

Correlation-Aware Graph Convolutional Networks for Multi-Label Node Classification

Yuanchen Bei
Zhejiang University
Hangzhou, China
yuanchenbei@zju.edu.cn

Weizhi Chen
Zhejiang University
Hangzhou, China
chenweizhi@zju.edu.cn

Hao Chen
The Hong Kong Polytechnic
University
HongKong SAR, China
sundaychenhao@gmail.com

Sheng Zhou*
Zhejiang University
Hangzhou, China
zhousheng_zju@zju.edu.cn

Carl Yang
Emory University
Atlanta, USA
j.carlyang@emory.edu

Jiapei Fan
Alibaba Group
Hangzhou, China
jiapei.fjp@@alibaba-inc.com

Longtao Huang
Alibaba Group
Hangzhou, China
kaiyang.hlt@@alibaba-inc.com

Jiajun Bu
Zhejiang University
Hangzhou, China
bjj@zju.edu.cn

Abstract

Multi-label node classification is an important yet under-explored domain in graph mining as many real-world nodes belong to multiple categories rather than just a single one. Although a few efforts have been made by utilizing Graph Convolution Networks (GCNs) to learn node representations and model correlations between multiple labels in the embedding space, they still suffer from the ambiguous feature and ambiguous topology induced by multiple labels, which reduces the credibility of the messages delivered in graphs and overlooks the label correlations on graph data. Therefore, it is crucial to reduce the ambiguity and empower the GCNs for accurate classification. However, this is quite challenging due to the requirement of retaining the distinctiveness of each label while fully harnessing the correlation between labels simultaneously. To address these issues, in this paper, we propose a **Correlation-aware Graph Convolutional Network (CorGCN)** for multi-label node classification. By introducing a novel Correlation-Aware Graph Decomposition module, CorGCN can learn a graph that contains rich label-correlated information for each label. It then employs a Correlation-Enhanced Graph Convolution to model the relationships between labels during message passing to further bolster the classification process. Extensive experiments on five datasets demonstrate the effectiveness of our proposed CorGCN.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '25, August 3–7, 2025, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1245-6/25/08
<https://doi.org/10.1145/3690624.3709197>

CCS Concepts

• **Information systems** → **Web mining**; *Social networks*; • **Human-centered computing** → *Social network analysis*.

Keywords

multi-label node classification, graph learning, graph neural networks, graph data mining

ACM Reference Format:

Yuanchen Bei, Weizhi Chen, Hao Chen, Sheng Zhou, Carl Yang, Jiapei Fan, Longtao Huang, and Jiajun Bu. 2025. Correlation-Aware Graph Convolutional Networks for Multi-Label Node Classification. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25)*, August 3–7, 2025, Toronto, ON, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3690624.3709197>

1 Introduction

Node classification serves as a cornerstone in the field of graph mining [7, 39]. Over the past decade, Graph Convolutional Networks (GCNs) have achieved remarkable success in this area by aggregating information from neighboring nodes, where edges in the graph often indicate similarities in the single-label space [18, 29]. However, graphs in real-world scenarios often entail multi-label nodes. For example, users in social networks usually exhibit broad interests and embody multiple labels [27, 45], and protein nodes in protein interaction networks usually carry several relevant gene ontology annotations [12, 44]. The diverse labels offer varied perspectives for delineating node characteristics while also introducing both challenges and opportunities for GCNs in the realm of multi-label node classification.

Inspired by the success of multi-label learning on non-relational data [21, 33, 35], such as images and text, there has been a growing interest in applying multi-label learning to relational data, including graphs [11, 38, 49]. Among the few existing methods for multi-label node classification, most employ GCNs to map nodes to low-dimensional representations. They then follow the traditional multi-label learning paradigm to model the relationships

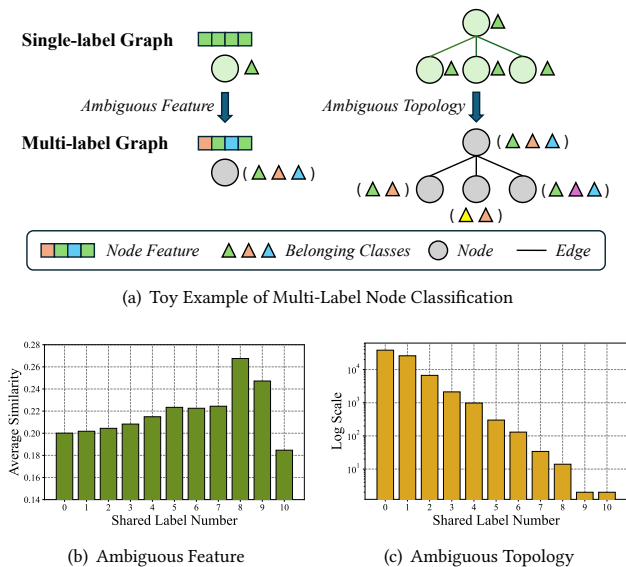


Figure 1: (a) A toy example of multi-label node classification challenges. (b)-(c) Illustrations of the ambiguous feature and ambiguous topology phenomena on the PCG dataset.

between nodes and labels, as well as the inter-label relationships in the representation space. However, this paradigm overlooks the characteristics of graph data in multi-label scenarios, as Figure 1-(a), leading to challenges for GCNs. (1) **Ambiguous Feature**: In the single-label setting, nodes can aggregate patterns specific to a particular type of label from one neighborhood node by transforming the features. However, in the multi-label setting, the feature of one node may be related to multiple labels and the patterns represented by the feature are ambiguous. As a result, aggregating information from such ambiguous features will significantly impair the discriminative ability of the representation. Figure 1-(b) illustrates that nodes that allocate similar features may share a diverse number of labels in real-world graphs. (2) **Ambiguous Topology**: In the single-label setting, connected nodes typically share the same single label (also known as the homophily assumption), so the pattern propagated along the edge is usually deterministic. However, in the multi-label setting, both the connected nodes have multiple labels, and the patterns propagated along the edge are often ambiguous. This ambiguity makes it difficult to determine from which connected nodes we should aggregate specific label information. Aggregating information directly from all neighbors would further accumulate ambiguity, compromising the discriminative power of the representations learned by GCNs and eventually impacting the inferring of specific labels. Figure 1-(c) illustrates that connected nodes may share diverse numbers of labels in real-world graphs. *It is essential to reduce the ambiguity so that the potential of GCNs in multi-label node classification can be fully released.*

Although important, addressing the above issues is non-trivial and meets the following challenges: (1) **Label Distinctiveness**: As previously discussed, node attributes and edges in a multi-label graph may be influenced by multiple labels simultaneously. Directly extracting information from such a graph with a mixture of node

labels can result in the loss of label distinctiveness and lead to inadequate exploration of each label. (2) **Label Correlation**: In the multi-label setting, a node’s association with multiple labels implies a correlation between these labels. The success of existing multi-label learning methods also demonstrates that fully leveraging these inter-label correlations can significantly enhance the quality of the representation [28, 48]. Consequently, merely extracting messages or neighbors for a single label only could forfeit these vital correlations, leading to suboptimal outcomes. This raises a pertinent question: *How can we retain the distinctiveness of each label while fully harnessing the correlation between labels to achieve more accurate multi-label learning?*

To address the aforementioned challenges, in this paper, we present (**CorGCN**), a **Correlation-Aware Graph Convolutional Network** for multi-label node classification. To tackle the first challenge, we propose a novel graph decomposition strategy where we learn an individual graph for each label while preserving its unique characteristics. To tackle the second challenge, we fully consider the correlation between labels during the graph decomposition process. Instead of indiscriminately discarding all information from other categories, we retain the information of related labels within each category’s graph. Lastly, based on the multiple graphs generated, we further exploit the correlation between labels using a newly designed correlation-enhanced GCN that takes category correlation into account, thereby enhancing the final multi-label classification results. We conduct extensive experiments on five datasets with comprehensive metrics. The results demonstrate the effectiveness of the proposed method. The primary contributions of this paper can be summarized as follows:

- We highlight the overlooked impact of ambiguous feature and ambiguous topology on GCNs in the multi-label node classification, which is fundamental in this critical task.
- We propose a Correlation-aware Graph Convolutional Network (CorGCN) to simultaneously retain label-distinct characters and label correlation information for multi-label node classification.
- Extensive experiments on five datasets demonstrate CorGCN significantly outperforms nine state-of-the-art baselines across seven metrics. Further in-depth analysis from diverse perspectives also demonstrates the strengths of CorGCN.

2 Related Works

2.1 Multi-Label Learning

Multi-label learning, where each instance is associated with a set of correlated labels rather than a single label, is widely applicable in various domains such as object detection [13] and text classification [37]. For example, an image may contain multiple relevant objects and a text can encompass various topics [5, 21, 22, 35]. The key to multi-label learning is to model the multi-label correlation. Early studies convert the multi-label learning problem into multiple single-label learning problems, which has limitations in its ability to explicitly model correlations [6, 24]. Recently, various studies have been proposed with deep representation learning. Among them, one type of method learns correlated representations with statistical analysis [32, 42]. The second type of method learns the label correlation on label embeddings with label sequence/relational modeling techniques, such as RNN and GNN [10, 31, 41]. Another type of

model proposes to co-learn label and instance representations with the auto-encoder architecture, which models the correlation during the latent encoding [1, 2, 48].

Despite significant progress, these methods are all focused on Euclidean data. In recent years, due to the development of GNNs, the problem of multi-label learning on non-Euclidean graphs has also shown its importance [47]. However, thus far, multi-label learning on graph data has not received much attention.

2.2 Multi-Label Node Classification

Node classification is a well-known fundamental task in graph data mining. Typically, node classification refers to the assignment of a unique label to each node [4, 14]. In recent years, multi-label node classification (MLNC) has posed significance due to the increasing realization that nodes on graphs often exhibit multiple associated categories simultaneously [11, 47]. For example, in social network analysis, a user often belongs to multiple interest groups [8, 26, 36]. This involves assigning non-unique and variable numbers of labels to each node. Currently, MLNC is still in its infancy, and only a few studies have focused on MLNC [26, 47]. Representatively, ML-GCN [11] learns distinct embeddings for each label and constructs node-label correlations to augment the GCN-encoded node embeddings. Subsequently, LANC [49] and LARN [38] further incorporate attention mechanisms to integrate label embeddings with GCN-encoded node embeddings more effectively.

Nonetheless, current approaches continue to utilize a unified graph convolution process across neighborhoods under various labels, which ignores the possibility that information among different multi-label nodes may become intermixed in a unified message passing process, including both correlated and uncorrelated labels.

2.3 Graph Structure Learning

Since the natural graph structure often contains noise information, it's not always directly applicable to a wide range of downstream tasks. This limitation has brought the field of Graph Structure Learning (GSL) into the spotlight, garnering increasing attention in recent years [3, 19]. Representatively, GRCN [43] adjusts edge weights in the graph, optimizing with downstream tasks. IDGL [9] iteratively improves the graph structure through node embeddings. And SUBLIME [23] uses a contrastive-based objective function to guide the learned graph. These GSL models generally strive to refine the graph structure by enhancing graph homophily or by augmenting task-specific useful information in the single-label space. Despite these advances, there is a noticeable absence of GSL approaches specifically tailored for MLNC, leaving a significant gap in the field.

In light of these challenges, our paper is dedicated to introducing a GSL-based methodology, uniquely designed for MLNC. This method aims to address the existing limitations by integrating the understanding of multi-label correlations and adapting the graph structure accordingly.

3 Preliminaries

Notations. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, consisting of two sets: a set of nodes \mathcal{V} and a set of edges \mathcal{E} . Here, \mathcal{V} represents the set of nodes in the graph, while $\mathcal{E} \subseteq \{(u, v) | u, v \in \mathcal{V}\}$ represents the

set of edges between the nodes. Then, the adjacency matrix of \mathcal{G} can be defined as $A \in \mathbb{R}^{n \times n}$ with m non-zero element in the matrix based on the edge set, where $n = |\mathcal{V}|$ and $m = |\mathcal{E}|$ are the node number and edge number, respectively. Each node $v_i \in \mathcal{V}$ in the graph \mathcal{G} is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^f$ and thus all the node feature vectors form the overall feature matrix $X \in \mathbb{R}^{n \times f}$ of the graph \mathcal{G} . To distinguish from the notation of decomposed graphs introduced in this paper, we represent the original graph adjacent matrix with the symbol A^0 .

Definition 1 (Graph with Multi-Label Nodes). In a multi-label graph, each node $v_i \in \mathcal{V}$ is associated with a set of labels $L_i \subseteq L$, where L is the set of all possible labels and $|L| = K$ is the total number of labels. Thus, each node can have multiple labels, i.e., $|L_i| \geq 1$. In practice, each node v_i will be assigned a multi-hot label vector $\mathbf{y}_i \in \mathbb{R}^{1 \times K}$, where each element $y_{i,j} \in [0, 1]$ in \mathbf{y}_i indicates whether node v_i belongs to class j .

Definition 2 (Multi-Label Node Classification). In the multi-label node classification scenario, we typically have a subset of nodes with known labels and another subset with unknown labels [47]. The node set \mathcal{V} in \mathcal{G} can be divided into two disjoint subsets: a labeled node set \mathcal{V}_L and an unlabeled node set \mathcal{V}_U . Thus, $\mathcal{V} = \mathcal{V}_L \cup \mathcal{V}_U$ and $\mathcal{V}_L \cap \mathcal{V}_U = \emptyset$. The goal of multi-label node classification is to learn a model $f : \mathcal{V} \rightarrow \{0, 1\}^K$ using the labeled nodes $v_i \in \mathcal{V}_L$ and the graph topology A . This model aims to predict the label set \mathbf{y}_j for unlabeled nodes $v_j \in \mathcal{V}_U$.

4 Methodology

Our proposed model for multi-label node classification comprises the correlation-aware graph learning module with node feature and graph topology decomposition (Sec. 4.1) and the correlation-aware graph convolution (Sec. 4.2 and Sec. 4.3). Further, we illustrate how CorGCN can be expanded to the scenario with large label space (Sec. 4.4). Figure 2 provides an overview of the architecture of our proposed CorGCN.

4.1 Correlation-Aware Graph Decomposition

Due to the ambiguity in node features and topological structures of graphs in multi-label scenarios, we aim to decompose them into multiple graphs. However, directly performing the decomposition would result in the loss of the critically important multi-label correlation property. Therefore, we need to carry out the decomposition based on representations that have encapsulated label correlation.

4.1.1 Correlation-Aware Feature Decomposition. This component primarily adopts an approach that first models the node-label correlation and label-label correlation, and then, based on this, performs the decomposition of node features across different label spaces. Summarizing these, we model the correlations with contrastive mutual information estimator ϕ and likelihood maximizing decoder θ , which can be defined as,

$$\arg \max_{\phi, \theta} \underbrace{I_{\phi}(E^X, E^L)}_{\text{Mutual Information}} + \underbrace{\mathcal{L}_{\theta}(\theta; E^X) + \mathcal{L}_{\theta}(\theta; E^L)}_{\text{Likelihood}}. \quad (1)$$

Specifically, to fully utilize the knowledge from labeled nodes for generalization to all nodes in the graph, we first assign a trainable

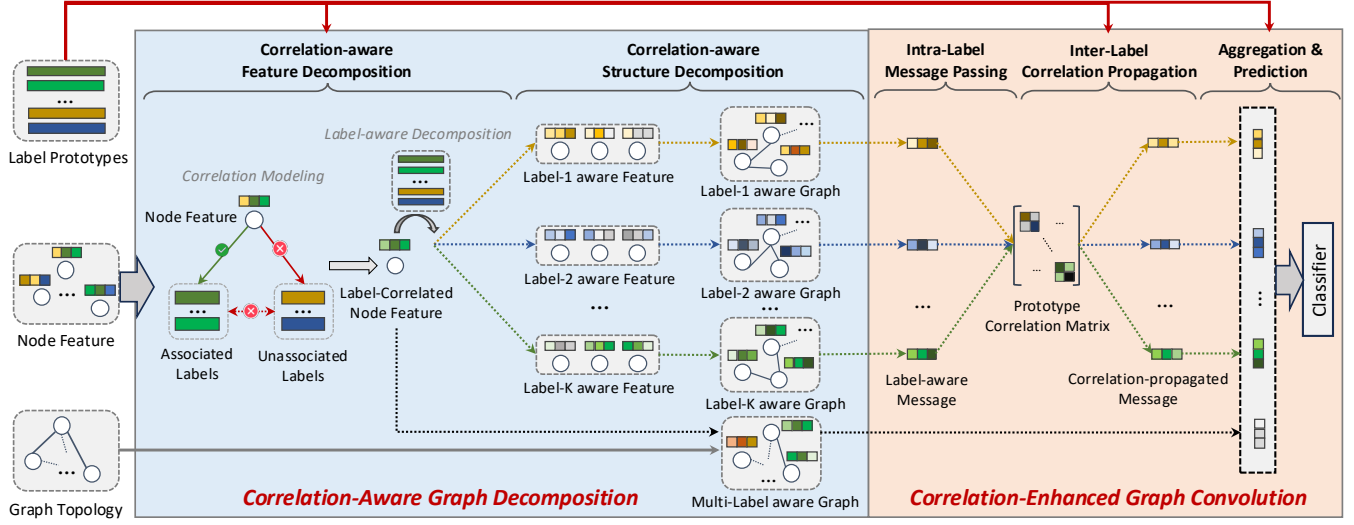


Figure 2: The overall architecture of CorGCN. (a) Correlation-Aware Graph Decomposition: it first learns label-correlated node features and decomposes them into multiple label-aware features. Then, based on the decomposed features, it decomposes multiple label-aware graphs. **(b) Correlation-Enhanced Graph Convolution:** each layer includes intra-label message passing of the neighborhood in each label-aware graph view and inter-label correlation propagation between label-aware messages.

prototype embedding to each label to describe the characteristics of this label:

$$E^l = [E_1^l, \dots, E_k^l, \dots, E_K^l], \quad (2)$$

where $1 \leq k \leq K$, and $E_k^l \in \mathbb{R}^d$ is the k -th label prototype. We then transform the original node features X to the same d -dimensional space using the linear transformation: $E^x = XW_t$, where $W_t \in \mathbb{R}^{f \times d}$ is a learnable transformation matrix, and each vector $E_i^x \in \mathbb{R}^d$ in E^x denotes the transformed feature of node v_i :

$$E^x = [E_1^x, \dots, E_i^x, \dots, E_n^x]. \quad (3)$$

Then, the mutual information estimator ϕ captures both node-label and label-label correlations based on those supervised nodes. Previous works for multi-label learning have successfully adopted contrastive learning to model multi-label correlations and inspired by [2], we utilize contrastive learning as ϕ here. Given the labeled nodes \mathcal{V}_L , the contrastive-based mutual information estimator with \mathcal{L}_{cmi} can be written as follows:

$$\mathcal{L}_{cmi} = -\frac{1}{|\mathcal{V}_L|} \sum_{i \in \mathcal{V}_L} \frac{1}{|\text{Pos}(\mathbf{y}_i)|} \sum_{p \in \text{Pos}(\mathbf{y}_i)} \log \frac{\exp(E_i^x \cdot E_p^l)}{\sum_{t \in \mathbf{y}_i} \exp(E_i^x \cdot E_t^l)}, \quad (4)$$

where $\text{Pos}(\mathbf{y}_i)$ denotes the label set that node v_i belongs to. Through this formulation, we directly map node features into a compact space that reflects multi-label correlations, while implicitly capturing the nuances of multi-label prototype relationships. Unlike single-label scenarios where classes are mutually exclusive, multi-label classification acknowledges the complexity of label correlations. Therefore, we do not enforce strict contrastive relations

among labels, maintaining their intrinsic associations [2, 48]. Instead, it utilizes feature embeddings as anchors and label embeddings as positive or negative examples, allowing frequently co-occurring labels to become more similar over time.

Subsequently, due to the above process mainly focuses on embedding space without explicit guidance, in order to ensure that different node features and label prototypes retain their own classification characteristics, we further constrain this through likelihood maximizing decoder with corresponding loss \mathcal{L}_{lm} as follows:

$$\rho_k = \frac{\sqrt{\frac{1}{\sum_{v_i \in \mathcal{V}_L} \mathbf{y}_{i,k}}}}{\sum_{j=1}^K \sqrt{\frac{1}{\sum_{v_i \in \mathcal{V}_L} \mathbf{y}_{i,j}}}}, \quad (5)$$

$$\mathcal{L}_{lm} = -\frac{1}{2|\mathcal{V}_L|} \sum_{v_i \in \mathcal{V}_L} \sum_{k=1}^K \rho_k [(1 - p_k(E_i^x | \theta))^\gamma \cdot \log p_k(E_i^x | \theta) + (1 - p_k(\mathbf{y}_i \cdot E^l | \theta))^\gamma \cdot \log p_k(\mathbf{y}_i \cdot E^l | \theta)], \quad (6)$$

where ρ_k is a statistical parameter to enhance the learning for classes with a small number of node samples and γ is a hyper-parameter to control the observation of hard node/label samples [20]. $\hat{p}_k(E_i^x | \theta)$ is the predicted result of the decoder θ , and if $y_{i,k} = 1$, $p_k(E_i^x | \theta) = \hat{p}_k(E_i^x | \theta)$. Otherwise, $p_k(E_i^x | \theta) = 1 - \hat{p}_k(E_i^x | \theta)$.

In the above process, we obtain the correlated node features E^x and label prototypes E^l . Each label prototype is associated with a specific label and its potentially correlated labels while each correlated node feature is still ambitious due to the mixture of multiple label information. Based on this, we decompose the correlated node features to each label-friendly feature space that *considers both the label-specific and its correlated label-specific information* with the

guidance of label prototypes. Specifically, given the label-correlated node feature E_i^x of node v_i and the label prototype of label k , the projection process can be obtained by:

$$w_{i,k} = \text{sim}(E_i^x, E_k^l) = \frac{E_i^x \cdot E_k^l}{\|E_i^x\| \|E_k^l\|}, \quad (7)$$

$$\begin{aligned} E_i^{proj} &= [E_{i,1}^{proj}, \dots, E_{i,k}^{proj}, \dots, E_{i,K}^{proj}] \\ &= [w_{i,1} E_i^x, \dots, w_{i,k} E_i^x, \dots, w_{i,K} E_i^x], \end{aligned} \quad (8)$$

where $w_{i,k}$ is the projection coefficient of node i to label k , $E_{i,k}^{proj}$ is the projected representation of node v_i towards the k -th label.

4.1.2 Correlation-Aware Structure Decomposition. Based on the correlation-aware decomposed node features E^{proj} , we aim to decompose the graph structure (message passing path) for each label with its correlated labels.

To preserve the graph topology patterns, we first aggregate the neighborhood projected features of each center node v_i in each label view k for structure decomposition:

$$E_{i,k}^{sd} = \text{Agg}(E^{proj}, A^0) = \frac{1}{|\mathcal{N}_i^0| + 1} \sum_{j \in \{\mathcal{N}_i^0 \cup \{v_i\}\}} E_{j,k}^{proj}, \quad (9)$$

where \mathcal{N}_i^0 is the neighborhood of node v_i in original graph adjacent matrix A^0 . We leverage the aggregated representations to measure the node similarity in each label view:

$$S_{i,j}^k = \text{sim}(E_{i,k}^{sd}, E_{j,k}^{sd}) = \frac{E_{i,k}^{sd} \cdot E_{j,k}^{sd}}{\|E_{i,k}^{sd}\| \|E_{j,k}^{sd}\|}, \quad (10)$$

where S^k is the similarity score matrix in label view- k , $S_{i,j}^k$ is the score between node v_i and node v_j .

Then, for each node v_i , the structure decomposing for different label-aware graphs can be described as:

$$A_{i,j}^k = \begin{cases} 1, & S_{i,j}^k \in \text{Top-}\lambda(S_i^k); \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

$$\mathcal{G}^k = (A^k, E_k^{proj}), \quad (12)$$

where A^k is the adjacent matrix of the k -th label view, and λ is the hyperparameter that controls the density of the decomposed graph.

Furthermore, with the correlated label-aware node feature and the original graph topology, the multi-label aware graph $\mathcal{G}^0 = (A^0, E^x)$ can be obtained to capture correlated structure patterns by message passing, and thus the learned correlation-aware decomposed graphs CDG can be obtained as follows:

$$CDG = \{\mathcal{G}^0, \mathcal{G}^1, \dots, \mathcal{G}^K\}. \quad (13)$$

4.2 Correlation-Enhanced Graph Convolution

Previous studies for multi-label node classification mainly conduct unified neighborhood message passing [11, 38, 49]. We contend this approach has two primary limitations: (1) passing ambiguous messages from the neighborhood in a unified manner; and (2) correlation ignorance, which overlooks the label correlation when passing the message. Therefore, we further equip the CDG with Correlation-Enhanced graph convolution.

4.2.1 Intra-Label Message Passing. Firstly, Correlation-Enhanced graph convolution conducts intra-label message passing within the graph \mathcal{G}^k of each label view k :

$$\begin{aligned} \hat{Z}_{[:,k]}^{(l)} &= \text{GCN}(Z_{[:,k]}^{(l-1)}, \tilde{A}^k) \\ &= \sigma(\tilde{A}^k Z_{[:,k]}^{(l-1)} \mathbf{W}^{(l)}), \end{aligned} \quad (14)$$

where $0 \leq k \leq K$, $\tilde{A}^k = \hat{D}^{k-\frac{1}{2}} A^{*k} \hat{D}^{k-\frac{1}{2}} \in \mathbb{R}^{n \times n}$ is the normalized adjacency matrix of label view- k , $\hat{D}^k \in \mathbb{R}^{n \times n}$ is the degree matrix of $A^{*k} = A^k + I$ where I is the identity matrix. $\hat{Z}^{(l)}$, $\mathbf{W}^{(l)}$ is the output features, trainable parameters in l -th message passing layer, respectively. The input of the first layer $Z^{(0)} = [E^x, E_1^{proj}, \dots, E_K^{proj}] \in \mathbb{R}^{n \times (K+1) \times d}$ is the feature matrix. Note that this is a GCN-like [18] message passing function, any other graph message passing functions [16, 29] can also be adopted here.

4.2.2 Inter-Label Correlation Propagation. After the above intra-label message passing, we obtain the neighborhood messages $\hat{Z}^{(l)} \in \mathbb{R}^{n \times K \times d}$ from each label-aware view. Then, to model their correlations, we further propose the inter-label correlation propagation between each label-aware graph view:

$$\text{Cor}_i = \text{Softmax}\left(\frac{(E^l \mathbf{W}_1)(\hat{Z}_{[i,1]}^{(l)} \mathbf{W}_2)^T}{\sqrt{d}}\right), \quad (15)$$

$$Z_{[i,1]}^{(l)} = \text{Cor}_i \hat{Z}_{[i,1]}^{(l)} \mathbf{W}_3, \quad (16)$$

where $\text{Cor}_i \in \mathbb{R}^{K \times K}$ is the label prototype correlation matrix for node v_i , $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3$ are the trainable parameters, and $Z_{[i,1]}^{(l)} \in \mathbb{R}^{K \times d}$ is the output of v_i in l -th graph convolution layer containing K views of inter-label correlated representations.

4.2.3 Aggregation & Prediction. After the multi-layer graph message passes, we aggregate these representations to obtain the final representation with both the node representations and label prototypes. The aggregation process can be written as:

$$Z_i^{cls} = [Z_{i,0} || (\sum_{k=1}^K \text{sim}(Z_{i,k}, E_k^l) \cdot Z_{i,k})], \quad (17)$$

$$\hat{y}_i = \sigma(\mathbf{W}^{cls} Z_i^{cls} + \mathbf{b}^{cls}), \quad (18)$$

where $Z_{i,k}$ is output features from final layers of correlation propagation for node v_i in the label view k , $Z_i^{cls} \in \mathbb{R}^{n \times 2d}$ is the final node representation for multi-label node classification of node v_i and $||$ is the concatenation operation, \mathbf{W}^{cls} and \mathbf{b}^{cls} are the trainable parameters in the classifier, and σ is the Sigmoid function.

4.3 Objective Function

To train and optimize the model parameters, we apply the binary cross-entropy loss as the model objective function. Formally, for each node $v_i \in \mathcal{V}_L$, the classification objective function can be expressed as:

$$\mathcal{L}_{cls} = -\frac{1}{|\mathcal{V}_L|} \sum_{i \in \mathcal{V}_L} \frac{1}{K} \sum_{j=1}^K y_{i,j} \log(\hat{y}_{i,j}) + (1-y_{i,j}) \log(1-\hat{y}_{i,j}), \quad (19)$$

where $\hat{y}_{i,j}$ is the predicted logits of node v_i towards label j and $y_{i,j}$ is the corresponding ground-truth label. Then, considering

the objective losses in each CorGCN module, the overall objective function of CorGCN is as follows:

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \cdot \mathcal{L}_{cmi} + \beta \cdot \mathcal{L}_{lm}, \quad (20)$$

where α and β are the controllable hyper-parameters.

4.4 Extension to Large Label Space

In some real-world graph structures, nodes may exist in a very large multi-label space, such as in large protein interaction networks, where individual protein nodes may have hundreds of characteristic labels simultaneously [17, 25]. In this subsection, we demonstrate that our proposed CorGCN can be readily extended to multi-label node classification with large label space for efficient learning.

Macro Label Prototypes. For the large label space, directly applying multi-label graph learning with K prototypes will obtain hundreds of label-aware graphs, which is not efficient enough for applications. Therefore, we aim to refine and cluster the original prototypes into macro-label prototypes, with the principle that correlated labels will form a macro prototype.

Specifically, we first pre-train the K label prototypes in Eq.(1) with \mathcal{L}_{cmi} and \mathcal{L}_{lm} (as Sec. 4.1). Then, we adopt K-means-based clustering to refine the K pre-trained label prototypes into K' macro prototypes [15]. Given the pre-trained label prototypes E^l , we randomly initialize K' macro prototype centroids $\{\mu_1, \dots, \mu_k, \dots, \mu_{K'}\}$, where $\mu_k \in \mathbb{R}^d$ is the centroids of macro label prototype C_k , $K' \ll K$ is the hyperparameter set as the macro prototype number. Then, we explore and assign the original label prototypes to the appropriate macro label prototype based on their representation similarity, and update the centroid of the macro label prototype iteratively. The process can be expressed as:

$$\mu_k = \frac{1}{|C_k|} \sum_{x_v=k, E_v^l \in C_k} E_v^l, \quad (21)$$

where $|C_k|$ is the number of labels within C_k , and x_v is the macro label prototype index that v is assigned to. Further, the optimization target of the macro label prototype generation is:

$$\begin{aligned} \min_{\substack{x_1, \dots, x_{K'} \\ \mu_1, \dots, \mu_{K'}}} J(x_1, \dots, x_{K'}; \mu_1, \dots, \mu_{K'}) \\ \triangleq \sum_{k=1}^{K'} \sum_{x_v=k, E_v^l \in C_k} \sqrt{(E_v^l - \mu_k)(E_v^l - \mu_k)^\top}, \end{aligned} \quad (22)$$

where J is the objective function of macro label prototype generation. In this way, we obtain the macro label prototype set:

$$E^{ml} = [\mu_1, \dots, \mu_k, \dots, \mu_{K'}], \quad (23)$$

where $1 \leq k \leq K' \ll K$. Therefore, we can reduce the original number of label prototypes to meet our needs. With the obtained E^{ml} , we can flexibly replace E^l with E^{ml} and procedure Sec. 4.1 - Sec. 4.3 in the same way. Model analyses are detailed in Appendix A.

5 Experiments

In this section, we conduct comprehensive experiments on benchmark datasets with CorGCN, aiming to answer the following questions: **RQ1:** How does CorGCN perform compared to representative and state-of-the-art models? **RQ2:** What is the effect of different components in CorGCN? **RQ3:** Can CorGCN be generalized and

adapted to different GNN message passing backbones? **RQ4:** How does the impact of CorGCN on each fine-grained class? **RQ5:** What is the efficiency of CorGCN on training and inference? **RQ6:** How do key hyper-parameters impact the performance of CorGCN?

5.1 Experimental Setup

5.1.1 Datasets. We conduct experiments on five widely used benchmark datasets: **Humloc** [47], **PCG** [47], **Blogcatalog** [49], **PPI** [44], and **Delve** [38] to verify the effectiveness of CorGCN. The statistics of these datasets are shown in Table 1. We randomly divided the labeled data of these datasets into training, validation, and test sets, following a 6:2:2 ratio, respectively [47]. The performance on PPI and Delve are conducted on the extension of CorGCN, with 20 and 10 macro label prototypes, respectively. Detailed descriptions of these datasets are illustrated in Appendix B.1.

Table 1: Statistics of the experimental datasets.

Dataset	# Nodes	# Edges	# Features	# Classes	Density
Humloc	3,106	18,496	32	14	0.3836%
PCG	3,233	37,351	32	15	0.7149%
Blogcatalog	10,312	333,983	100	39	0.6282%
PPI	14,755	225,270	50	121	0.2070%
Delve	1,229,280	4,322,275	300	20	0.0006%

5.1.2 Baselines. To evaluate the effectiveness of CorGCN, we compare it with nine representative and state-of-the-art models. (i) Traditional GNN for node classification: **GCN** [18]. (ii) Graph structure learning methods: **GRCN** [43], **IDGL** [9], and **SUBLIME** [23]. (iii) Multi-label node classification methods: **GCN-LPA** [30], **ML-GCN** [11], **LANC** [49], **SMLG** [26], and **LARN** [38]. Note that we equip all models with the GCN backbone for fair comparisons. Due to the traditional workflow of GCN being unsuitable for multi-label node classification, we replace the last activation function from softmax to sigmoid, with the activation value for the classification threshold set to 0.5 for each label. The details of these baselines are left in Appendix B.2.

5.1.3 Implementation Setting. For all models, the embedding size is fixed to 64 for fair node classification. The batch size is set to 1024 for all models and the Adam optimizer is used. Experiments are conducted on the Centos system with NVIDIA RTX-3090 GPUs. The detailed experiment implementations of CorGCN are posed in Appendix B.3. Source code is available at <https://github.com/YuanchenBei/CorGCN>.

5.1.4 Evaluation Metrics. We evaluate the model under the setting for semi-supervised multi-label learning on graphs as previous works [26]. We evaluate the models with seven widely-adopted metrics including *Ranking Loss*, *Hamming Loss*, *Macro-AUC*, *Micro-AUC*, *Macro-AP*, *Micro-AP*, and *Label Ranking AP* (abbreviated as *Ranking*, *Hamming*, *Ma-AUC*, *Mi-AUC*, *Ma-AP*, *Mi-AP*, and *LPAP* in the following subsections, respectively). All these multi-label node classification metrics are clearly defined in [46]. Note that we run all the experiments *five* times with different random seeds and report the average results with standard deviation.

Table 2: Multi-label node classification comparison results in percentage over five trial runs (↑: the higher, the better; ↓: the lower, the better). The best and second-best results in each column are highlighted in bold font and underlined. OOM denotes out-of-memory during the model training. The average ranking is calculated based on the numerical result of each model.

Dataset	Metrics	GCN	GRCN	IDGL	SUBLIME	GCN-LPA	ML-GCN	LANC	SMLG	LARN	CorGCN
Humloc	Ranking (↓)	13.29 ± 1.01	16.84 ± 0.09	19.04 ± 0.13	14.44 ± 0.10	19.58 ± 4.96	<u>13.28 ± 0.70</u>	14.66 ± 0.36	17.82 ± 0.43	13.62 ± 0.46	12.57 ± 0.31
	Hamming (↓)	7.52 ± 0.13	8.37 ± 0.03	8.40 ± 0.02	8.53 ± 0.03	8.30 ± 0.13	<u>7.46 ± 0.22</u>	7.82 ± 0.16	8.55 ± 0.26	8.04 ± 0.29	7.37 ± 0.17
	Ma-AUC (↑)	69.10 ± 2.25	56.27 ± 0.23	50.26 ± 0.94	72.42 ± 0.40	57.91 ± 2.39	72.41 ± 2.47	68.73 ± 1.76	<u>72.64 ± 1.65</u>	71.63 ± 3.30	77.31 ± 1.58
	Mi-AUC (↑)	85.39 ± 1.30	82.85 ± 0.02	78.48 ± 0.17	85.63 ± 0.07	80.97 ± 5.29	<u>87.63 ± 0.52</u>	86.06 ± 0.47	79.06 ± 0.38	86.63 ± 0.50	88.57 ± 0.37
	Ma-AP (↑)	<u>24.65 ± 1.29</u>	14.11 ± 0.12	9.65 ± 0.35	20.14 ± 0.13	13.70 ± 1.72	23.78 ± 1.17	20.94 ± 1.69	21.17 ± 0.61	20.60 ± 0.34	26.91 ± 1.67
	Mi-AP (↑)	46.46 ± 2.06	35.11 ± 0.10	26.31 ± 0.34	36.61 ± 0.13	31.88 ± 7.37	<u>47.76 ± 1.59</u>	42.10 ± 2.61	35.53 ± 0.18	41.30 ± 1.22	48.97 ± 1.30
PCG	LRAP (↑)	64.66 ± 0.80	55.45 ± 0.19	52.04 ± 0.13	59.71 ± 0.25	52.50 ± 3.97	<u>64.91 ± 1.14</u>	61.02 ± 1.45	54.27 ± 0.38	62.46 ± 1.46	65.40 ± 0.83
	Ranking (↓)	27.50 ± 0.58	27.30 ± 0.04	27.97 ± 0.10	27.31 ± 0.16	28.55 ± 0.34	<u>27.11 ± 0.44</u>	28.26 ± 0.51	28.15 ± 0.56	28.45 ± 0.53	25.97 ± 0.57
	Hamming (↓)	13.22 ± 0.30	12.96 ± 0.00	13.20 ± 0.01	13.16 ± 0.01	12.59 ± 0.20	<u>12.56 ± 0.31</u>	12.59 ± 0.18	16.28 ± 0.16	12.65 ± 0.14	12.41 ± 0.23
	Ma-AUC (↑)	62.92 ± 0.84	48.03 ± 0.36	45.60 ± 0.27	57.32 ± 0.41	53.84 ± 0.90	61.48 ± 0.84	58.75 ± 1.05	<u>63.02 ± 0.97</u>	57.00 ± 1.14	64.86 ± 1.05
	Mi-AUC (↑)	71.74 ± 0.38	67.61 ± 0.13	62.86 ± 0.13	70.89 ± 0.09	70.09 ± 0.38	<u>72.03 ± 0.40</u>	71.31 ± 0.56	71.08 ± 0.62	70.02 ± 0.47	74.16 ± 0.61
	Ma-AP (↑)	<u>23.49 ± 0.38</u>	13.12 ± 0.05	12.68 ± 0.12	18.45 ± 0.29	16.01 ± 0.69	20.95 ± 0.91	19.49 ± 0.85	22.73 ± 0.75	16.23 ± 0.66	24.64 ± 0.90
Blogcatalog	Mi-AP (↑)	<u>30.04 ± 0.54</u>	23.43 ± 0.09	21.32 ± 0.28	28.04 ± 0.29	25.50 ± 0.95	29.33 ± 1.21	27.54 ± 0.98	24.77 ± 0.85	24.87 ± 0.59	31.91 ± 1.07
	LRAP (↑)	<u>48.03 ± 1.25</u>	46.42 ± 0.13	46.76 ± 0.24	<u>48.23 ± 0.26</u>	45.37 ± 1.03	48.21 ± 0.57	46.62 ± 1.27	47.58 ± 0.86	45.76 ± 0.99	49.04 ± 0.83
	Ranking (↓)	25.67 ± 0.20	25.69 ± 0.02	25.95 ± 0.01	<u>25.48 ± 0.01</u>	42.19 ± 3.87	25.68 ± 0.25	<u>25.48 ± 0.12</u>	26.61 ± 0.86	25.67 ± 0.27	25.42 ± 0.23
	Hamming (↓)	3.58 ± 0.03	3.57 ± 0.00	<u>3.56 ± 0.00</u>	3.59 ± 0.00	3.58 ± 0.03	3.58 ± 0.03	3.55 ± 0.03	3.66 ± 0.04	3.58 ± 0.03	3.58 ± 0.03
	Ma-AUC (↑)	50.59 ± 0.51	48.19 ± 0.04	47.08 ± 0.01	50.52 ± 0.02	50.19 ± 2.06	50.94 ± 1.17	<u>53.29 ± 0.52</u>	51.70 ± 0.32	50.52 ± 1.07	54.48 ± 0.52
	Mi-AUC (↑)	73.80 ± 0.20	69.22 ± 0.05	65.28 ± 0.01	72.37 ± 0.01	56.74 ± 2.32	73.85 ± 0.17	<u>74.11 ± 0.06</u>	71.75 ± 0.86	73.73 ± 0.47	74.15 ± 0.21
PPI	Ma-AP (↑)	4.16 ± 0.14	3.67 ± 0.01	3.67 ± 0.01	3.91 ± 0.02	4.10 ± 0.24	4.16 ± 0.14	5.07 ± 0.20	4.13 ± 0.13	4.15 ± 0.48	<u>4.61 ± 0.15</u>
	Mi-AP (↑)	9.41 ± 0.15	7.97 ± 0.01	7.13 ± 0.02	8.89 ± 0.02	4.72 ± 0.53	9.42 ± 0.15	11.71 ± 0.55	5.54 ± 1.33	9.38 ± 1.22	<u>9.65 ± 0.21</u>
	LRAP (↑)	27.88 ± 0.15	27.30 ± 0.00	27.82 ± 0.04	27.81 ± 0.00	17.58 ± 3.93	27.87 ± 0.18	28.94 ± 1.00	25.37 ± 0.27	28.10 ± 0.92	<u>28.32 ± 0.29</u>
	Ranking (↓)	<u>18.32 ± 0.16</u>	25.73 ± 0.01	OOM	25.44 ± 0.01	25.33 ± 0.08	20.05 ± 0.62	18.44 ± 0.14	18.38 ± 0.30	19.85 ± 0.18	16.17 ± 0.22
	Hamming (↓)	22.63 ± 0.14	26.28 ± 0.00	OOM	26.24 ± 0.00	25.60 ± 0.18	23.83 ± 0.14	<u>21.85 ± 0.28</u>	31.06 ± 0.32	23.41 ± 0.13	20.79 ± 0.29
	Ma-AUC (↑)	73.06 ± 0.21	44.65 ± 0.02	OOM	51.82 ± 0.12	58.86 ± 0.29	70.26 ± 1.12	<u>74.24 ± 0.42</u>	74.04 ± 0.41	70.21 ± 0.24	77.54 ± 0.42
Delve	Mi-AUC (↑)	80.19 ± 0.14	67.84 ± 0.01	OOM	70.24 ± 0.02	72.48 ± 0.08	78.37 ± 0.72	<u>81.15 ± 0.26</u>	80.46 ± 0.26	78.50 ± 0.13	83.35 ± 0.29
	Ma-AP (↑)	54.70 ± 0.42	29.34 ± 0.01	OOM	34.07 ± 0.01	39.06 ± 0.30	50.69 ± 1.54	<u>57.19 ± 0.15</u>	55.71 ± 0.70	51.16 ± 0.61	61.33 ± 0.67
	Mi-AP (↑)	67.64 ± 0.43	55.66 ± 0.01	OOM	56.52 ± 0.01	58.68 ± 0.37	64.98 ± 1.32	<u>69.59 ± 0.27</u>	63.11 ± 0.62	65.57 ± 0.54	72.35 ± 0.58
	LRAP (↑)	68.82 ± 0.39	62.19 ± 0.01	OOM	62.28 ± 0.01	62.18 ± 0.40	66.68 ± 1.01	<u>68.87 ± 0.35</u>	63.68 ± 0.61	67.96 ± 0.36	71.40 ± 0.34
	Ranking (↓)	3.46 ± 0.04	OOM	OOM	6.10 ± 0.53	OOM	3.33 ± 0.01	OOM	5.07 ± 0.02	1.46 ± 0.02	<u>2.41 ± 0.02</u>
	Hamming (↓)	3.16 ± 0.01	OOM	OOM	30.26 ± 0.51	OOM	2.99 ± 0.01	OOM	7.50 ± 0.01	1.67 ± 0.07	<u>2.48 ± 0.01</u>
Average Rank	Ma-AUC (↑)	94.98 ± 0.14	OOM	OOM	58.37 ± 0.74	OOM	95.71 ± 0.04	OOM	95.26 ± 0.09	97.85 ± 0.05	<u>96.84 ± 0.10</u>
	Mi-AUC (↑)	96.55 ± 0.06	OOM	OOM	57.58 ± 0.04	OOM	97.29 ± 0.02	OOM	95.11 ± 0.04	98.39 ± 0.07	<u>98.05 ± 0.02</u>
	Ma-AP (↑)	64.28 ± 0.39	OOM	OOM	11.60 ± 0.07	OOM	65.20 ± 0.07	OOM	65.67 ± 0.31	80.16 ± 0.32	<u>72.80 ± 0.26</u>
	Mi-AP (↑)	78.72 ± 0.07	OOM	OOM	10.12 ± 0.07	OOM	81.78 ± 0.09	OOM	68.97 ± 0.25	90.44 ± 0.12	<u>86.12 ± 0.03</u>
	LRAP (↑)	85.57 ± 0.10	OOM	OOM	29.44 ± 0.40	OOM	86.12 ± 0.05	OOM	76.47 ± 0.27	92.09 ± 0.21	<u>89.20 ± 0.01</u>
	Average Rank	3.86	7.79	8.52	6.17	7.79	3.57	<u>3.50</u>	5.94	4.46	1.37

5.2 Main Results (RQ1)

In this subsection, we compare our proposed CorGCN with nine state-of-the-art baseline models on the five experimental datasets. The comparison results on seven metrics are reported in Table 2, and corresponding Friedman statistics are shown in Table 3, with the following observations:

- **CorGCN can achieve significant improvements over state-of-the-art methods on all experimental datasets.** From the table, we can observe that our proposed CorGCN can generally achieve the best performance among the seven metrics on average. Specifically, CorGCN outperforms the best baseline by 6.43%, 2.92%, 2.23%, and 4.45% for Macro-AUC on Humloc, PCG, Blogcatalog, and PPI, respectively. The significance of model gains was further validated by Friedman statistics and the Bonferroni-Dunn test. These results verify that designing the correlation decomposed graph and the equipped CorGCN contributes to achieving better multi-label node classification performance.

- **Graph structure learning models do not bring about performance gains.** Comparing three graph structure learning models (GRCN, IDGL, SUBLIME) with the simply adopted GCN, we can find that these methods fail to yield substantial benefits in multi-label node classification scenarios and, in many instances, may even lead to negative outcomes. For example, all three models negatively impact the GCN model performance in terms of Micro-AUC, Macro-AP, and Micro-AP. This indicates that directly learning graph structures without considering multi-label correlations is less effective for the multi-label node classification downstream task.
- **The models with multi-label modeling generally become the best baselines.** We can observe from the comparison results that LANC and ML-GCN achieve the top-2 performance among the other baseline models. Their shared advantage lies in the explicit modeling of multi-label relationships. However, due to the unified message passing with ambitious information, they are less optimal compared to the CorGCN. Our proposed CorGCN

Table 3: Friedman statistics F_F on each evaluation metric with corresponding p -value (all < 0.05 level).

Evaluation Metric	F_F	p -value
Ranking Loss	17.59	0.0015
Hamming Loss	21.97	0.0002
Macro-AUC	11.80	0.0189
Micro-AUC	13.41	0.0094
Macro-AP	14.85	0.0050
Micro-AP	16.54	0.0024
Label Ranking AP	16.54	0.0024

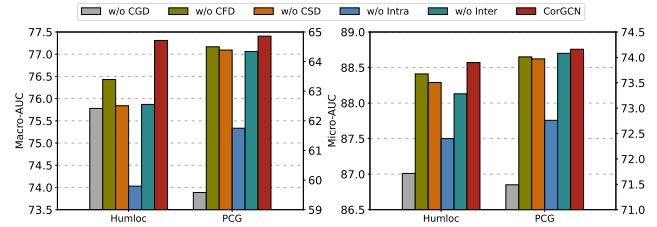
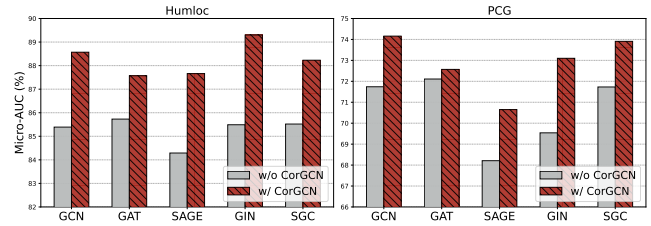
explicitly considers the label correlation and addresses the ambitious message passing with enhanced information aggregation over the correlation-aware decomposed graph, which achieves the best performance.

We have also assessed the convergence of CorGCN, with the results detailed in Appendix C. The results indicate that our model adeptly balances effectiveness and efficiency, also exhibiting faster convergence compared to the pure backbone.

5.3 Ablation Study (RQ2)

To verify the effectiveness of the key components in CorGCN, we conduct ablation studies by comparing CorGCN with its five variants: (1) *w/o CGD* replaces the decomposed graphs by the origin input graph. (2) *w/o CFD* removes the correlation-aware decomposition of node features, replacing them with original features for graph learning. (3) *w/o CSD* excludes the correlation-aware structure decomposition, replacing them with correlated node features and origin structure for CorGCN message passing. (4) *w/o Intra* ignores the intra-label message passing and merges all views of graphs for unified message passing. (5) *w/o Inter* ignore the inter-label correlation propagation without considering message correlations. From Figure 3, we can have following observations:

- **Effectiveness of the correlation decomposed graph.** The severe performance decline of *w/o CGD*, along with the diminished efficacy of *w/o CFD* and *w/o CSD*, demonstrates that the original graph topology and node features are suboptimal to multi-label node classification. The proposed correlation decomposed graph (CDG) with both correlation decomposed node features and topology structure can effectively benefit the classification.
- **Effectiveness of the correlation decomposed graph convolution.** Based on the effective correlation decomposed graph, the lack of carefully designed decomposed graph convolution modules (*w/o Intra* and *w/o Inter*) both results in inferior classification performance than CorGCN, which demonstrates that modeling intra-label message and passing inter-label correlation is meaningful. Further, removing the intra-label message passing with unified message passing has a greater impact on CorGCN, indicating that the issue of feature and topology ambitiousness indeed brings negative effects during the message passing.

**Figure 3: Ablation study on CorGCN with its five variants.****Figure 4: Generalization study results of CorGCN in different GNN message passings.**

5.4 Generalization Study (RQ3)

To investigate the generalization of the proposed CorGCN method, we evaluate the performance of CorGCN with different GNN backbones. Specifically, we utilize the representative GCN-like [18], GAT-like [29], SAGE-like [16], GIN-like [40], and SGC-like [34] for CorGCN in Eq.(14) of message passing, respectively. The study results are illustrated in Figure 4.

From the results, we can find that equipped with CorGCN (*w/ CorGCN*) is consistently significantly better than message passing with vanilla backbones (*w/o CorGCN*), which brings 3.50% and 3.15% average Micro-AUC improvements on Humloc and PCG, respectively. The improvements brought by CorGCN to other backbones are greater compared to GAT. This is because GAT implicitly differentiates the neighborhoods by considering the importance of different nodes. However, the message passing paradigms in these backbones all remain ambitious in multi-label node classification. CorGCN explicitly addresses this issue and thus exhibits better performance. The results demonstrate both the effectiveness and generalization of CorGCN in different message passing functions.

5.5 Case Study (RQ4)

To further analyze the performance of our CorGCN with the graph decomposition, we conduct a case study that analyzes the performance of each fine-grained class. The comparison between the original GCN and the top-performed baseline MLGCN over the Humloc dataset is as follows. From the results in Figure 5, we can have the following findings: (i) Under the decomposition paradigm, CorGCN can achieve significantly better performance among all the fine-grained classes than the pure backbone. (ii) Although MLGCN is the top-performing baseline on Humloc, due to the lack of ambiguous discrimination ability with a unified convolution process, it will lead to 5/14 classes with performance dropping than the original GCN.

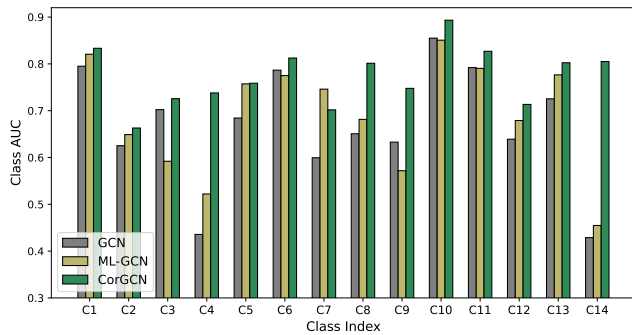


Figure 5: Case study of class AUC performance on Humloc.

Therefore, based on the case study results, we can find that the decomposition paradigm of CorGCN is helpful in better understanding each class characteristic over the multi-label graph with ambiguous features and topology.

5.6 Efficiency Study (RQ5)

To further evaluate whether CorGCN can balance the effectiveness and efficiency, we conduct the efficiency study of CorGCN compared with the top-3 performed baselines on both training and inference stage in Table 4.

From the results, we can find that our CorGCN can generally balance effectiveness and efficiency. For the training stage, the end-to-end label correlation modeling strategy employed by CorGCN yields an efficiency that is second to that of pure GCN. The negative sampling process costs additional time for ML-GCN during training. Then, during the inference phase, CorGCN achieves computational efficiency that is lower than top-3 performing baselines due to the multi-label structure decomposition while demonstrating a good enhancement in performance compared to these models.

Table 4: Efficiency study results on Humloc and PCG datasets each epoch in seconds.

Dataset	Stage	GCN	ML-GCN	LANC	CorGCN
Humloc	Training	0.37	15.79	1.18	0.76
	Inference	0.37	0.49	0.48	0.57
PCG	Training	0.39	18.15	1.21	0.84
	Inference	0.30	0.62	0.39	0.64

5.7 Parameter Study (RQ6)

5.7.1 Effect of Parameter λ . We investigate the effect of λ for density controlling in correlation decomposed structure learning with the range from 1 to 13 with a step size of 2 as Figure 6. From the results, we can observe that a too-small value of λ will cause poor performance, which indicates the information is too sparse and insufficient. Furthermore, the suitable value of λ for Humloc is larger than the value for PCG, one possible reason is that the λ of CorGCN should be set larger on more sparse graphs.

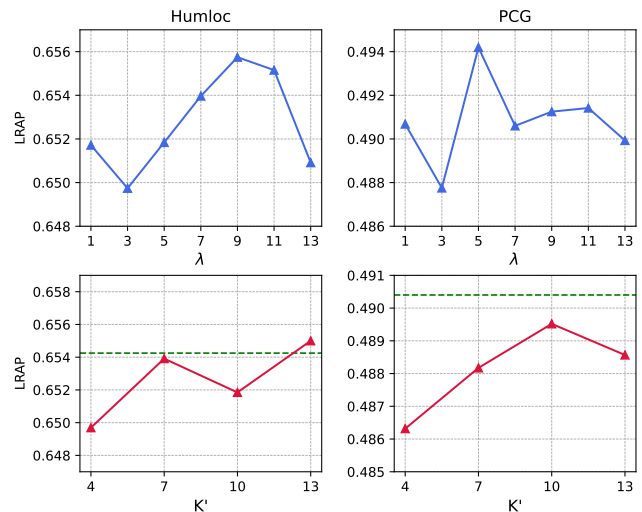


Figure 6: Parameter study results of λ and K' .

5.7.2 Effect of Macro Label Prototype Number K' . We also evaluate the impact of different macro label numbers K' in Sec. 4.4 to CorGCN. From Figure 6, we can find that the too-small cluster number will lead to too coarse label segmentation and result in relatively poor results. In addition, choosing a relatively appropriate number of macro labels, such as 13 for Humloc and 10 for PCG, can achieve model performance close to or even equivalent to the performance of original label numbers (the green dashed line). More hyperparameter studies on γ for CorGCN can be found in Appendix D.

6 Conclusion

In this paper, we propose Correlation-Aware Graph Convolutional Networks (CorGCN) for multi-label node classification. Specifically, aiming to address the ambitious feature and topology in the original graph in multi-label scenarios, we introduce more suitable Correlation-aware Decomposed Graphs (CDG). Then, based on the CDG, we further design a novel Correlation-enhanced Graph Convolution with inter-label message passing and intra-label correlation propagation. Extensive experiments on five multi-label node classification datasets and in-depth analyses in multiple perspectives illustrate the effectiveness and strengths of CorGCN. Future works can be conducted on the optimization of model inference efficiency and the exploration of downstream tasks of multi-label node classification, such as multi-label graph anomaly detection.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62476245, No. 62372408), Zhejiang Provincial Natural Science Foundation of China (Grant No. LTGG23F030005). Carl Yang was not supported by any grants from China.

References

- [1] Junwen Bai, Shufeng Kong, and Carla Gomes. 2021. Disentangled variational autoencoder based multi-label classification with covariance-aware multivariate

- probit model. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 4313–4321.
- [2] Junwen Bai, Shufeng Kong, and Carla P Gomes. 2022. Gaussian mixture variational autoencoder with contrastive learning for multi-label classification. In *International Conference on Machine Learning*. PMLR, 1383–1398.
 - [3] Yuanchen Bei, Sheng Zhou, Qiaoyu Tan, Hao Xu, Hao Chen, Zhao Li, and Jiajun Bu. 2023. Reinforcement neighborhood selection for unsupervised graph anomaly detection. In *2023 IEEE International Conference on Data Mining (ICDM)*. IEEE, 11–20.
 - [4] Smriti Bhagat, Graham Cormode, and S Muthukrishnan. 2011. Node classification in social networks. *Social network data analytics* (2011), 115–148.
 - [5] Jasmin Bogatinovski, Ljupčo Todorovski, Sašo Džeroski, and Drago Kocev. 2022. Comprehensive comparative study of multi-label classification methods. *Expert Systems with Applications* 203 (2022), 117215.
 - [6] Matthew R Boutell, Jiebo Luo, Xipeng Shen, and Christopher M Brown. 2004. Learning multi-label scene classification. *Pattern recognition* 37, 9 (2004), 1757–1771.
 - [7] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. 2018. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE transactions on knowledge and data engineering* 30, 9 (2018), 1616–1637.
 - [8] Hao Chen, Yuanchen Bei, Qijie Shen, Yue Xu, Sheng Zhou, Wenbing Huang, Feiran Huang, Senzhang Wang, and Xiao Huang. 2024. Macro graph neural networks for online billion-scale recommender systems. In *Proceedings of the ACM Web Conference 2024*. 3598–3608.
 - [9] Yu Chen, Lingfei Wu, and Mohammed Zaki. 2020. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. *Advances in neural information processing systems* 33 (2020), 19314–19326.
 - [10] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. 2019. Learning a deep convnet for multi-label classification with partial labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 647–657.
 - [11] Kaisheng Gao, Jing Zhang, and Cangqi Zhou. 2019. Semi-supervised Graph Embedding for Multi-label Graph Node Classification. In *Web Information Systems Engineering – WISE 2019*. Springer International Publishing, Cham, 555–567.
 - [12] Ziqi Gao, Chenran Jiang, Jiawen Zhang, Xiaosen Jiang, Lanqing Li, Peilin Zhao, Huanming Yang, Yong Huang, and Jia Li. 2023. Hierarchical graph learning for protein–protein interaction. *Nature Communications* 14, 1 (2023), 1093.
 - [13] Weifeng Ge, Sibe Yang, and Yizhou Yu. 2018. Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1277–1286.
 - [14] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
 - [15] Greg Hamerly and Charles Elkan. 2003. Learning the k in k-means. *Advances in neural information processing systems* 16 (2003).
 - [16] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, Vol. 30.
 - [17] Xiaoyang Jing and Jinbo Xu. 2021. Fast and effective protein model refinement using deep graph neural networks. *Nature computational science* 1, 7 (2021), 462–469.
 - [18] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
 - [19] Zhixun Li, Xin Sun, Yifan Luo, Yanqiao Zhu, Dingshuo Chen, Yingtao Luo, Xi-angxin Zhou, Qiang Liu, Shu Wu, Liang Wang, et al. 2024. GSLB: the graph structure learning benchmark. *Advances in Neural Information Processing Systems* 36 (2024).
 - [20] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*. 2980–2988.
 - [21] Jingzhou Liu, Wei-Cheng Chang, Yuexin Wu, and Yiming Yang. 2017. Deep learning for extreme multi-label text classification. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*. 115–124.
 - [22] Weiwei Liu, Haobo Wang, Xiaobo Shen, and Ivor W Tsang. 2021. The emerging trends of multi-label learning. *IEEE transactions on pattern analysis and machine intelligence* 44, 11 (2021), 7955–7974.
 - [23] Yixin Liu, Yu Zheng, Daokun Zhang, Hongxu Chen, Hao Peng, and Shirui Pan. 2022. Towards unsupervised deep graph structure learning. In *Proceedings of the ACM Web Conference 2022*. 1392–1403.
 - [24] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. 2009. Classifier chains for multi-label classification. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2009, Bled, Slovenia, September 7–11, 2009, Proceedings, Part II* 20. Springer, 254–269.
 - [25] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. 2020. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems* 33 (2020), 12559–12571.
 - [26] Zixing Song, Ziqiao Meng, Yifei Zhang, and Irwin King. 2021. Semi-supervised multi-label learning for graph-structured data. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1723–1733.
 - [27] Peijie Sun, Le Wu, and Meng Wang. 2018. Attentive recurrent social recommendation. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 185–194.
 - [28] Adane Nega Tarekegn, Mohib Ullah, and Faouzi Alaya Cheikh. 2024. Deep learning for multi-label learning: A comprehensive survey. *arXiv preprint arXiv:2401.16549* (2024).
 - [29] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *The International Conference on Learning Representations*.
 - [30] Hongwei Wang and Jure Leskovec. 2021. Combining graph convolutional neural networks and label propagation. *ACM Transactions on Information Systems (TOIS)* 40, 4 (2021), 1–27.
 - [31] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. 2016. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2285–2294.
 - [32] Kaixiang Wang, Ming Yang, Wanqi Yang, and YiLong Yin. 2018. Deep correlation structure preserved label space embedding for multi-label classification. In *Asian Conference on Machine Learning*. PMLR, 1–16.
 - [33] Ya Wang, Dongliang He, Fu Li, Xiang Long, Zhichao Zhou, Jinwen Ma, and Shilei Wen. 2020. Multi-label classification with label graph superimposing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12265–12272.
 - [34] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*. PMLR, 6861–6871.
 - [35] Jian Wu, Victor S Sheng, Jing Zhang, Hua Li, Tetiana Dadakova, Christine Leon Swisher, Zhiming Cui, and Pengpeng Zhao. 2020. Multi-label active learning algorithms for image classification: Overview and future promise. *ACM Computing Surveys (CSUR)* 53, 2 (2020), 1–35.
 - [36] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37.
 - [37] Lin Xiao, Xin Huang, Boli Chen, and Liping Jing. 2019. Label-specific document representation for multi-label text classification. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 466–475.
 - [38] Lin Xiao, Pengyu Xu, Liping Jing, Uchenna Akujuobi, and Xiangliang Zhang. 2022. Semantic guide for semi-supervised few-shot multi-label node classification. *Information Sciences* 591 (2022), 235–250.
 - [39] Shunxin Xiao, Shiping Wang, Yuanfei Dai, and Wenzhong Guo. 2022. Graph neural networks in node classification: survey and evaluation. *Machine Vision and Applications* 33 (2022), 1–19.
 - [40] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*.
 - [41] Vacit Oguz Yazici, Abel Gonzalez-Garcia, Arnau Ramisa, Bartłomiej Twardowski, and Joost van de Weijer. 2020. Orderless recurrent models for multi-label classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13440–13449.
 - [42] Chih-Kuan Yeh, Wei-Chieh Wu, Wei-Jen Ko, and Yu-Chiang Frank Wang. 2017. Learning deep latent space for multi-label classification. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 31.
 - [43] Donghan Yu, Ruohong Zhang, Zhengbao Jiang, Yuexin Wu, and Yiming Yang. 2021. Graph-revised convolutional network. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part III*. Springer, 378–393.
 - [44] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *International Conference on Learning Representations*.
 - [45] Jiawei Zhang and Philip S Yu. 2018. Broad learning: An emerging area in social network analysis. *ACM SIGKDD Explorations Newsletter* 20, 1 (2018), 24–50.
 - [46] Min-Ling Zhang and Zhi-Hua Zhou. 2013. A review on multi-label learning algorithms. *IEEE transactions on knowledge and data engineering* 26, 8 (2013), 1819–1837.
 - [47] Tianqi Zhao, Thi Ngan Dong, Alan Hanjalic, and Megha Khosla. 2023. Multi-label Node Classification On Graph-Structured Data. *Transactions on Machine Learning Research* (2023).
 - [48] Wenting Zhao, Shufeng Kong, Junwen Bai, Daniel Fink, and Carla Gomes. 2021. Hot-vae: Learning high-order label correlation for multi-label classification via attention-based variational autoencoders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 15016–15024.
 - [49] Cangqi Zhou, Hui Chen, Jing Zhang, Qianmu Li, Dianming Hu, and Victor S Sheng. 2021. Multi-label graph node classification with label attentive neighborhood convolution. *Expert Systems with Applications* 180 (2021), 115063.

A Methodology Details

A.1 Analysis of the Objective Function on your Multi-Label Estimator

The mutual information estimator ϕ and likelihood maximizing decoder θ work collaboratively for node-label correlation and label-label correlation modeling as

$$\arg \max_{\phi, \theta} \underbrace{I_{\phi}(E^x, E^l)}_{\text{Mutual Information}} + \underbrace{\mathcal{L}_{\theta}(\theta; E^x) + \mathcal{L}_{\theta}(\theta; E^l)}_{\text{Likelihood}}. \quad (24)$$

Analysis of the effectiveness of mutual information estimator ϕ for multi-label correlation modeling can be found in previous works [2, 48]. Then, The likelihood maximizing decoder θ can be utilized as a likelihood estimator of node feature and label prototype with the entropy KL divergence with focal enhancement [20] for target labels as follows:

$$\mathcal{L}_{lm} = \frac{1}{2|\mathcal{V}_L|} \sum_{v_i \in \mathcal{V}_L} \sum_{k=1}^K \rho_k [\underbrace{(1 - p_k(E_i^x | \theta))^y \log p_k(E_i^x | \theta)}_{\text{Likelihood Estimation with Node Feature}} + \underbrace{(1 - p_k(\mathbf{y}_i \cdot E^l | \theta))^y \log p_k(\mathbf{y}_i \cdot E^l | \theta)}_{\text{Likelihood Estimation with Label Prototype}}]. \quad (25)$$

A.2 Complexity Analysis

The overall complexity of CorGCN arises from the following stages. During the correlation-aware graph decomposition, the calculation involves the projection process, incurring a computational cost of $\Theta(K \cdot n^2)$. For the Correlation-Enhanced graph convolution, we postulate that the message passing process within a graph comprising n nodes entails a cost of $\Theta_{mp}(n)$, which varies according to the specific GCN backbone employed. During the intra-label message passing, due to the process can be paralleled across K label views, the overall expense of message passing is quantified as $\Theta_{mp}(n)$. Then, the cost associated with inter-label correlation propagation is $\Theta(K^2 \cdot n)$. Thus, the overall time complexity of the process can be expressed as $T(n) = \Theta(K(n^2 + \Theta_{mp}(n) + K \cdot n))$.

As outlined in Sec. 4.4, this complexity can be reduced to $T(n) = \Theta(K'(n^2 + \Theta_{mp}(n) + K'n))$ for the large label space, where $K' \ll K$, denoting a significant reduction in computational cost. In practice, the entire process can be batch-processed, where the similarity is calculated among the nodes within each batch, thus making the time complexity of each batch B to $T(B) = \Theta(K'(B^2 + \Theta_{mp}(B) + K'B))$.

A.3 Pseudocode of CorGCN

The overall pipeline of CorGCN includes (1) correlation-aware graph learning with feature and topology decompositions, and (2) correlation-aware graph convolution with intra-label message passing and inter-label correlation propagation. The general process of this pipeline is described as the pseudocode in Algorithm 1.

B Experimental details

B.1 Dataset Details

We adopt five publicly available datasets for comprehensive evaluation of our proposed CorGCN. The detailed description is as follows:

Algorithm 1 The overall workflow of CorGCN.

Input: Feature matrix X ; Label prototype E^l ; Graph topology A^0 ; Number of GNN layers N_{gnn} .

Output: Predicted labels $\hat{\mathbf{y}}$.

- 1: $E^x \leftarrow X$; // Compute transformed features
 - 2: $\mathcal{L}_{cmi} \leftarrow \{E^x, E^l\}$ using Eq.(4); // Compute contrastive loss
 - 3: $\mathcal{L}_{lm} \leftarrow \{E^x, E^l, \mathbf{y}\}$ using Eq.(5) and Eq.(6); // Compute likelihood loss
 - 4: $E^{proj} \leftarrow \{E^x, E^l\}$ using Eq.(7), (8); // Compute projected features as node representations
 - 5: $E^{sd} \leftarrow \{E^{proj}, A\}$ using Eq.(9); // Aggregate the node representations from neighborhood
 - 6: $CDG \leftarrow \{E^{sd}, E^{proj}, A\}$ using Eq.(10) - Eq.(13); // Learn the graphs by aggregated representations
 - 7: **for** $epoch \in 1, 2, \dots, N_{gnn}$ **do**
 - 8: $\hat{Z}^{(l)} \leftarrow \{Z^{(l-1)}, CDG\}$ using Eq.(14); // Update the node representations by message passing
 - 9: $Z^{(l)} \leftarrow \{\hat{Z}^{(l)}, E^l\}$ using Eq.(15) - (16); // Update the node representations by correlation propagation
 - 10: **end for**
 - 11: $\hat{\mathbf{y}} \leftarrow \{Z^{(E)}, E^l\}$ using Eq.(17) - (18); // Predict the probabilities
 - 12: $\mathcal{L}_{cls} \leftarrow \{\hat{\mathbf{y}}, \mathbf{y}\}$ using Eq.(19); // Compute the classification loss
 - 13: $\mathcal{L} \leftarrow \mathcal{L}_{cls} + \alpha \mathcal{L}_{cmi} + \beta \mathcal{L}_{lm}$;
 - 14: **if** *Training* **then**
 - 15: Back-propagate \mathcal{L} to update model parameters.
 - 16: **end if**
-

Humloc¹ [47] is a human protein subcellular location prediction dataset, consisting of 3,106 nodes and 18,496 edges. Each node may have one or more labels in 14 possible locations. The edge data is derived from protein-protein interactions obtained from the open-sourced database. An edge exists between two nodes in the graph if there is an interaction between the respective proteins. **PCG**¹ [47] is a protein phenotype prediction dataset with 3,233 nodes into 15 classes and 37,351 edges between them. Each node represents a protein with 32-dimensional features. Each edge between a pair of protein nodes is their functional interaction. The multi-label that each node is associated with is its correlated phenotypes, of which a phenotype is any observable characteristic or trait of a disease. The correspondence edge between protein and phenotype is retrieved from the open-sourced database. **Blogcatalog**² [49] is a social network dataset with 10,312 nodes and 333,983 edges, which each node represents a blogger at the platform and each edge represents the contact relationship between a pair of nodes. The labels denote the social interest groups a blogger node is a part of, forming a total of 39 interest groups. **PPI**³ [44] serves as a dataset with a large label space for protein-protein interaction networks, comprising 14,755 nodes (representing proteins) and 225,270 edges (indicating functional interactions). The dataset is annotated with gene ontology categories as labels, encompassing a total of 121 unique sets. **Delve**⁴ [38] is a large-scale citation network dataset

¹<https://github.com/Tianqi-py/MLGNC>

²https://figshare.com/articles/dataset/BlogCatalog_dataset/11923611

³<https://github.com/GraphSAINT/GraphSAINT>

⁴<https://www.dropbox.com/sh/tg8nclx5gpcbtmo/AADQc0YFpuVeqXhaNuUK4MMba?dl=0>

with 1,229,280 nodes and 4,322,275 edges, which each node is a paper and each edge represents a citation relationship between a pair of papers. The labels denote the research fields that papers belong to, forming a total of 20 categories.

B.2 Baseline Details

We compare our proposed CorGCN with nine representative state-of-the-art GNN models as follows: **GCN** [18] is a widely recognized model in the domain of graph neural networks, grounded in spectral theory. **GRCN** [43] incorporates a module based on GCN specifically engineered for the prediction of missing edges and the adjustment of edge weights, which is optimized with downstream tasks. **IDGL** [9] mutually enhances the graph structure and node embeddings: the premise is to iteratively improve the graph structure through enhanced node embeddings and, reciprocally, to refine node embeddings leveraging an optimized graph structure. **SUBLIME** [23] uses a contrastive loss to maximize the agreement between the anchor graph and the learned graph, which is further enhanced with a bootstrapping mechanism. **GCN-LPA** [30] unifies GCNs and label propagation in the same learning framework, which we utilized to propagate the multi-label simultaneously. **ML-GCN** [11] employs a strategy to generate and utilize a label matrix for capturing node-label and label-label correlations through a relaxed skip-gram model in a unified vector space. **LANC** [49] develops a label attention module that employs an additive attention mechanism to synergize input and output contextual representations. **SMLG** [26] comprises a joint variational representation module that generates node and label embeddings and a confidence-rated margin ranking module that detects second-order label correlations and refines the decision boundaries. **LARN** [38] designs a novel label correlation scanner that adaptively captures label relationships, extracting vital encoded information for the generation of comprehensive representations.

B.3 Implementation Details

We explore the combination space of the following hyper-parameters lr , λ , K' , γ and list the values finally chosen in Table 5 as the setting of performing the main results. Specifically, the learning rate lr is selected from $\{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.02, 0.05\}$; λ that controls the density of graphs is tuned from 1 to 20; the macro label prototype number K' is searched from 1 to 20; and γ that controls the focusing factor in \mathcal{L}_{lm} is chosen from $\{1.0, 1.5, 2.0, 2.5, 3.0\}$. Besides, to compare fairly, the hidden dim d is set to 64, and the dropout rate of the network is set to 0.3. To balance the proportions between \mathcal{L}_{cls} , \mathcal{L}_{cmi} and \mathcal{L}_{lm} , we calculate the loss balancing parameters $\alpha = |\frac{\mathcal{L}_{cls}}{3\mathcal{L}_{cmi}}|$ and $\beta = |\frac{\mathcal{L}_{cls}}{3\mathcal{L}_{lm}}|$ adaptively.

Table 5: The values of hyper-parameters used in CorGCN.

Dataset	lr	λ	K'	γ
Humloc	0.001	7	\	2.0
PCG	0.001	19	\	2.0
Blogcatalog	0.001	5	\	2.0
PPI	0.02	5	20	2.0
Delve	0.001	5	10	2.0

C Convergence Analysis

To further analyze the convergence of our CorGCN, we provide the figures of training classification loss changing over epochs and validation metrics changing over epochs as Figure 7. From the results, we can find that CorGCN can achieve a faster and better convergence than the pure backbone for multi-label node classification with lower training loss and higher validation performance. This is due to the design of CorGCN with the correlation-aware multi-label graph decomposition and convolution can help the backbone better and faster understand the multi-label relationships and important information during message passing in the training process.

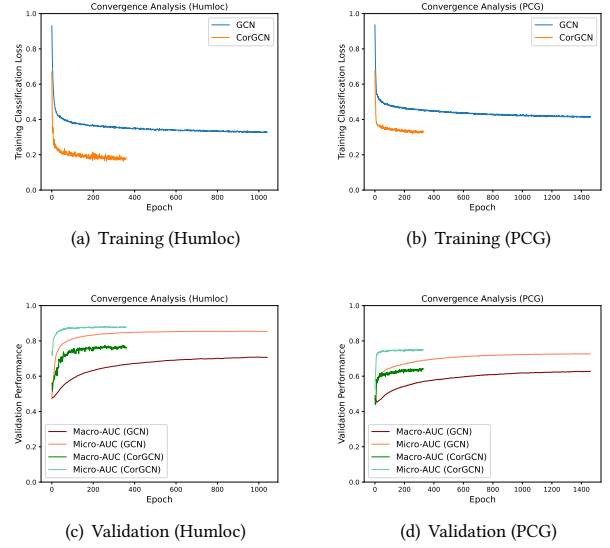


Figure 7: Convergence analysis on (a)-(b) the training classification loss. (c)-(d) the validation set performance.

D More Parameter Study Results

We explored the impact of γ here. Parameter γ amplifies the loss weight of misclassified samples, thereby increasing the model’s attention to them. As shown in Figure 8, it can be observed that the model achieves the best performance on the LARP metric when γ is set to 1.5. However, as γ further increases, it may cause the model to overly focus on misclassified samples, neglecting the well-classified samples, resulting in a slight decrease in model performance.

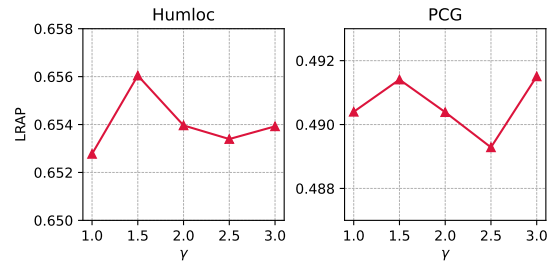


Figure 8: Parameter study results on γ of CorGCN.