



RAHNet: Retrieval Augmented Hybrid Network for Long-tailed Graph Classification

Zhengyang Mao*

National Key Laboratory for
Multimedia Information Processing,
School of Computer Science,
Peking University
Beijing, China
zhengyang.mao@stu.pku.edu.cn

Wei Ju*

National Key Laboratory for
Multimedia Information Processing,
School of Computer Science,
Peking University
Beijing, China
juwei@pku.edu.cn

Yifang Qin

National Key Laboratory for
Multimedia Information Processing,
School of Computer Science,
Peking University
Beijing, China
qinyifang@pku.edu.cn

Xiao Luo[†]

Department of Computer Science,
University of California, Los Angeles
Los Angeles, USA
xiaoluo@cs.ucla.edu

Ming Zhang[†]

National Key Laboratory for
Multimedia Information Processing,
School of Computer Science,
Peking University
Beijing, China
mzhang_cs@pku.edu.cn

ABSTRACT

Graph classification is a crucial task in many real-world multimedia applications, where graphs can represent various multimedia data types such as images, videos, and social networks. Previous efforts have applied graph neural networks (GNNs) in balanced situations where the class distribution is balanced. However, real-world data typically exhibit long-tailed class distributions, resulting in a bias towards the head classes when using GNNs and limited generalization ability over the tail classes. Recent approaches mainly focus on re-balancing different classes during model training, which fails to explicitly introduce new knowledge and sacrifices the performance of the head classes. To address these drawbacks, we propose a novel framework called Retrieval Augmented Hybrid Network (RAHNet) to jointly learn a robust feature extractor and an unbiased classifier in a decoupled manner. In the feature extractor training stage, we develop a graph retrieval module to search for relevant graphs that directly enrich the intra-class diversity for the tail classes. Moreover, we innovatively optimize a category-centered supervised contrastive loss to obtain discriminative representations, which is more suitable for long-tailed scenarios. In the classifier fine-tuning stage, we balance the classifier weights with two weight regularization techniques, *i.e.*, Max-norm and weight decay. Experiments on various popular benchmarks verify the superiority of the proposed method against state-of-the-art approaches.

*Both authors contributed equally to this research.

[†]Corresponding authors.

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](https://www.acm.org/permissions).

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0108-5/23/10...\$15.00

<https://doi.org/10.1145/3581783.3612360>

CCS CONCEPTS

• Information systems → Data mining; • Computing methodologies → Neural networks.

KEYWORDS

Graph Classification, Long-tailed Learning, Graph Neural Networks, Graph Retrieval

ACM Reference Format:

Zhengyang Mao, Wei Ju, Yifang Qin, Xiao Luo, and Ming Zhang. 2023. RAHNet: Retrieval Augmented Hybrid Network for Long-tailed Graph Classification. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*, October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3581783.3612360>

1 INTRODUCTION

Graph classification, which captures the graph-level properties to predict classes for graphs, is a fundamental task in data mining with many real-world multimedia applications across various domains [26, 45]. For example, social networks [69], images [8], and knowledge graphs [4] can all be represented as graphs, and graph classification can be used in multimedia tasks such as object recognition, semantic segmentation, and recommender systems. As a powerful technique for graph representation learning, graph neural networks (GNNs) [13, 17, 25, 54] have achieved remarkable success in various applications, such as graph classification [19, 21, 35, 36], novel drug discovery [1, 18, 37, 48], traffic forecasting [10, 29, 72], and recommender systems [20, 46, 47, 61, 62]. GNNs are designed to propagate and aggregate messages on a graph, where each node obtains messages from all of its neighbors and then performs neighborhood aggregation and representation combination iteratively. Finally, node representations can be integrated into graph representations by pooling operations, and thus both the structural information and attributive knowledge of individual nodes can be implicitly merged into the graph-level representation [11].

Although existing GNNs have achieved excellent performance, these methods typically concentrate on a balanced data-split setting. However, realistic data tends to follow the long-tailed class distribution [44, 55]. It is primarily evident from the fact that a minority of dominant classes (i.e. head classes) usually occupy a large amount of data, while the majority of classes (i.e. tail classes) each contains very few data samples. In such data distributions, naive GNNs often degrade and result in sub-optimal classification performance on tail classes due to two major limitations: (i) models learned from the long-tailed distribution can be easily biased towards head classes [71], and (ii) models fail to generalize well for tail classes because of insufficient training data. Undoubtedly, the long-tailed characteristics immensely restrict the practical use of GNNs, hence it is essential to develop tailored GNN methods for realistic long-tailed graph data.

Recently, long-tailed learning has received widespread attention in the context of neural networks [2, 27, 28, 32, 56, 70, 74], there exist three active strands of work: class re-balancing, information augmentation, and decoupling training. Class re-balancing attempts to balance the training data of different classes during model training, either by re-sampling the data [12, 16] or using cost-sensitive learning [3, 43, 60]. Information augmentation, on the other hand, seeks to improve model performance by introducing additional knowledge through a transfer learning approach [40, 43]. Moreover, recent work has shown that decoupling feature learning and classifier learning can lead to improved performance compared to traditional end-to-end training methods [22]. While there are several long-tailed algorithms designed for node-level classification on graphs [33, 41, 52, 67], the problem of long-tailed recognition for graph-level classification is still largely unexplored and challenging, which is the main focus of this paper.

Despite achieving encouraging performance, most of the existing algorithms still suffer from two key limitations. Firstly, most class re-balancing approaches are inclined to improve the performance of tail classes, but this often comes at the expense of head classes [71]. Secondly, previous research has focused on balancing either classifier learning or representation learning individually, neglecting the importance of simultaneously addressing both aspects. Therefore, it is desirable to jointly optimize representation learning and classifier learning and avoid damaging the head class performance to develop effective and robust long-tailed classification algorithms.

We identify insufficient training samples and a lack of within-class variability in tail classes as the most significant reason that hinders long-tailed graph classification. Therefore, we propose a novel framework named Retrieval Augmented Hybrid Network (RAHNet), which adopts a decoupled training procedure to boost the learning of both the feature extractor and the classifier. In contrast to existing class re-balancing approaches, our framework aims to explicitly introduce additional knowledge to enrich tail classes and learn the balanced weights of the classifier. For the learning of feature extractors, we jointly train a standard base encoder, an additional retrieval branch, and a contrastive learning based branch. Specifically, the retrieval branch makes use of training graphs as retrieval keys and returns the most relevant corpus graphs by sub-graph matching techniques, which directly enriches the intra-class diversity for the tail classes. We also explore supervised contrastive learning strategies and tailor them to balance the lower bound of

loss value among head and tail classes by introducing a set of category centers. By this means, our balanced supervised contrastive learning (BSCL) aggregates the samples of each category to their center and facilitates better representation learning from long-tailed data. Moreover, we investigate different weight regularizers and use them to fine-tune the classifier, which avoids biased weight norms toward head classes. The contributions of this work can be summarized as follows:

- We elaborate on a hybrid network RAHNet for long-tailed graph classification, which comprises a retrieval augmentation branch and a balanced contrastive learning module to explicitly introduce intra-class variability and enhance feature learning.
- To balance classifier weights without harming representation learning, we propose to decouple feature extractor learning and classifier training, and fine-tune the classifier with multiple weight regularization techniques.
- RAHNet is compared with multiple imbalance handling methods on a variety of benchmarks with different imbalance factors. Experimental results show the superiority of our approach in long-tailed graph classification.

2 RELATED WORK

2.1 Long-tailed learning

Long-tailed learning approaches can be roughly divided into three main categories: class re-balancing, transfer learning, and decoupling training. Generally, the re-balancing strategy aims to equalize the contribution of different classes and emphasize the tail classes from various angles. Specifically, re-sampling techniques [24, 68] balance the classes from the perspective of training data, which achieve a more balanced data distribution across classes. Cost-sensitive learning approaches [15, 53] introduce balance from a different perspective, which adjust loss values for each class during training. However, the use of re-balancing techniques may lead to under-fitting for head classes and reduced accuracy [58]. For transfer learning approaches, head-to-tail knowledge transfer seeks to transfer knowledge from data-abundant classes to data-poor classes [31, 42], while [57, 63] propose to learn a unified student model using adaptive knowledge distillation from the multiple teacher experts. Besides, decoupling training methods divide the training procedure into multiple stages. For instance, [22] proposes to decouple the learning procedure into the encoder learning and the classifier training, and [73] suggests learning from both the instance-balanced sampling branch and reversed sampling branch. While some recent work [59] has applied imbalance handling techniques to imbalanced graph classification, the effectiveness of these methods for multi-class long-tailed graph classification remains uncertain. Our proposed framework simultaneously enhances feature learning with transfer learning and strengthens the classifier training with weight regularization in a decoupled manner to tackle the long-tailed graph classification problem.

2.2 Graph retrieval

Given a query graph, graph retrieval systems aim to determine the most similar graph among a set of corpus graphs. Recently, leveraging neural architectures to tackle the task of graph matching has received increasing attention. For example, GMN [30] computes a




Figure 1: Overview of the proposed RAHNet architecture. Our RAHNet is a hybrid network with three primary branches and corresponding loss values, *i.e.*, the base supervised learning branch (\mathcal{L}_{base}), the retrieval augmentation branch (\mathcal{L}_{ret}), and the balanced contrastive learning branch (\mathcal{L}_{con}), the model is trained by the losses summed up from each branch with different weight coefficients.

similarity score for a pair of graphs through a cross-graph attention mechanism to associate nodes across graphs. More recently, ISONet [50] is also built upon MPNNs to obtain edge embeddings, which are subsequently used to learn an edge alignment network to approximate the underlying correspondence between graphs.

2.3 Contrastive learning

Contrastive learning has attracted increasing attention in the field of unsupervised representation learning, which aims to aggregate semantically similar samples and obtain discriminative representations by comparing positive and negative pairs. Among them, SimCLR [6] and MoCo [14] are two classical approaches to self-supervised contrastive learning. Supervised contrastive learning (SCL) [23] is an extension to contrastive learning by leveraging the label information to compose positive and negative samples, which leads to remarkable performance for classification.

3 METHODOLOGY

This section begins by introducing the problem setting and providing necessary background information. Subsequently, we introduce the proposed RAHNet framework (as shown in Figure 1) in detail, including the retrieval module, BSC module and classifier regularization. Furthermore, the training pipeline and time complexity analysis of RAHNet are thoroughly explained.

3.1 Preliminaries

Problem Formulation. Let $\mathcal{G} = \{(V_i, E_i, y_i)\}_{i=1}^N$ denote the long-tailed graph dataset, where each graph consists of a node set V_i , an edge set E_i , along with the corresponding label $y_i \in \{1, 2, \dots, C\}$. The total number of training samples over C classes is $N = \sum_{c=1}^C n_c$, where n_c denotes the data number of class c . Without loss of generality, we assume $n_1 \geq n_2 \geq \dots \geq n_C$, and $n_1 \gg n_C$ after sorting the classes by cardinality in decreasing order, and then the imbalance factor (IF) can be defined as n_1/n_C . In the task of long-tailed graph

classification, we aim to learn an unbiased model on a long-tailed training dataset that generalizes well on a balanced test dataset.

Base Encoder. In our hybrid framework, we adopt a unified GNN encoder shared by multiple branches, which leverages both node attributes and graph topology to obtain embeddings. Let $f(\cdot)$ be a L -layer GNN encoder, and $\mathcal{A}_\theta^{(l)}$ and $\mathcal{C}_\theta^{(l)}$ denote the aggregation and combination functions at the l -th layer, then the propagation of the l -th layer is formulated as:

$$\mathbf{h}_v^{(l)} = \mathcal{C}_\theta^{(l)} \left(\mathbf{h}_v^{(l-1)}, \mathcal{A}_\theta^{(l)} \left(\left\{ \mathbf{h}_u^{(l-1)} \right\}_{u \in \mathcal{N}(v)} \right) \right), \quad (1)$$

where $\mathbf{h}_v^{(l)}$ is the embedding of the node v at the l -th layer, $\mathcal{N}(v)$ refers to the neighbors of v . After l iterations, the final graph embedding $f(G)$ is integrated from embedding vectors at all layers through a READOUT function:

$$\mathbf{h}_i = f(G_i) = \text{READOUT}(\{\mathbf{h}_v^{(l)} : v \in \mathcal{V}_i, l \in L\}). \quad (2)$$

3.2 Retrieval Augmented Branch

In the task of graph retrieval, we are given a set of query graphs G_q and a set of corpus graphs G_c , the goal is to design a neural distance function $d(G_q|G_c)$ that accurately predicts the similarity between them. We first pad G_q with $|V_c| - |V_q|$ dummy nodes, and the augmented node adjacency matrices of size W can be represented as \mathbf{A}_q and \mathbf{A}_c . The objective function for the node alignment approach is defined as:

$$\arg \min_S \sum_{i,j \in [W] \times [W]} [(\mathbf{A}_q - \mathbf{S} \mathbf{A}_c \mathbf{S}^T)_+]_{i,j}, \quad (3)$$

where $(\cdot)_+ = \max\{0, \cdot\}$, $[\cdot]_{i,j}$ represents the element located at the i -th row and j -th column of the matrix, and \mathbf{S} is the node permutation matrix. The node alignment methods seek to approximate Eq. (3) with node embeddings, which is still a hard quadratic assignment problem. Thus, one feasible way of node alignment [30] is directly aggregating node embeddings to graph embedding and performing

graph matching, which may fail to capture edge correspondences. To tackle this problem, we perform subgraph matching with edge alignment by rewriting the objective of Eq. 3 as:

$$\arg \min_S \sum_{k \in [W^2]} [(\text{vec}(\mathbf{A}_q) - (\mathbf{S} \otimes \mathbf{S}) \text{vec}(\mathbf{A}_c))_+]_k, \quad (4)$$

where \otimes is the Kronecker product and $\text{vec}(\mathbf{A}_\cdot)$ denotes the vector obtained by column-wise concatenation. In practice, we approximate Eq. 4 with edge embeddings. Let $\mathbf{r}_e \in \mathbb{R}^D$ denote the embedding of edge e , and the edge embedding matrices can be represented as $\mathbf{R}_q, \mathbf{R}_c \in \mathbb{R}^{|E| \times D}$ collectively. Therefore, we can write the distance measure as:

$$d(G_q|G_c) = \min_P \sum_e [(\mathbf{R}_q - \mathbf{P}\mathbf{R}_c)_+]_e, \quad (5)$$

where \mathbf{P} is an edge permutation without a quadratic term.

We develop our edge alignment-based retrieval branch by approximate $d(G_q|G_c)$ in Eq. 5 with two neural networks as $d_{\phi, \varphi}(G_q|G_c)$: (i) The first network is a message passing framework with parameter ϕ to model the edge embedding matrices \mathbf{R}_q and \mathbf{R}_c , (ii) The second network with parameter φ is the Gumbel-Sinkhorn operator, which aims to provide a differentiable solution \mathbf{P}^* to address the optimization problem. Therefore, the search for discrete edge permutation matrix P is replaced with the relaxation:

$$d_{\phi, \varphi}(G_q|G_c) = \sum_{i,j} [(\mathbf{R}_q - F_\varphi(\mathbf{R}_q, \mathbf{R}_c) \mathbf{R}_c)_+]_{i,j}. \quad (6)$$

Finally, we follow [50] to train the above two networks with parameters of ϕ and φ through a hinge loss:

$$\min_{\phi, \varphi} \sum_{G_q} \sum_{G_{c\oplus}, G_{c\ominus}} [\gamma + d_{\phi, \varphi}(G_q|G_{c\oplus}) - d_{\phi, \varphi}(G_q|G_{c\ominus})]_+, \quad (7)$$

where $\gamma \in \mathbb{R}^+$ is a margin parameter, and $G_{c\oplus}$ and $G_{c\ominus}$ represent the relevant and irrelevant graphs for G_q , respectively. We use a breadth-first search-based technique to sample $G_{c\oplus}$ and $G_{c\ominus}$ from \mathcal{G} , employing the VF2 exact match algorithm [34] to determine the ground-truth training targets for any given G_q .

With the pre-trained retrieval network $d_{\phi, \varphi}(\cdot|\cdot)$, we set both G_q and G_c to be the training dataset, and retrieve each query graph for top- q relevant graphs. In the implementation, we augment the original data batch B to \hat{B} with retrieved graphs, \hat{B} is then fed into the encoder to obtain both the base representation \mathbf{h}_{base} and representations of retrieved graphs $\hat{\mathbf{h}} \in \mathbb{R}^{q \times D}$, where D is the dimension of embedding vectors. The retrieval feature \mathbf{h}_{ret} are then adaptively combined by the representations of q retrieved graphs, each with a different weight coefficient:

$$\mathbf{h}_{ret} = \mathbf{a}^\top \hat{\mathbf{h}} = \sum_{i=1}^q \mathbf{a}_i \hat{\mathbf{h}}_i, \quad (8)$$

where $\mathbf{a} \in \mathbb{R}^q$ represents the attention coefficients over each retrieved graph derived from the base feature. Specifically, we use a fully-connected layer \mathbf{W}_a with a softmax activation function to obtain the attention coefficients: $\mathbf{a} = \text{Softmax}(\mathbf{W}_a \mathbf{h}_{base})$. The obtained embeddings \mathbf{h}_{base} and \mathbf{h}_{ret} are forwarded by a classifier layer to obtain the prediction probabilities. Finally, the prediction probabilities for both the supervised learning branch and retrieval augmented branch are used to calculate the standard cross-entropy loss, which can be denoted as \mathcal{L}_{base} and \mathcal{L}_{ret} , respectively.

3.3 Balanced Supervised Contrastive Learning

In this subsection, we discuss how to leverage contrastive learning to produce discriminative and robust representations under the long-tailed setting. Following [65], we first obtain different positive views of graphs by involving four fundamental data augmentation strategies that preserve intrinsic structural and attribute information: (1) *Edge permutation* (2) *Attribute masking* (3) *Node dropping* (4) *Subgraph*. During training, we are given a two-viewed graph batch $B = (B_{v1}, B_{v2})$ and label y , the graphs in B_{v2} are transformed by aforementioned data augmentation strategies. Each view of graphs is fed into the shared GNN encoder to obtain the embeddings \mathbf{H}_{v1} and \mathbf{H}_{v2} , respectively. Afterward, the embeddings are fed into the projection network $g(\cdot)$ which projects H to another latent space where the contrastive loss is calculated. After obtaining projected representation \mathbf{Z}_{v1} and \mathbf{Z}_{v2} , we define the positives $P(i)$ for any anchor graph G_i in B_{v1} as:

$$\begin{aligned} A(i) &= \{\mathbf{z}_k \in \mathbf{Z}_{v1} \cup \mathbf{Z}_{v2}\} \setminus \{\mathbf{z}_k \in \mathbf{Z}_{v1} : k = i\}, \\ P(i) &= \{\mathbf{z}_k \in A(i) : y_k = y_i\}. \end{aligned} \quad (9)$$

However, the scarcity of tail classes will result in the lack of positive pairs, which consequently leads to performance deterioration of instance-level contrastive loss on the long-tailed dataset. Therefore, we incorporate a set of category centers \mathbf{O} to balance the loss value among head and tail classes. More specifically, the category centers are a set of learnable parameters, and the dimension of \mathbf{O} is $\mathbb{R}^{C \times D}$, where C is the number of classes and D is the embedding dimension. Formally, the loss of BSCL can be written as:

$$\mathcal{L}_i^{con} = \sum_{\mathbf{z}_+ \in P(i) \cup \mathbf{O}_y} -w(\mathbf{z}_+) \log \frac{\exp(\mathbf{z}_+ \cdot g(\mathbf{h}_i)/\tau)}{\sum_{\mathbf{z}_k \in A(i) \cup \mathbf{O}} \exp(\mathbf{z}_k \cdot g(\mathbf{h}_i)/\tau)}, \quad (10)$$

where $\tau \in \mathbb{R}^+$ is a scalar temperature parameter, and

$$w(\mathbf{z}_+) = \begin{cases} 1.0, & \mathbf{z}_+ \in \mathbf{O}_y \\ \alpha, & \mathbf{z}_+ \in P(i) \end{cases}. \quad (11)$$

Moreover, we instantiate $g(\cdot)$ as either a multi-layer perceptron $M(\cdot)$ or an identity mapping layer $I(\cdot)$, i.e., $I(x) = x$. Specifically, different projection head is chosen as follows:

$$\mathbf{z} \cdot g(\mathbf{h}_i) = \begin{cases} \mathbf{z} \cdot I(\mathbf{h}_i), & \mathbf{z} \in \mathbf{O} \\ \mathbf{z} \cdot M(\mathbf{h}_i), & \mathbf{z} \in A(i) \end{cases}. \quad (12)$$

To demonstrate the superiority of BSCL in long-tailed learning, we further conduct an analysis of the loss value between supervised contrastive learning and BSCL. Suppose K_{y_i} is the expected number of positive pairs with respect to given graph G_i and its label y_i , which can be calculated as:

$$K_{y_i} = (2 * \text{batchsize} - 1) * \pi_{y_i}, \quad (13)$$

where π_{y_i} is the class frequency over the whole dataset, i.e., n_{y_i}/N . When supervised contrastive loss achieves minimum, the optimal value for the probability that two graph samples are a true positive pair is $1/K_y$, where y is the corresponding label. In that case, the head classes will have a higher lower bound of loss value and contribute significantly more importance than tail classes during training. However, in regard to BSCL loss, the optimal value for the probability that two graph samples are a true positive pair is $\alpha/(1 + \alpha \cdot K_y)$, and optimal value for the probability that a graph

Algorithm 1: Training algorithm of RAHNet**Input:** Training set \mathcal{G} **Parameter:** GNN encoder $f(\cdot, \theta_f)$, pre-trained retrieval network $d_{\phi, \varphi}(\cdot|\cdot)$, contrastive learning module parameter $\psi(\cdot, \theta_b)$, classifier (\cdot, θ_c) , epoch number T , epoch number for fine-tuning T_f .**Output:** The proposed RAHNet

- 1: **for** $t = 1$ to T **do**
- 2: Sample batch B from the instance-balanced sampler.
- 3: Augment B to a two-viewed batch $B = (B_{v1}, B_{v2})$.
- 4: Retrieve for relevant graphs G_r using $d_{\phi, \varphi}(\cdot|\cdot)$.
- 5: Expand B_{v1} to \hat{B}_{v1} with G_r .
- 6: Forward propagation \hat{B}_{v1} via $f(\cdot, \theta_f)$ and (\cdot, θ_c) .
- 7: Compute \mathcal{L}_{base} and \mathcal{L}_{ret} .
- 8: Forward propagation $B = (B_{v1}, B_{v2})$ via $\psi(\cdot, \theta_b)$.
- 9: Compute contrastive loss \mathcal{L}_{con} using Eq. 10.
- 10: Sum up the loss values using Eq. 17.
- 11: Update θ_f , θ_b , and θ_c by back propagation.
- 12: **end for**
- 13: Fix the weights of GNN encoder θ_f .
- 14: **for** $t = 1$ to T_f **do**
- 15: Sample batch B using the class-balanced sampler.
- 16: Forward propagation B via $f(\cdot, \theta_f)$ and (\cdot, θ_c) .
- 17: Compute the objective function of cross-entropy and weight decay.
- 18: Update θ_c by back propagation.
- 19: Balance θ_c using equation Eq. 15.
- 20: **end for**

is closest to its category center \mathbf{o}_y is $1/(1 + \alpha \cdot K_y)$. Therefore, the head-to-tail optimal value gap is reduced from $1/K_{y_{head}} \rightarrow 1/K_{y_{tail}}$ to $1/(1 + \alpha \cdot K_{y_{head}}) \rightarrow 1/(1 + \alpha \cdot K_{y_{tail}})$. As the value of α decreases, the optimal value from head to tail becomes increasingly balanced, which enhances the representation quality of BSCL in the long-tailed settings.

3.4 Re-balancing the Classifier

Long-tailed class distribution often results in larger classifier weight norms for head classes [64], which makes the classifier easily biased to dominant head classes. To address this problem, we regularize the weights of the classifier with the trained feature extractor frozen. Specifically, we investigate two different weight regularization approaches to balance weights with respect to norms, *i.e.*, Max-norm and weight decay.

Max-norm. Max-norm constrains the weights to have a norm less than or equal to a specific value. More formally, Max-norm caps weight norms within an L2-norm ball:

$$\Theta^* = \arg \min_{\Theta} F(\Theta; \mathcal{G}), \quad \text{s.t.} \quad \|\theta_k\|_2^2 \leq \delta^2, \quad \forall k, \quad (14)$$

where δ is the radius. Moreover, we adopt projected gradient descent (PGD) to solve Eq. 14, which efficiently projects out-of-ranged weights onto the L2-normball. Concretely, it applies a renormalization step at each iteration. After each batch update, PGD computes

Table 1: Statistics of the datasets used in the experiments.

Dataset	# Graphs	# Classes	# Avg. Nodes	# Avg. Edges
Synthie	400	4	95.00	172.93
ENZYMES	600	6	32.63	62.14
MNIST	60,000	10	75	1393.27
Letter-High	2250	15	4.67	4.50
Letter-low	2250	15	4.67	4.50
COIL-DEL	3900	100	21.54	54.24

an updated θ_k and projects it onto the L2-norm ball with radius δ :

$$\theta_k \leftarrow \min(1, \delta / \|\theta_k\|_2) * \theta_k. \quad (15)$$

Unlike L2-normalization, which forces weights to be unit-length in norm, Max-norm loosens this constraint, allowing the weights to fluctuate within the norm-ball during training.

Weight Decay. Weight decay is a widely used type of regularization, which is utilized to constrain the growth of the network weights. Weight decay typically penalizes the network weights according to their L2-norm:

$$\Theta^* = \arg \min_{\Theta} F(\Theta; \mathcal{G}) + \lambda \sum_k \|\theta_k\|_2^2, \quad (16)$$

where the hyperparameter λ is used to control the impact of weight decay. It enhances the model’s generalization performance by punishing large weights and encouraging learning small balanced weights which consequently decreases the complexity of the network to prevent overfitting.

Moreover, jointly applying Max-norm and weight decay leads to better performance because of their complementing benefits. Generally, Max-norm caps large weights within the specific radius and prevents them from dominating the training, while weight decay on the small weights still improves the overall generalization and avoids overfitting.

3.5 Training Pipeline

We train our RAHNet framework in a two-stage manner, which decouples representation learning from classifier learning. We also utilized different sampling strategies for different training stages: (i) instance-balanced sampler, in which each data sample has the same probability of being sampled; (ii) class-balanced sampler, in which each class is sampled uniformly and each instance is sampled uniformly within it.

In the first stage, we use an instance-balanced sampler to learn a better feature extractor that preserves better generalizability. Given a data batch B , we first transform B to a two-viewed graph batch $B = (B_{v1}, B_{v2})$ through data augmentation. The original view of batch B_{v1} is fed into the base branch and retrieval augmented branch to deduce the standard cross-entropy loss \mathcal{L}_{base} and the retrieval loss \mathcal{L}_{ret} , respectively. While $B = (B_{v1}, B_{v2})$ is fed into the BSCL module to deduce the contrastive loss \mathcal{L}_{con} . Finally, RAHNet sums up those losses by:

$$\mathcal{L}_{total} = \mathcal{L}_{base} + \eta_{ret} \cdot \mathcal{L}_{ret} + \eta_{con} \cdot \mathcal{L}_{con}, \quad (17)$$

where η_{ret} and η_{con} are hyperparameters that control the contribution of different branches.

Table 2: Long-tailed graph classification accuracy on six benchmark datasets with various IFs (best performance in bold).

Model	Synthie		ENZYMES		MNIST		Letter-high		Letter-low		COIL-DEL	
	IF=15	IF=30	IF=15	IF=30	IF=50	IF=100	IF=25	IF=50	IF=25	IF=50	IF=10	IF=20
GraphSAGE	34.74	30.25	30.66	25.16	68.67	63.46	51.06	42.16	86.00	84.32	38.80	31.32
Up-sampling	35.25	33.50	32.33	28.50	64.69	59.78	53.62	44.20	88.48	86.72	39.20	26.96
CB loss	34.75	30.75	32.19	26.83	68.85	63.40	53.76	45.06	87.46	85.44	41.72	32.34
LACE loss	33.25	30.85	31.16	25.50	69.72	64.59	47.46	38.94	87.89	84.69	41.96	32.18
Augmentation	39.37	35.37	32.08	26.75	72.18	68.17	49.28	42.36	88.32	86.40	38.18	30.80
G ² GNN _n	38.08	27.94	35.00	29.17	70.91	66.73	58.91	51.12	89.49	87.98	38.32	27.98
G ² GNN _e	40.19	37.53	35.83	29.50	73.69	70.31	58.85	49.96	89.84	87.80	39.18	31.06
GraphCL	40.25	36.25	36.66	29.83	69.37	65.12	57.34	48.93	89.28	87.89	42.02	33.19
SupCon	40.34	37.25	37.08	30.67	69.76	64.88	57.29	48.93	89.12	87.36	42.93	34.20
RAHNet	42.35	36.76	38.50	32.17	75.12	71.98	59.20	50.37	89.65	88.69	43.04	36.80
RAHNet _{dec}	40.75	39.00	39.02	34.16	75.79	72.60	59.79	52.90	90.19	89.28	45.32	38.48

In the second stage, a class-balanced sampler together with Max-norm and weight decay regularization are used to learn a balanced classifier. To avoid representation damage because of the balanced sampler, we fix the feature encoder and only fine-tune the classifier layers with cross-entropy loss. The overall training pipeline is illustrated in Algorithm 1.

3.6 Time Complexity Analysis

Denote the batch size B , and the average number of nodes in the input graphs is $|V|$. The time complexity of acquiring embeddings from the GNN encoder for a batch is $O(BLD|V|)$, where L is the number of layers of the encoder and D denotes the embedding dimension. With the pre-trained retrieval model, the complexity of obtaining the retrieval feature is $O(qB)$, where q is the number of retrievals. For BSCL, we calculate the loss for a batch in $O(DB^2)$ time, and for the classifier re-balancing, the time complexity of max-norm and weight decay is $O(BC)$, where C is the number of classes. Overall, the total time complexity for RAHNet is $O(B(LD|V| + q + DB + C))$.

4 EXPERIMENT

4.1 Experimental Setups

Benchmark Datasets. We evaluate our proposed RAHNet on six publicly accessible datasets in various fields, including synthetic (Synthie [39]), bioinformatics (ENZYMES [51]), and computer vision (MNIST [9], Letter-high [49], Letter-low [49], and COIL-DEL [49]). We follow Zipf’s law to process the original datasets into long-tailed datasets. Moreover, we split the dataset into train/val/test sets in a ratio of 60%/20%/20%, respectively.

Competing Models. We carry out comprehensive comparisons with methods from four categories, which include: (a) data re-balancing method, *i.e.*, up-sampling [5], (b) loss re-weighting approaches, *i.e.*, class-balanced (CB) loss [7], and logit adjustment cross-entropy (LACE) loss [38], (c) information augmentation, *i.e.*,

data augmentation [66], and G²GNN [59], and (d) contrastive learning, *i.e.*, GraphCL [65], and SupCon [23].

Implementation Details. For the proposed RAHNet, we use GraphSAGE [13] as the base GNN encoder, and empirically set the embedding dimension to 64, the number of epochs to 1000, and batch size to 32. The softmax temperature γ is set to 0.2, the contrast weight α is set to 0.05, and weight decay of the second stage is 0.1. Moreover, we tune η_{ret} , η_{con} and δ for different datasets, and average accuracy over 10 runs is used as the evaluation metric.

4.2 Results and Analysis

The experimental results on the six benchmarks with different imbalance factors (IFs) are shown in Table 2. Our experiment involves comparing the baseline methods with two variants of our RAHNet model: one without weight regularization (RAHNet) and the other with decoupled classifier training (RAHNet_{dec}). We can make the following observations:

- The classification performance of both baseline methods and our proposed RAHNet suffers a sharp decrease in all six datasets when the long-tailedness between head and tail classes increases, indicating that GNNs often degrade and result in sub-optimal classification performance under the long-tailed setting.
- Among the four categories of baselines, information augmentation approaches surpass data re-sampling and loss re-weighting baselines on most datasets. This may be due to the fact that it is crucial for information augmentation methods to explicitly introduce new knowledge to enrich the tail classes, leading to better representation ability for the tail. Moreover, contrastive learning-based approaches exhibit relatively stable performance across all six datasets.
- Our RAHNet and its variant (RAHNet_{dec}) consistently outperform other baselines on all six datasets under different imbalance settings, especially under harsh imbalance settings such as COIL-DEL dataset with IF=20, where over 60 classes only contain less

Table 3: Ablation study for the primary components of RAHNet (RA: Retrieval Augmentation, SCL: Supervised Contrastive Learning, BSCL: Balanced Supervised Contrastive Learning, and WR: Weight Regularization).

RA	SCL	BSCL	WR	Letter-high	ENZYMES
✓				48.83	36.67
	✓			48.93	37.08
		✓		49.51	37.78
✓	✓			49.58	37.83
	✓		✓	50.76	37.61
✓		✓		50.37	38.50
✓		✓	✓	52.90	39.02

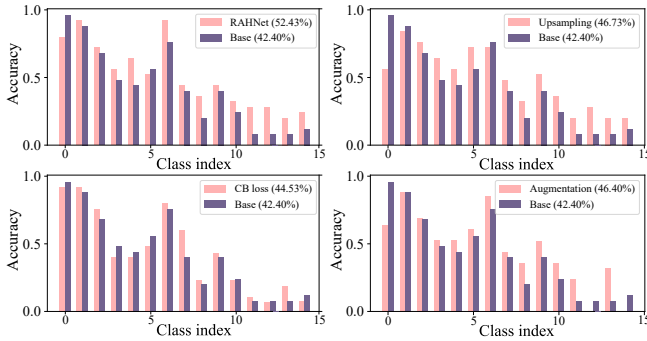


Figure 2: Per-class classification accuracy on the Letter-high dataset with an IF of 50.

than 3 training samples. Moreover, even without re-balancing the classifier, our RAHNet still outperforms the others in most scenarios, which implies that the retrieval branch and BSCL module jointly learn discriminative representations and play a significant role in alleviating the long-tailedness problem.

4.3 Ablation Study

We perform several ablation studies to characterize the proposed RAHNet. First, we compare the performance of the primary components of the framework in Table 3, where retrieval augmentation (RA), balanced supervised contrastive learning (BSCL), and weight regularization (WR) are three main components in RAHNet framework. Moreover, we investigate the effectiveness of BSCL by substituting the supervised contrastive learning (SCL) module for BSCL. All experiments are performed on Letter-high (IF=50) and ENZYMES (IF=15). First of all, we can clearly observe that using both retrieval augmentation and balanced contrastive learning solely can enhance the overall performance, and jointly training the model with both of the components can significantly boost the performance. Secondly, combining BSCL with other components (RA + BSCL) outperforms RA + SCL. The results indicate that SCL is not applicable to long-tailed datasets due to the lack of

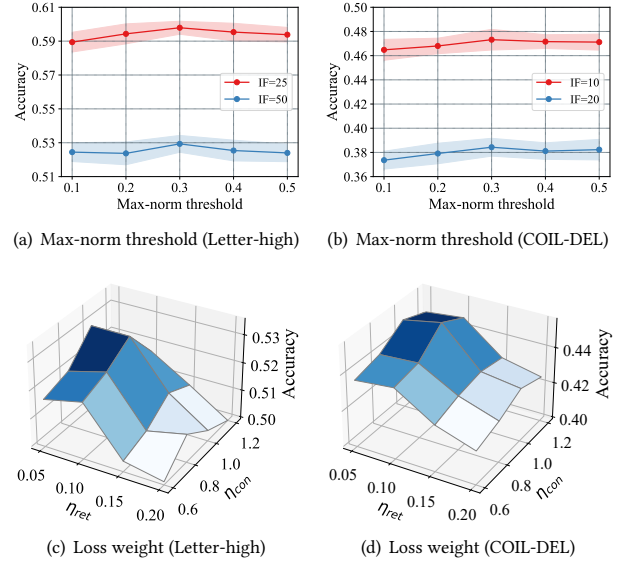


Figure 3: Hyper-parameter sensitivity study of RAHNet.

positive pairs for tail classes, and our proposed BSCL effectively balanced the training in long-tailed distributions. Finally, our complete framework achieves better performance, demonstrating that weight regularization helps to learn an unbiased classifier under the long-tailed setting.

We further study the impact of decoupling training by comparing the per-class accuracy, the results are shown in Figure 2. It can be observed that all the approaches, perform well on data-rich head classes but significantly worse on the tail classes. Our RAHNet framework enhances both head and tail classes classification accuracy with little loss of head-class performance. Without explicitly re-balancing the data samples or loss contribution during training, decoupling training with weight regularization still balances the performance between head and tail classes well.

4.4 Hyper-parameter Study

We further study the sensitivity of hyper-parameters in RAHNet. First, we examine the effect of Max-norm threshold δ by varying δ in $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ with all other hyper-parameters fixed. The experimental results on the Letter-high and COIL-DEL datasets are shown in Figure 3(a) and 3(b), respectively. We can clearly observe that the classification accuracy rises as the Max-norm threshold increases from 0.1 to 0.3, and the performance saturates as the threshold reaches 0.3 on both datasets with various IFs, which validates that larger δ offers more free space within the norm ball to let weights grow and also caps the weights to prevent the head class from dominating the training.

Moreover, we study the effect of loss weight hyper-parameters of the retrieval loss η_{ret} and the contrastive loss η_{con} , which are critical to the learning of the feature extractor. We conduct experiments on the Letter-high (IF=50) and COIL-DEL (IF=10) datasets by varying $\eta_{ret} \in \{0.05, 0.1, 0.15, 0.2\}$ and $\eta_{con} \in \{0.6, 0.8, 1.0, 1.2\}$. In Figure 3(c) and 3(d), it can be found that the accuracy reaches the

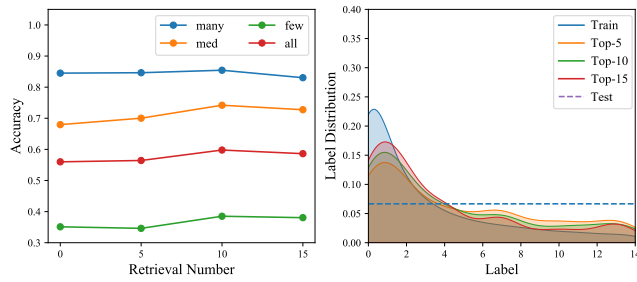


Figure 4: Analysis of retrieval module on Letter-high (IF=25). Left: Comparison of using different retrieval number q ; Right: Label distributions of retrieval augmented training data.

peak when η_{ret} is set to 0.1 and η_{con} is set to 1.0. As η_{ret} rises from 0.05 to 0.1, the performance improves accordingly. These results indicate that increasing the retrieval branch contribution of loss in an appropriate range can be beneficial to long-tailed learning because tail classes are augmented with diverse retrieval features. The BSCL module serves the purpose of representation enhancement and highlights the tail classes during training. Therefore, the effectiveness of BSCL is maximized when the contrastive loss shares an equal contribution of the cross-entropy loss.

4.5 Retrieval Branch Analysis

As shown in Figure 4, we analyze the effectiveness of the retrieval branch by varying the retrieval number q from 0 to 15. We first divide the different classes into many-shot, med-shot, and few-shot categories according to their number of samples, and record the corresponding accuracy when q changes. It can be concluded that the performance on med-shot and few-shot categories consistently increases when q increases from 0 to 10, indicating the retrieval module is able to emphasize tail classes during training. Long-tailed learning can also be formulated as a label shift problem where the training and testing label distributions are different. Therefore, we visualize the label distribution of original training data, testing data, and retrieval augmented data. As can be seen from the figure, retrieval augmentation alleviates the class long-tailedness by diminishing the dominance of the head class and making the label distribution closely aligned with the distribution of the test data. The top-10 distribution among various levels of augmentation exhibited the lowest oscillation amplitude in the tail and demonstrated the best overall performance.

4.6 Case Study

We investigate the power of our retrieval augmentation branch to show the superiority of explicitly introducing extra knowledge via a transfer learning approach. We visualize the feature distribution of Letter-low with t-SNE in Figure 5, and we focus on a specified class ID-15 in this case study. In (a), we first visualize the feature distribution when ID-15 is a head class, and has the exact same number of samples as other classes. We can observe that ID-15 generalizes well and has a clear decision boundary with neighboring classes due to abundant training samples. In (b), we limit the training samples of ID-15, so it becomes a tail class. The feature

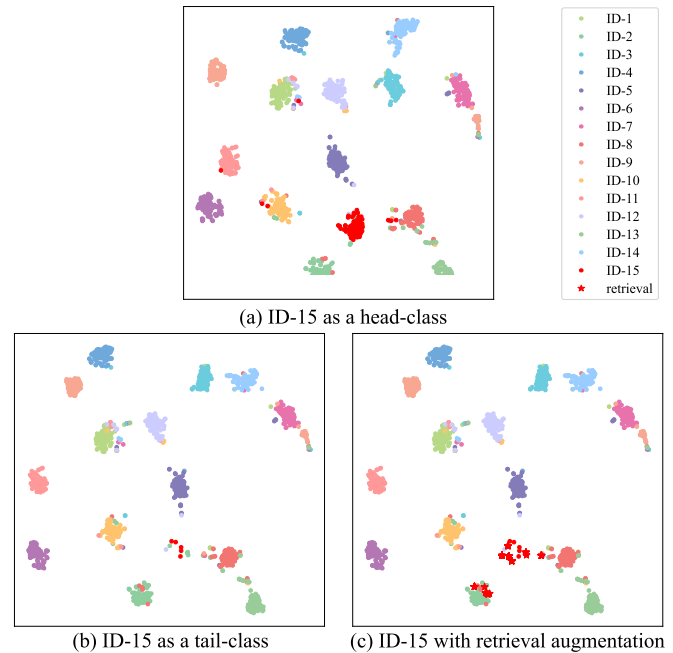


Figure 5: Visualize the effectiveness of retrieval augmentation on the COIL-DEL dataset with t-SNE.

distribution of ID-15 collapses into a relatively small scope due to the lack of intra-class variability. In (c), we augment the tail class of ID-15 with relevant graphs collected from our retrieval branch. The retrieval features are scattered around the original features of ID-15 in the high-dimensional embedding space, which improve the variety of tail class samples and promote the generalization ability in the tail class.

5 CONCLUSION

This paper presents a retrieval augmented hybrid network (RAHNet), which combines transfer learning and decoupling training to boost the learning of both the feature extractor and the classifier simultaneously. RAHNet leverages a retrieval augmented branch and a balanced supervised contrastive learning module to jointly learn a balanced feature extractor under the long-tailed setting. Moreover, we re-balancing the classifier weight in norms by applying two weight regularization techniques, which prevent the head classes from dominating the training and improves the tail classes' performance. Extensive empirical studies on six benchmarks show that our approach outperforms competitive baselines on long-tailed data. A potential direction for future work is to enhance the retrieval module by incorporating multi-modal knowledge.

ACKNOWLEDGEMENTS

This paper is partially supported by the China Postdoctoral Science Foundation with Grant No. 2023M730057 as well as the National Natural Science Foundation of China with Grant No. 62276002.

REFERENCES

- [1] Pietro Bongini, Monica Bianchini, and Franco Scarselli. 2021. Molecular generative graph neural networks for drug discovery. *Neurocomputing* 450 (2021), 242–252.
- [2] Jiarui Cai, Yizhou Wang, Hung-Min Hsu, Jenq-Neng Hwang, Kelsey Magrane, and Craig S Rose. 2022. LUNA: Localizing Unfamiliarity Near Acquaintance for Open-Set Long-Tailed Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 131–139.
- [3] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems* 32 (2019).
- [4] Xianshuai Cao, Yuliang Shi, Jihu Wang, Han Yu, Xinjun Wang, and Zhongmin Yan. 2022. Cross-modal knowledge graph contrastive learning for machine learning method recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*. 3694–3702.
- [5] Nitesh V Chawla. 2003. C4. 5 and imbalanced data sets: investigating the effect of sampling method, probabilistic estimate, and decision tree structure. In *Proceedings of the ICML*, Vol. 3. CIBC Toronto, ON, Canada, 66.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [7] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*. 9268–9277.
- [8] Yanni Dong, Quanwei Liu, Bo Du, and Liangpei Zhang. 2022. Weighted feature fusion of convolutional neural network and graph attention network for hyperspectral image classification. *IEEE Transactions on Image Processing* 31 (2022), 1559–1572.
- [9] Vijay Prakash Dwivedi, Chaitanya K Joshi, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. 2020. Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982* (2020).
- [10] Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. 2021. Spatial-temporal graph ode networks for traffic flow forecasting. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*. 364–373.
- [11] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*. PMLR, 1263–1272.
- [12] Hao Guo and Song Wang. 2021. Long-tailed multi-label visual recognition by collaborative training on uniform and re-balanced samplings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15089–15098.
- [13] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9729–9738.
- [15] Yin-Yin He, Peizhen Zhang, Xiu-Shen Wei, Xiangyu Zhang, and Jian Sun. 2022. Relieving Long-tailed Instance Segmentation via Pairwise Class Balance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7000–7009.
- [16] Ziyu Jiang, Tianlong Chen, Ting Chen, and Zhangyang Wang. 2021. Improving contrastive learning on imbalanced data via open-world sampling. *Advances in Neural Information Processing Systems* 34 (2021), 5997–6009.
- [17] Wei Ju, Zheng Fang, Yiyang Gu, Zequn Liu, Qingqing Long, Ziyue Qiao, Yifang Qin, Jianhao Shen, Fang Sun, Zhiping Xiao, et al. 2023. A Comprehensive Survey on Deep Graph Representation Learning. *arXiv preprint arXiv:2304.05055* (2023).
- [18] Wei Ju, Zequn Liu, Yifang Qin, Bin Feng, Chen Wang, Zhihui Guo, Xiao Luo, and Ming Zhang. 2023. Few-shot molecular property prediction via Hierarchically Structured Learning on Relation Graphs. *Neural Networks* 163 (2023), 122–131.
- [19] Wei Ju, Xiao Luo, Meng Qu, Yifan Wang, Chong Chen, Minghua Deng, Xian-Sheng Hua, and Ming Zhang. 2023. TGNN: A joint semi-supervised framework for graph-level classification. *arXiv preprint arXiv:2304.11688* (2023).
- [20] Wei Ju, Yifang Qin, Ziyue Qiao, Xiao Luo, Yifan Wang, Yanjie Fu, and Ming Zhang. 2022. Kernel-based substructure exploration for next POI recommendation. In *2022 IEEE International Conference on Data Mining (ICDM)*. IEEE, 221–230.
- [21] Wei Ju, Junwei Yang, Meng Qu, Weiping Song, Jianhao Shen, and Ming Zhang. 2022. Kgnn: Harnessing kernel-based networks for semi-supervised graph classification. In *Proceedings of the fifteenth ACM international conference on web search and data mining*. 421–429.
- [22] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2020. Decoupling representation and classifier for long-tailed recognition. *ICLR* (2020).
- [23] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems* 33 (2020), 18661–18673.
- [24] Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. 2020. Imbalanced continual learning with partitioning reservoir sampling. In *European Conference on Computer Vision*. Springer, 411–428.
- [25] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *ICLR* (2017).
- [26] Ling Lei, Jianfeng Li, Tong Chen, and Shigang Li. 2020. A novel graph-tcn with a graph structured representation for micro-expression recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*. 2237–2245.
- [27] Bolian Li, Zongbo Han, Haining Li, Huazhu Fu, and Changqing Zhang. 2022. Trustworthy long-tailed classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6970–6979.
- [28] Jun Li, Zichang Tan, Jun Wan, Zhen Lei, and Guodong Guo. 2022. Nested collaborative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6949–6958.
- [29] Mengzhang Li and Zhanxing Zhu. 2021. Spatial-temporal fusion graph neural networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 4189–4196.
- [30] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. 2019. Graph matching networks for learning the similarity of graph structured objects. In *International conference on machine learning*. PMLR, 3835–3845.
- [31] Jialun Liu, Yifan Sun, Chuchu Han, Zhaopeng Dou, and Wenhui Li. 2020. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2970–2979.
- [32] Xialei Liu, Yu-Song Hu, Xu-Sheng Cao, Andrew D Bagdanov, Ke Li, and Ming-Ming Cheng. 2022. Long-Tailed Class Incremental Learning. In *European Conference on Computer Vision*. Springer, 495–512.
- [33] Zemin Liu, Trung-Kien Nguyen, and Yuan Fang. 2021. Tail-gnn: Tail-node graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1109–1119.
- [34] Zhaoyu Lou, Jiaxuan You, Chengtao Wen, Arquimedes Canedo, Jure Leskovec, et al. 2020. Neural subgraph matching. *arXiv preprint arXiv:2007.03092* (2020).
- [35] Xiao Luo, Wei Ju, Meng Qu, Chong Chen, Minghua Deng, Xian-Sheng Hua, and Ming Zhang. 2022. Dualgraph: Improving semi-supervised graph classification via dual contrastive learning. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 699–712.
- [36] Xiao Luo, Yusheng Zhao, Yifang Qin, Wei Ju, and Ming Zhang. 2023. Towards Semi-supervised Universal Graph Classification. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [37] Changsheng Ma and Xiangliang Zhang. 2021. GF-VAE: a flow-based variational autoencoder for molecule generation. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1181–1190.
- [38] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. 2020. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314* (2020).
- [39] Christopher Morris, Nils M Kriege, Kristian Kersting, and Petra Mutzel. 2016. Faster kernels for graphs with continuous attributes via hashing. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 1095–1100.
- [40] Sarah Parisot, Pedro M Esperança, Steven McDonagh, Tamas J Madarasz, Yongxin Yang, and Zhenguo Li. 2022. Long-tail Recognition via Compositional Knowledge Transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6939–6948.
- [41] Joonhyung Park, Jaeyun Song, and Eunho Yang. 2021. GraphENS: Neighbor-Aware Ego Network Synthesis for Class-Imbalanced Node Classification. In *International Conference on Learning Representations*.
- [42] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoon Yun, and Jin Young Choi. 2022. The Majority Can Help The Minority: Context-rich Minority Oversampling for Long-tailed Classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6887–6896.
- [43] Seulki Park, Jongin Lim, Younghun Jeon, and Jin Young Choi. 2021. Influence-balanced loss for imbalanced visual classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 735–744.
- [44] Zhiliang Peng, Wei Huang, Zonghao Guo, Xiaosong Zhang, Jianbin Jiao, and Qixiang Ye. 2021. Long-tailed distribution adaptation. In *Proceedings of the 29th ACM International Conference on Multimedia*. 3275–3282.
- [45] Zhihao Peng, Hui Liu, Yuheng Jia, and Junhui Hou. 2021. Attention-driven graph clustering network. In *Proceedings of the 29th ACM international conference on multimedia*. 935–943.
- [46] Yifang Qin, Wei Ju, Hongjun Wu, Xiao Luo, and Ming Zhang. 2023. Learning Graph ODE for Continuous-Time Sequential Recommendation. *arXiv preprint arXiv:2304.07042* (2023).
- [47] Yifang Qin, Yifan Wang, Fang Sun, Wei Ju, Xuyang Hou, Zhe Wang, Jia Cheng, Jun Lei, and Ming Zhang. 2023. DisenPOI: Disentangling Sequential and Geographical Influence for Point-of-Interest Recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 508–516.
- [48] Prakash Chandra Rathi, R Frederick Ludlow, and Marcel L Verdonk. 2019. Practical high-quality electrostatic potential surfaces for drug discovery using a graph-convolutional deep neural network. *Journal of medicinal chemistry* 63, 16 (2019), 8778–8790.

- [49] Kaspar Riesen and Horst Bunke. 2008. IAM graph database repository for graph based pattern recognition and machine learning. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 287–297.
- [50] Indradyumna Roy, Venkata Sai Baba Reddy Velugoti, Soumen Chakrabarti, and Abir De. 2022. Interpretable Neural Subgraph Matching for Graph Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 8115–8123.
- [51] Ida Schomburg, Antje Chang, Christian Ebeling, Marion Gremse, Christian Heldt, Gregor Huhn, and Dietmar Schomburg. 2004. BRENDA, the enzyme database: updates and major new developments. *Nucleic acids research* 32, suppl_1 (2004), D431–D433.
- [52] Jaeyun Song, Joonhyung Park, and Eunho Yang. 2022. TAM: Topology-Aware Margin Loss for Class-Imbalanced Node Classification. In *International Conference on Machine Learning*. PMLR, 20369–20383.
- [53] Jingru Tan, Xin Lu, Gang Zhang, Changqing Yin, and Quanquan Li. 2021. Equalization loss v2: A new gradient balance approach for long-tailed object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1685–1694.
- [54] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [55] Dong Wang, Yicheng Liu, Liangji Fang, Fanhua Shang, Yuanyuan Liu, and Hongying Liu. 2022. Balanced Gradient Penalty Improves Deep Long-Tailed Learning. In *Proceedings of the 30th ACM International Conference on Multimedia*. 5093–5101.
- [56] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X Yu. 2020. Long-tailed recognition by routing diverse distribution-aware experts. *arXiv preprint arXiv:2010.01809* (2020).
- [57] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X Yu. 2021. Long-tailed recognition by routing diverse distribution-aware experts. *ICLR* (2021).
- [58] Yidong Wang, Bowen Zhang, Wenxin Hou, Zhen Wu, Jindong Wang, and Takahiro Shinozaki. 2021. Margin calibration for long-tailed visual recognition. *arXiv preprint arXiv:2112.07225* (2021).
- [59] Yu Wang, Yuying Zhao, Neil Shah, and Tyler Derr. 2022. Imbalanced graph classification via graph-of-graph neural networks. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2067–2076.
- [60] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. 2017. Learning to model the tail. *Advances in neural information processing systems* 30 (2017).
- [61] Yinwei Wei, Xiang Wang, Qi Li, Liqiang Nie, Yan Li, Xuanping Li, and Tat-Seng Chua. 2021. Contrastive learning for cold-start recommendation. In *Proceedings of the 29th ACM International Conference on Multimedia*. 5382–5390.
- [62] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM international conference on multimedia*. 1437–1445.
- [63] Liuyu Xiang, Guiguang Ding, and Jungong Han. 2020. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*. Springer, 247–263.
- [64] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. 2019. Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5704–5713.
- [65] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems* 33 (2020), 5812–5823.
- [66] Shuo Yu, Huafei Huang, Minh N Dao, and Feng Xia. 2022. Graph augmentation learning. In *Companion Proceedings of the Web Conference 2022*. 1063–1072.
- [67] Sukwon Yun, Kibum Kim, Kanghoon Yoon, and Chanyoung Park. 2022. LTE4G: Long-Tail Experts for Graph Neural Networks. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2434–2443.
- [68] Yuhang Zang, Chen Huang, and Chen Change Loy. 2021. Fasa: Feature augmentation and sampling adaptation for long-tailed instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3457–3466.
- [69] Yanfu Zhang, Shangqian Gao, Jian Pei, and Heng Huang. 2022. Improving social network embedding via new second-order continuous graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2515–2523.
- [70] Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. 2022. Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. *Advances in Neural Information Processing Systems* 35 (2022), 34077–34090.
- [71] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. 2021. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596* (2021).
- [72] Yusheng Zhao, Xiao Luo, Wei Ju, Chong Chen, Xian-Sheng Hua, and Ming Zhang. 2023. Dynamic Hypergraph Structure Learning for Traffic Flow Forecasting. *ICDE*.
- [73] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. 2020. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9719–9728.
- [74] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. 2022. Balanced Contrastive Learning for Long-Tailed Visual Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6908–6917.