{L}_d , the conditional distributions of weakly augmented node embeddings given the original node embeddings are employed to supervise the corresponding conditional distributions of strongly augmented node embeddings given the original node embeddings. The original node embeddings provide useful supervision information to ensure the consistency of the intrinsic semantic information between weak and strong augmentations. Consequently, this alignment enables the CDL model preserve intrinsic semantic information when both weak and strong augmentations are applied to graph-structured data.

The intrinsic semantic information within graph-structured data remains susceptible to distortion, even though weak augmentation is applied. We focus on how $p(\mathbf{h}_i^w | \mathbf{h}_i)$ provides valuable supervision information for $p(\mathbf{h}_i^s | \mathbf{h}_i)$. GCL is considered an effective strategy for mitigating the potential risk of distortion during weak augmentation. Unfortunately, there is a conflict between the message-passing mechanism of GNNs and the contrastive learning of negative pairs in augmented views, as previously discussed. For example, a node embedding \mathbf{h}_i^w is aggregated from the neighboring nodes $\{\mathbf{h}_j^w : j \in \mathcal{N}(i)\}$ in the weakly augmented view. However, these neighboring nodes are considered negative for \mathbf{h}_i^w when the loss function in Eq. (4) is employed. To evaluate the similarity between a positive pair $(\mathbf{h}_i^w, \mathbf{h}_i)$, the similarity loss is formulated as follows:

$$\mathcal{L}_s = -\frac{1}{n_g} \sum_{i=1}^{n_g} \log \frac{\exp(\text{sim}(\mathbf{p}_i^w, \mathbf{p}_i)/\tau)}{\sum_{j \neq i} \exp(\text{sim}(\mathbf{p}_i^w, \mathbf{p}_j)/\tau)} \quad (10)$$

where \mathbf{p}_i^w , \mathbf{p}_i and \mathbf{p}_j are computed via Eq. (7). Considering a pair of negative node embeddings $(\mathbf{h}_i^w, \mathbf{h}_j)$ ($j \neq i$) in Eq. (10), \mathbf{h}_i^w does not participate in the message-passing of the CDL model involving with \mathbf{h}_j , as \mathbf{h}_j originates from the original view. This similarity loss helps mitigate the aforementioned conflict in graph representation learning, making it beneficial for preventing the CDL model from overfitting to graph representations. The loss function in Eq. (10) is employed during the pretraining stage.

Algorithm 1: Optimization Procedure for CDL

Input: The training graphs $\mathcal{G}_r = \{(\mathcal{G}_i)\}_{i=1}^{n_g}$ and the testing graphs $\mathcal{G}_t = \{(\mathcal{G}_i)\}_{i=1}^{n_t}$.

Parameter: The number of iterations *epochs*, parameters $\alpha > 0$ and $\beta > 0$.

Output: The predicted labels of the testing graphs \mathcal{G}_t .

- 1: Construct a weakly augmented view \mathcal{G}_w and a strongly augmented view \mathcal{G}_s from \mathcal{G}_r ;
 - 2: **for** $t = 1$ to *epochs* **do**
 - 3: Compute \mathbf{P} and \mathbf{P}_w via Eq. (7) using the unlabeled parts of \mathcal{G}_r and \mathcal{G}_w , respectively;
 - 4: Update the network by minimizing \mathcal{L}_s in Eq. (10)
 - 5: **end for**
 - 6: **for** $t = 1$ to *epochs* **do**
 - 7: Compute \mathbf{q}_i via Eq. (3) for each labeled graph in \mathcal{G}_r ;
 - 8: Compute \mathbf{P} and \mathbf{P}_w via Eq. (7) using the labeled parts of \mathcal{G}_r and \mathcal{G}_w , respectively;
 - 9: Compute $p(\mathbf{h}_i^w|\mathbf{h}_i)$ and $p(\mathbf{h}_i^s|\mathbf{h}_i)$ via Eq. (8) using the labeled parts of \mathcal{G}_r , \mathcal{G}_s and \mathcal{G}_w ;
 - 10: Update the network by minimizing \mathcal{L} in Eq. (12);
 - 11: **end for**
 - 12: Compute \mathbf{q}_i via Eq. (3) for each testing graph in \mathcal{G}_t .
-

Applying strong augmentations to graph-structured data may distort its intrinsic semantic information. Similarly, weak augmentations can also introduce subtle distortions. Eq. (9) ensures consistency between weak and strong augmentations. Furthermore, Eq. (10) aims to preserve the complete intrinsic semantic information of graph-structured data as much as possible under weak augmentations. This preservation makes the alignment of conditional distributions in Eq. (9) both meaningful and effective.

Fine-Tuning Stage During the fine-tuning stage, a small number of labeled graphs are employed to train the CDL model. We employ the cross-entropy loss function for graph classification, i.e.,

$$\mathcal{L}_c = -\frac{1}{n_g} \sum_{i=1}^{n_g} \sum_{j=1}^c y_{ij} \log q_{ij} \quad (11)$$

where $\mathbf{y}_i \in \mathbb{R}^c$ denotes the one hot encoding for the label of the i th graph, and q_{ij} can be computed via Eq. (3). We can obtain the graph-level representation \mathbf{H}_s from a strongly augmented view via the shared GNN encoder module, similar to \mathbf{H} and \mathbf{H}_w . The loss function of the CDL model in Eq. (9) is introduced in the fine-tuning stage. Instead of maximizing the mutual information between the strongly augmented view and the original graph-structured data, the weakly augmented view can provide useful information to bridge the gap by minimizing the distribution divergence \mathcal{L}_d . Therefore, the overall loss of the proposed CDL model consists of three main components:

$$\mathcal{L} = \mathcal{L}_c + \alpha \mathcal{L}_s + \beta \mathcal{L}_d \quad (12)$$

where α and β are tradeoff hyperparameters. Algorithm 1 summarizes the entire optimization procedure of the proposed CDL method.

Theoretical Justification

Why is Pretraining Needed in CDL? Let two random variables $\mathbf{U} \in \mathbb{R}^{d'}$ and $\mathbf{V} \in \mathbb{R}^{d'}$ be node embeddings corresponding to a weakly augmented view and the original view, respectively, where d' is the dimensionality of the node embeddings. The loss function of the CDL model in Eq. (10) is a lower bound of mutual information (MI), i.e.,

$$I(\mathbf{U}; \mathbf{V}) \geq \log(n_g - 1) - \mathcal{L}_s. \quad (13)$$

Therefore, minimizing the loss function of the CDL model in Eq. (10) is equivalent to maximizing a lower bound on $I(\mathbf{U}; \mathbf{V})$. The proof is similar to that of minimizing the InfoNCE loss (Oord, Li, and Vinyals 2018). By pretraining, the intrinsic semantic information is preserved as much as possible when applying weak augmentation to the graph-structured data. This explains why $p(\mathbf{h}_i^w|\mathbf{h}_i)$ can provide valuable supervision information for $p(\mathbf{h}_i^s|\mathbf{h}_i)$ in CDL.

Generalization Bound of the Distribution Divergence

We analyze the generalization bound of the distribution divergence in \mathcal{L}_d in Eq. (9). The proof of Theorem 1 is provided in the supplementary material. According to Theorem 1, \mathcal{L}_d has a specific lower bound in Eq. (9) if two conditions are satisfied, i.e., $\text{sim}(\mathbf{h}_i, \mathbf{h}_k^w) = 0$ ($1 \leq k \leq K$) and $K \geq e^{\frac{1}{\tau}} - 1$. In particular, we have $\mathcal{L}_d \geq 0$ if $K \geq e^{\frac{1}{\tau}} - 1$. Therefore, a lower bound is theoretically guaranteed when minimizing \mathcal{L} in Eq. (12).

Theorem 1 Assume that there are n_g graphs and K negative samples for each graph, where $\text{sim}(\mathbf{h}_i, \mathbf{h}_k^w) = 0$ ($1 \leq k \leq K$). Given two conditional distributions \mathbf{h}_i^w and \mathbf{h}_i^s relative to \mathbf{h}_i , denoted as $p(\mathbf{h}_i^w|\mathbf{h}_i)$ and $p(\mathbf{h}_i^s|\mathbf{h}_i)$, respectively, and the distribution divergence \mathcal{L}_d in Eq. (9), the inequality $\mathcal{L}_d \geq \log(K + 1) - \frac{1}{\tau}$ holds if K satisfies the condition, i.e., $K \geq e^{\frac{1}{\tau}} - 1$, where τ is a parameter.

Experiments

Experimental Settings

We employed eight benchmark graph datasets from the TUDataset (Morris et al. 2020), including MUTAG, PROTEINS, IMDB-B, NCI1, RDT-B, RDT-M5K, COLLA and GITHUB, for evaluation. The statistics of these datasets and the parameter settings are provided in the supplementary material. Following the experimental settings in (Yue et al. 2022), we evaluated all competing algorithms via 10-fold cross-validation. Each dataset was randomly shuffled and evenly divided into 10 parts. One part was employed as the test set, another as the validation set, and the remaining parts were used for training. The training set contains two types of graphs. Three different percentages of the dataset, 30%, 50% and 70%, are selected from each graph set as labeled training graphs, while the labels of the remaining graphs are unknown during training. Classification accuracy was employed to evaluate the performance of all competing algorithms. We reported the average results and the standard deviations of the classification accuracy. A higher average accuracy on graph classification tasks indicates better preservation of the intrinsic semantic information. The source code is available at <https://github.com/chenjie20/SSCDL>.

Label	Methods	MUTAG	PROTEINS	IMDB-B	NCI1	RDT-B	RDT-M5K	COLLAB	GITHUB
30%	GCN	81.52±11.12	67.84±3.46	71.90±5.59	71.48±1.65	86.65±1.86	51.81±2.81	71.04±2.98	65.65±1.07
	GAT	82.02±9.58	63.25±3.31	72.60±4.40	64.94±2.83	79.30±4.18	49.83±1.85	71.42±5.82	65.28±0.83
	GCL	84.65±8.32	74.95±5.54	72.00±3.59	76.08±2.12	90.40±1.74	55.01±1.39	77.56±2.49	64.70±1.23
	GLIA	<u>87.25±8.44</u>	<u>75.12±4.54</u>	<u>73.30±4.45</u>	<u>76.72±2.20</u>	<u>90.35±1.63</u>	<u>55.09±1.45</u>	<u>77.94±1.81</u>	<u>68.61±1.24</u>
	G-Mixup	82.49±8.49	72.06±3.63	72.40±5.58	75.88±2.35	89.50±2.90	51.45±2.78	73.92±1.37	65.56±1.59
	GCMAE	70.23±7.46	69.46±3.41	72.70±4.47	61.97±1.73	76.60±2.03	-	72.14±1.66	-
	GRDL	83.45±9.12	69.55±3.80	72.50±4.32	68.49±0.12	-	-	71.64±6.77	-
	CDL	89.36±6.14	76.74±4.51	74.60±4.25	79.37±1.62	91.15±1.76	55.31±1.65	79.44±1.82	70.15±1.27
50%	GCN	83.10±8.83	69.27±3.17	72.40±4.99	72.80±1.93	87.30±3.54	52.59±3.19	73.98±1.58	66.55±1.17
	GAT	84.06±8.64	67.12±5.09	73.10±3.18	66.62±2.29	80.00±4.06	51.23±2.68	73.88±3.35	65.82±1.34
	GCL	85.09±8.19	75.30±3.50	72.80±4.47	76.72±1.61	90.80±1.7	55.97±2.23	79.36±2.04	64.46±1.54
	GLIA	<u>89.39±10.21</u>	<u>75.03±3.78</u>	<u>74.30±4.22</u>	<u>76.28±2.11</u>	<u>90.90±2.07</u>	<u>56.11±1.60</u>	<u>79.66±1.22</u>	<u>69.75±1.23</u>
	G-Mixup	81.43±8.32	72.15±3.45	72.90±5.58	77.02±2.20	89.80±1.70	52.27±2.09	74.30±1.83	66.14±1.45
	GCMAE	74.97±7.21	71.71±3.06	72.90±3.56	63.63±1.57	77.15±3.02	-	72.20±1.82	-
	GRDL	84.59±9.23	70.36±2.48	73.60±4.72	72.51±1.50	-	-	73.04±11.74	-
	CDL	89.94±8.76	76.10±2.80	74.90±3.70	79.08±1.86	92.05±2.14	56.51±2.32	80.96±1.29	70.83±1.13
70%	GCN	84.21±11.37	69.28±4.72	72.80±4.61	75.09±2.29	88.50±2.59	53.59±2.19	74.84±1.72	67.50±1.75
	GAT	84.68±8.32	68.46±4.89	73.30±4.00	67.62±2.74	81.05±3.83	51.53±2.71	72.18±1.43	66.04±1.38
	GCL	88.80±5.87	76.20±4.32	73.70±5.60	78.47±2.60	91.25±1.72	56.81±1.55	80.82±1.53	65.54±1.03
	GLIA	<u>89.42±7.45</u>	<u>75.75±4.82</u>	<u>74.50±3.08</u>	<u>79.29±2.13</u>	<u>91.40±1.96</u>	<u>56.09±1.56</u>	<u>81.18±1.67</u>	<u>69.95±0.77</u>
	G-Mixup	84.06±6.63	72.07±5.25	73.10±3.60	79.71±2.48	89.70±2.63	52.31±2.22	75.72±2.78	67.58±1.76
	GCMAE	76.05±5.92	72.37±2.45	73.60±5.10	64.89±1.99	78.35±3.03	-	72.18±1.43	-
	GRDL	85.67±7.02	71.53±2.85	73.30±5.83	72.51±1.50	-	-	74.06±10.33	-
	CDL	89.91±7.30	77.27±3.62	74.90±5.63	82.36±1.52	92.35±1.63	56.65±1.87	82.36±1.72	71.06±1.17

Table 1: The graph classification results (average accuracy (%) ± standard deviation (%)) with eight benchmark graph datasets.

To evaluate the performance of graph classification, we compared the results of the proposed CDL method with those of several state-of-the-art graph classification methods, including GCL (You et al. 2020), GLIA (Yue et al. 2022), G-Mixup (Han et al. 2022), the graph contrastive masked autoencoder (GCMAE) (Wang et al. 2024) and graph reference distribution learning (GRDL) (Wang and Fan 2024). We employed the GCN as a backbone in GCL. The graph classification results convey specific semantic information about the predicted labels of the graphs.

Weak and Strong Augmentation Settings The same graph augmentation strategy, node attribute masking, was utilized for both the weak and strong augmentations applied to the graph datasets. Specifically, a relatively small masking ratio of the node attributes is regarded as slight perturbations and is treated as a weak augmentation of graph-structured data. The primary difference between the weak and strong augmentations lies in the node masking ratio. The masking ratio of the node attributes for the weak augmentation was selected from the set $\{0.1, 0.2, 0.3, 0.35\}$, whereas the ratio for the strong augmentation was set to be twice that of the weak augmentation. By adjusting the magnitude of node attribute corruption through the masking ratio, we aim to investigate different levels of semantic information distortion while preserving the integrity of the overall graph topology.

Performance Evaluation

The graph classification results for competing methods across the eight graph datasets are reported in Table 1. The

best and second-best results are highlighted in bold and underlined, respectively. Table 1 shows that the CDL method consistently outperforms the other competing methods. For example, the CDL method achieves performance improvements of approximately 2.11%, 0.55% and 0.49% over the second-best method (GLIA) at labels ratios of 30%, 50% and 70%, respectively, with the MUTAG dataset. The gap in the classification results between the CDL and GLIA is small when the label ratio is 50% or 70% with the MUTAG dataset, due to the small number of graphs in this dataset. However, the gap in the classification results widened with the other datasets as the number of graphs increased. Moreover, we observe the same advantages of the CDL method with different label ratios.

GCL and GLIA are two representative graph contrastive learning-based methods used in the experiments. The GLIA algorithm achieves better graph classification results than the other competing methods do, such as G-Mixup, GCMAE and GRDL. GCL also achieves satisfactory results in the experiments. This finding highlights the significant benefits of employing contrastive learning on graphs. The symbol ‘-’ in Table 1 indicates an out-of-memory condition. GCMAE and GRDL often fail to perform graph classification for several of the large datasets, e.g., the RDT-M5K and GITHUB datasets, because of memory constraints.

Ablation Study

To investigate the impact of the loss components \mathcal{L}_s in Eq. (12), we performed ablation studies on graph classification tasks without the pretraining stage. In particular, we consid-

Label	Methods	\mathcal{L}_s	\mathcal{L}_c	\mathcal{L}_d	MUTAG	PROTEINS	IMDB-B	NCI1	RDT-B	RDT-M5K	COLLAB	GITHUB
30%	CDL _{cl}	✓			87.22±8.21	74.49±4.24	70.70±4.22	78.25±1.63	90.15±1.45	54.55±1.76	78.64±1.86	68.44±0.56
	CDL _{ft}	✓	✓		87.81±8.66	74.75±3.86	71.20±4.64	78.64±1.62	90.20±1.36	54.25±1.75	78.70±0.94	68.04±1.26
	CDL	✓	✓	✓	89.36±6.14	76.74±4.51	74.60±4.25	79.37±1.62	91.15±1.76	55.31±1.65	79.44±1.82	70.15±1.27
50%	CDL _{cl}	✓			89.36±7.89	74.58±3.45	71.3±5.56	77.91±2.65	90.25±1.27	55.19±2.02	79.98±1.16	69.04±1.06
	CDL _{ft}	✓	✓		88.80±9.48	75.21±3.06	72.4±4.01	78.03±1.74	91.05±1.77	55.29±1.75	80.08±2.18	69.22±1.24
	CDL	✓	✓	✓	89.94±8.76	76.10±2.80	74.90±3.70	79.08±1.86	92.05±2.14	56.51±2.32	80.96±1.29	70.83±1.13
70%	CDL _{cl}	✓			89.39±7.83	75.03±4.34	73.00±4.78	81.44±1.97	90.55±1.48	55.27±1.42	81.04±1.54	69.34±1.16
	CDL _{ft}	✓	✓		89.30±7.70	76.37±3.83	73.10±6.21	81.68±1.68	91.55±1.76	56.19±2.02	81.24±1.10	69.42±1.05
	CDL	✓	✓	✓	89.91±7.30	77.27±3.62	74.90±5.63	82.36±1.52	92.35±1.63	56.65±1.87	82.36±1.72	71.06±1.17

Table 2: Ablation study of the main components in Eq. (12) with eight benchmark graph datasets.




Figure 2: Graph classification results with different node masking ratios across different percentages of training samples. The datasets: (1) MUTAG, (2) PROTEINS, (3) IMDB-B, (4) NCI1, (5) RDT-B, (6) RDT-M5K, (7) COLLAB, and (8) GITHUB.

ered two scenarios in the experiments. We first employed the loss components \mathcal{L}_c and \mathcal{L}_s in Eq. (12), which represents a semisupervised graph classification model with pretraining. This model retains only positive pairs of node representations for contrastive learning, which measures the similarity between the original features and the corresponding weakly augmented features. Then, we only conducted the fine-tuning stage of CDL, which involves the two loss components, \mathcal{L}_c and \mathcal{L}_d , in Eq. (12). The invariants of the methods corresponding to these two scenarios are referred to as CDL_{cl} and CDL_{ft}, respectively.

Table 2 shows ablation study results for CDL_{cl} and CDL_{ft}. On the one hand, CDL_{ft} is superior to CDL_{cl} in most cases, even without the pretraining stage. This finding further highlights the effectiveness of the conditional distribution in \mathcal{L}_s . On the other hand, CDL achieves superior graph classification results compared with those of CDL_{ft}. This finding indicates that pretraining helps preserve intrinsic semantic information, which significantly improves the performance of the proposed CDL method. These experimental results demonstrate the effectiveness of our proposed semisupervised training scheme for graph classification.

Empirical Investigation

We empirically investigated the impact of varying node masking ratios on the performance of the proposed CDL method. The node masking ratio for weak augmentation was selected from the set $\{0.1, 0.2, 0.3, 0.35\}$. The node masking ratio for strongly augmented features is twice that of weak augmentation. Fig. 2 shows the graph classification

results with different node masking ratios across different percentages of training samples with eight benchmark graph datasets. The graph classification result often increases when the node masking ratio gradually increases from 0.1 to 0.3 with most datasets. This finding demonstrates that conditional distribution learning can effectively support graph representation learning. Additionally, the optimal range for the node masking ratio is relatively narrow, making it easier to determine this parameter in practice. In contrast, the graph classification result often decreases when the node masking ratio increases to 0.35. The primary reason for the decline in graph classification performance is the severe degradation of graph structure information in the strongly augmented view, where the node masking ratio is 0.7.

Conclusion

In this paper, we introduce the CDL method for semisupervised graph classification tasks. The CDL model is designed to align the conditional distributions of weakly and strongly augmented features over the original features. The proposed CDL method leverages the diversity and abundance provided by graph-structured data augmentations while preserving intrinsic semantic information. Moreover, CDL resolves the conflict between the message-passing mechanism of GNNs and the contrastive learning of negative pairs within intraviews. Additionally, we present a semisupervised learning scheme consisting of pretraining and fine-tuning stages. Extensive experimental results with various graph datasets demonstrate the superior performance of the CDL method.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China (NSFC) under Grants 62576231, 62572331 and 72374089, in part by the Fundamental Research Funds for the Central Universities under Grant 2462025YJRC038.

References

- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020a. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3438–3445. New York, USA.
- Chen, J.; Mao, H.; Woo, W. L.; and Peng, X. 2023. Deep multiview clustering by contrasting cluster assignments. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, 16752–16761. Paris, France.
- Chen, T.; Bian, S.; and Sun, Y. 2019. Are powerful graph neural nets necessary? A dissection on graph classification. *arXiv preprint arXiv:1905.04579*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, 1597–1607.
- Fan, W.; Ma, Y.; Li, Q.; Wang, J.; Cai, G.; Tang, J.; and Yin, D. 2022. A graph neural network framework for social recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 34(5): 2033–2047.
- Feng, J.; Chen, Y.; Li, F.; Sarkar, A.; and Zhang, M. 2022. How powerful are k -hop message passing graph neural networks. In *Proceedings of the 36th Advances in Neural Information Processing Systems*, 4776–4790. New Orleans, USA.
- Han, X.; Jiang, Z.; Liu, N.; and Hu, X. 2022. G-Mixup: Graph data augmentation for graph classification. In *Proceedings of the 37th International Conference on Machine Learning*, 8230–8248. Baltimore, Maryland USA.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, 448–456. Lille, France.
- Ji, C.; Huang, Z.; Sun, Q.; Peng, H.; Fu, X.; Li, Q.; and Li, J. 2024. ReGCL: rethinking message passing in graph contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8544–8552. Vancouver, Canada.
- Ju, W.; Mao, Z.; Yi, S.; Qin, Y.; Gu, Y.; Xiao, Z.; Wang, Y.; Luo, X.; and Zhang, M. 2024. Hypergraph-enhanced dual semi-supervised graph classification. In *Proceedings of the 41th International Conference on Machine learning*, 22594–22604. Vienna, Austria.
- Kazemi, S. M.; Goel, R.; Jain, K.; Kobayez, I.; Sethi, A.; Forsyth, P.; and Poupart, P. 2020. Representation learning for dynamic graphs: A survey. *Journal of Machine Learning Research*, 21(70): 1–73.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of 5th International Conference on Learning Representations*, 1–14. Toulon, France.
- Li, M.; and Zhu, Z. 2021. Spatial-temporal fusion graph neural networks for traffic flow forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4189–4196. Vancouver, Canada.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond*.
- Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *Proceedings of the 30th Advances in Neural Information Processing Systems*, 1857–1865. Barcelona, Spain.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Li, P.; and Bengio, Y. 2018. Graph attention networks. In *Proceedings of the 6th International Conference on Learning Representations*, 1–12. Vancouver, BC, Canada.
- Vignac, C.; Loukas, A.; and Frossard, P. 2020. Building powerful and equivariant graph neural networks with structural message-passing. In *Proceedings of the 34th Advances in Neural Information Processing Systems*, 14143–14155. Online.
- Wang, X.; and Qi, G.-J. 2023. Contrastive learning with stronger augmentations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5): 5549–5560.
- Wang, Y.; Wang, J.; Cao, Z.; and Farimani, A. B. 2022. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3): 279–287.
- Wang, Y.; Yan, X.; Hu, C.; Xu, Q.; Yang, C.; Fu, F.; Zhang, W.; Wang, H.; Du, B.; and Jiang, J. 2024. Generative and contrastive paradigms are complementary for graph self-supervised learning. In *IEEE 40th International Conference on Data Engineering*, 3364–3378. Geneva, Switzerland.
- Wang, Z.; and Fan, J. 2024. Graph classification via reference distribution learning theory and practice. In *Proceedings of 38th Advances in Neural Information Processing Systems*, 29350–29361. Vancouver, Canada.
- Wu, L.; Lin, H.; Tan, C.; Gao, Z.; and Li, S. Z. 2023. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE Transactions on Knowledge and Data Engineering*, 35: 4216–4235.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. In *Proceedings of the 34th Advances in Neural Information Processing Systems*, 5812–5823. Online.
- Yue, H.; Zhang, C.; Zhang, C.; and Liu, H. 2022. Label-invariant augmentation for semi-supervised graph classification. In *Proceedings of 36th Advances in Neural Information Processing Systems*, 29350–29361. New Orleans, Louisiana, USA.

Zeng, J.; and Xie, P. 2021. Contrastive self-supervised learning for graph classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10824–10832. Vancouver, Canada.

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*, 2069–2080. Ljubljana, Slovenia.