# Less is More: Using Buffer Nodes to Reduce Excessive Majority Node Influence in Class Imbalance Graphs

Proceedings Track Submission

#### Anonymous Author(s)

Anonymous Affiliation Anonymous Email

Abstract

Graph Neural Networks (GNNs), despite success in node classification, struggle with class-imbalanced graphs, leading to minority node misclassification. Existing methods that synthesize minority nodes often overlook how majority nodes propagate misleading information through majority-minority edges; our analysis confirms this negative impact. To address this, we propose BufferGraph, a framework that inserts buffer nodes on such edges. These nodes act as controlled bottlenecks to reduce excessive majority node influence. And we theoretically demonstrate they reduce minority node feature distortion. Experiments on five real-world datasets show BufferGraph improves accuracy by up to 2% over state-of-the-art methods, excelling in imbalanced settings and for minority classes with high heterophily. Code is available at https://anonymous.4open.science/r/BufferGraph-C257.

#### 1 Introduction

6

8

9

10

11

12 13

20

21

23

24

27

28

34

35

36

37

38

Recent years have witnessed the rapid development of graph representations, especially GNNs [1–6]. Among various graph tasks, node classification has proven important and socially beneficial in real-world applications, such as identifying influencers on social networks [7–9] and detecting fraudsters in financial activities [10–12]. Naturally, the misclassification of minority classes can have harmful societal impacts. For example, minor fraudulent accounts that are not identified will continue to affect the reputation of e-commerce platforms and lead to loss of platform users [13, 14]. Despite the critical importance of accurate minority class prediction in real-world applications, current GNN approaches struggle with class-imbalanced graphs, creating a significant gap between model capabilities and practical needs.

Standard GNN models assume class balance and neighborhood homogeneity when aggregating features [1–3]. However, real-world graphs typically exhibit significant class imbalances—fraudulent accounts form a tiny minority in the Ethereum network [12, 15], and computer vision papers greatly outnumber computer architecture papers in citation networks [16]. These imbalances bias model performance toward majority classes [4, 17, 18], resulting in poor minority class prediction when GNNs are applied directly to imbalanced graphs.

To address this problem, previous methods have synthesized fake minority nodes and connected them to the original graph [19–24]. These approaches include techniques like GraphSMOTE [19], GraphENS [20], TAM [21], and GraphSHA [22], which generate synthetic minority nodes through various interpolation and mixing strategies. However, these approaches overlook a critical issue: majority nodes continue to propagate information to minority nodes through majority-minority edges, which we define as *misleading edges*. These edges distort minority node representations during message passing. Despite generating new nodes to balance classes globally, local neighborhoods often remain imbalanced, with minority nodes surrounded by majority nodes. As shown in Figure 1, higher percentages of majority neighbors correlate with lower prediction accuracy for minority classes, which restricts the performance of baselines.

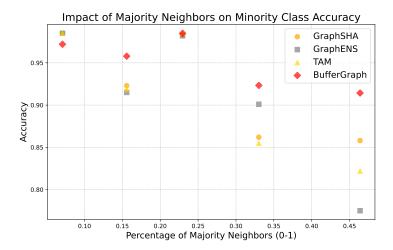


Figure 1: Impact of majority-class neighborhood composition on minority class accuracy on the Amazon-Computers dataset. Each point represents a minority class, plotting its classification accuracy against the percentage of its majority-class neighbors. While baseline methods struggle with accuracy degradation as majority-class neighbor percentage increases, our proposed BufferGraph maintains consistently superior performance across all neighborhood compositions, effectively addressing the challenge of majority-class interference in graph-based classification.

To address this fundamental limitation, we propose BufferGraph—a framework that explicitly regulates cross-class edge message propagation through adaptive edge intervention.

Our approach enhances minority class classification by decreasing message propagation along *misleading* edges—connections often dominated by majority nodes, which can distort the inherent features of minority instances as shown in Figure 2.

To achieve this goal, we first use predicted labels from the GNN pre-training stage to identify potential *misleading edges*. Edges connecting to predicted minority nodes are classified as non-problematic as their predictions have shown high precision [22]. For other potentially *misleading edges*, we insert buffer nodes to modulate message passing, allowing information to flow through both the original path and a new path via the buffer node. These

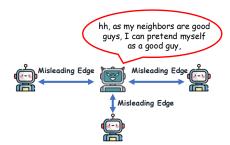
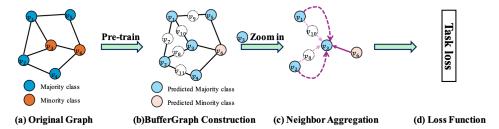


Figure 2: Misleading edge example.

unlabeled buffer nodes act as message passing modulators. The proportion of message passing through the original edge is determined by the difference in predicted labels of connected nodes, helping maintain prediction accuracy for majority-majority node pairs while reducing majority nodes' misleading influences on minority nodes. We provide the overview of our method in Figure 3.

We evaluate BufferGraph on five real-world datasets: Amazon-Photos, Amazon-Computers [25], Coauthor-CS, Coauthor-Physics [16], and WikiCS [26]. Experiments are conducted under both natural and artificially imbalanced settings to assess BufferGraph's effectiveness across varying imbalance ratios, comparing against state-of-the-art baselines.

Experiment results demonstrate that BufferGraph consistently outperforms the baseline models [19–22] in most configurations and datasets. For instance, under the random splitting scenario, BufferGraph exhibits its superiority by achieving a 2% increase in accuracy (from 88.39% to 90.22%), a 2% enhancement in balanced accuracy (from 89.23% to 91.85%), and a 2% boost in F1-score (from 87.55% to 89.30%) on the Amazon-Computers dataset compared to the second-best outcomes. Conversely, within the imbalanced splitting framework, BufferGraph continues to excel, marking a 2% improvement in accuracy, a 1% gain in balanced accuracy, and a 1.5% increase in F1-score on the



**Figure 3:** BufferGraph overview where  $v_1$ ,  $v_2$ ,  $v_4$ ,  $v_5$  are of the major class and  $v_3$ ,  $v_6$  are of the minor class. The input graph is shown in (a). After the pre-training using the GCN model, the nodes are predicted into majority or minority classes. Subsequently, we introduce a buffer node into each edge except those minority nodes' neighboring edges, as depicted in (b). The feature of the buffer node is a mixup of the features from the two nodes connected by the edge. Then, we zoom in  $v_3$  to show the neighbor aggregation of BufferGraph in (c). For edges with buffer nodes, messages pass both through the buffer node and directly to  $v_3$ , with direct edge weights determined by the predicted label differences between nodes. Loss calculation during BufferGraph neighbor aggregation is shown in (d).

Amazon-Computers dataset relative to the runner-up results. These findings underscore BufferGraph's adaptability and efficacy in addressing class imbalances. We make the following contributions:

- Heterophily-Aware Buffer Framework: We identify a critical issue overlooked by previous work the negative impact of majority-minority edges on message passing in imbalanced graphs. To address this, we propose BufferGraph, which strategically introduces buffer nodes to regulate information flow through misleading connections.
- Theoretical Analysis: We establish the mathematical foundation explaining how buffer nodes preserve minority class features by controlling eigenvalue decay during message passing, providing insights into why our approach effectively addresses class imbalance through message regulation.
- **Experimental Validation**: We conduct comprehensive experiments on multiple real-world datasets. Results demonstrate BufferGraph's consistent superiority over state-of-the-art methods, particularly in improving performance on minority classes across varying imbalance ratios. We also conduct comprehensive ablation studies on BufferGraph to assure each component is necessary.

#### 2 Proposed Method

Figure 3 illustrates our approach, which consists of three key components: (1) pre-training a GCN to identify potential misleading edges, (2) inserting buffer nodes into these edges to create controlled information bottlenecks, and (3) implementing an adaptive message-passing mechanism that adjusts information flow based on node similarity.

#### 2.1 Background

72

74

75

76

77

78

79

80

81

82

83

84

85

86

89

Notations. We address the challenge of class-imbalanced node classification on an unweighted, undirected graph  $\mathcal{G}=(\mathcal{V},\mathcal{E})$ , where  $\mathcal{V}=\{v_1,\cdots,v_N\}$  represents the set of N nodes, and  $\mathcal{E}\subseteq\mathcal{V}\times\mathcal{V}$  denotes the edges. The graph structure is captured by an adjacency matrix  $\mathbf{A}\in\{0,1\}^{N\times N}$ , with  $\mathbf{A}_{ij}=1$  indicating an edge between nodes  $v_i$  and  $v_j$ . Node features are represented by a matrix  $\mathbf{X}\in\mathbb{R}^{N\times d}$ , where each row  $\mathbf{X}_i\in\mathbb{R}^d$  corresponds to the d-dimensional features of node  $v_i$ . Each node  $v_i$  is labeled with one of  $v_i$  classes,  $v_i$  corresponds to the  $v_i$  corresponds to the  $v_i$  corresponds in class  $v_i$  corresponds to the  $v_i$  corresp

**Definition.** *Majority-minority edges*, also referred to as *misleading edges*, are defined as the edges connecting nodes from the majority class to nodes from the minority class in a graph. Formally, given a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with a label function  $\mathbf{Y} : \mathcal{V} \to \{1, \cdots, C\}$ , let M denote the set of nodes in the

majority class and m denote the set of nodes in the minority class. An edge  $e=(v_i,v_j)\in\mathcal{E}$  is a majority-minority edge if  $v_i\in M$  and  $v_j\in m$  or  $v_i\in m$  and  $v_j\in M$ .

#### 2.2 Buffer Node Generation

104

118

122

124

125

126

Buffer nodes represent our key innovation for addressing class imbalance in GNNs by modulating message passing between majority and minority classes through strategically introduced buffer nodes, rather than generating new minority nodes.

A buffer node, denoted as  $v_{\text{buf}}$ , does not possess a label but is characterized by a feature vector  $X_{\text{buf}}$ .

This vector is derived by interpolating the features of two adjacent nodes,  $v_{\text{a}}$  and  $v_{\text{b}}$ , connected by an undirected edge. The interpolation is governed by the equation [27]:

$$\boldsymbol{X}_{\text{buf}} = \alpha \boldsymbol{X}_{\text{a}} + (1 - \alpha) \boldsymbol{X}_{\text{b}}, \quad \alpha \in [0, 1]. \tag{1}$$

In this formula,  $\alpha$  acts as the mixup coefficient, influencing the degree to which the features of node  $v_a$  and node  $v_b$  affect the buffer node's features. A lower value of  $\alpha$  biases the feature vector  $\boldsymbol{X}_{buf}$  towards the features of  $v_b$ , whereas a higher value biases it towards those of  $v_a$ .

The buffer nodes integrated into edges serve primarily to regulate message passing rather than to enhance feature representation. Given this focused role, we initially set  $\alpha=1/2$  uniformly across all edges, prioritizing methodological consistency in our preliminary investigation. A comprehensive analysis of  $\alpha$ 's impact through systematic parameter variation is presented in Section 3.5.

#### 2.3 BufferGraph Framework

To address the challenges of message passing across majority-minority edges within class-imbalanced graphs, BufferGraph implements a dynamic message passing mechanism that precisely modulates the flow of information. We illustrate the details in Figure 3 as follows:

**Pre-training Stage.** We first pre-train a GCN model using training sets label information to predict the labels of all nodes in the original graph. These predicted labels serve as a guide for identifying potential problematic edges, which are then targeted for buffer node insertion. According to our algorithm, edges connected to predicted minority nodes are considered non-problematic and left unchanged, as previous work [22] has shown high precision in minority node predictions.

**Difference Score.** After pre-training, we compute a difference score  $s_{(v_i,v_j)}$  for each edge  $e_{(v_i,v_j)}$ , based on the Manhattan distance between the predicted label distributions of the connected nodes:

$$s_{(v_i,v_j)} = \sum_{k} |\hat{y}_{i,k} - \hat{y}_{j,k}| \tag{2}$$

where  $\hat{y}_{i,k}$  and  $\hat{y}_{j,k}$  are the predicted label distributions for nodes  $v_i$  and  $v_j$  respectively. This score quantifies the dissimilarity between nodes' predicted class distributions.

Adaptive Message Passing. For edges in the original graph, we modulate the message passing by adjusting the edge weight  $w_{(v_i,v_j)}$  according to:

$$w_{(v_i,v_j)} = 1 - s_{(v_i,v_j)} \tag{3}$$

For edges connected to buffer nodes, we maintain their weights at 1. This design ensures that when there is a large label distribution difference (high  $s_{(v_i,v_j)}$ ) between connected nodes, more information flows through the buffer node path rather than the direct edge. This mechanism effectively reduces the direct influence of majority nodes on their minority neighbors while preserving necessary information exchange through the buffer nodes.

Neighbor Aggregation. The GNN neighbor aggregation mechanism we utilize follows the standard form [22]:

$$oldsymbol{H}_{t}^{(l)} \leftarrow ext{Transform} \left( ext{Propagate} \left( oldsymbol{H}_{s}^{(l-1)}; oldsymbol{H}_{t}^{(l-1)} \right) 
ight),$$
 (4)

where  $m{H}_t^{(l)}$  represents node embedding of node  $v_t$  in the l-th layer.

A GCN can be formalized with the following operations [1]:

**Propagate:** 
$$M_t^{(l)} = \hat{A}H_t^{(l-1)},$$
 (5)

**Transform:** 
$$H_t^{(l)} = \text{ReLU}(M_t^{(l)}W_t^{(l)}),$$
 (6)

where  $W^{(l)}$  is the learnable weight matrix of the l-th layer, and ReLU(·) is the activation function.

Training Objective. During training, we calculate the validation loss every epoch and re-evaluate the 143 difference scores if the current validation loss improves over the best previous loss. Our loss function 145 combines two components:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pred}} + \lambda \cdot \mathcal{L}_{\text{diff}}. \tag{7}$$

where  $\lambda$  is a hyperparameter that balances the prediction loss and difference loss. While  $\lambda$  can 146 theoretically take any positive real value, we primarily consider  $\lambda \in [0.1, 1.0]$  to maintain a balanced 147 contribution between the two loss terms. We also explore larger values of  $\lambda$  in our parameter 148 sensitivity analysis (Section 3.5). 149

Algorithm 1 details the complete BufferGraph implementation. And complexity analysis of Buffer-Graph can be found in Appendix D. 151

#### **Theoretical Analysis** 2.4

152

This section provides theoretical justification for how buffer node insertion enhances minority node classification by modifying graph spectral properties and improving information propagation. First, 154 we establish the mathematical foundation of GNN propagation and identify the key challenge of 155 eigenvalue decay. Second, we demonstrate how buffer nodes address this challenge through controlled 156 message passing. Third, we prove that this mechanism effectively preserves minority class features.

**• GNN Propagation and Eigenvalue Decay**: The feature update rule in standard GNNs for a single layer  $\ell$  is: 159

$$H^{(\ell+1)} = \sigma\left(\hat{S}H^{(\ell)}W^{(\ell)}\right),\tag{8}$$

where  $\hat{S} = \hat{D}^{-1/2} \hat{A} \hat{D}^{-1/2}$  is the symmetrically normalized adjacency matrix,  $H^{(\ell)}$  represents node 160 features, and  $W^{(\ell)}$  is a trainable weight matrix. Multi-layer propagation effectiveness is governed by 161 the eigenvalues ( $\mu$ ) of operator  $\hat{S}$ . A critical challenge in deep GNNs is that smaller eigenvalues of  $\hat{S}$ , 162 essential for preserving minority class features, decay exponentially through multiple propagation steps ( $\mu_{\rm eff}^{(k)} \sim \mu^{2k+1}$  after k layers) [28]. This rapid decay leads to feature smoothing and minority 163 164 information loss. 165

- **2** Buffer Node Impact on Propagation: When a buffer node b is inserted between nodes (u, v), it 166 creates two new edges (u, b) and (b, v), modifying the graph structure in two key ways:
- 1) Modified Propagation Operator: The augmented adjacency matrix  $\tilde{A} = \begin{bmatrix} A & B \\ B^{\top} & 0 \end{bmatrix}$  introduces 168

buffer connections through block matrix B, altering the graph Laplacian  $\dot{L}$  and propagation operator 169  $\tilde{S}$  while preserving original node degrees.

- 2) Controlled Information Flow: Buffer nodes  $b_{ij}$  between  $(v_i, v_j)$  create two-hop paths with attenuation factor  $\alpha_{ij}$ , regulating message passing via  $\tilde{h}_{b_{ij}}^{(l)} = \alpha_{ij}(W^{(l)}h_i^{(l-1)} + W^{(l)}h_j^{(l-1)})$ . 172
- **The Example 2** Feature Preservation Mechanism: BufferGraph creates a modified propagation operator  $\hat{S}'$  with 173 more favorable spectral properties, achieving an effect similar to specialized GNN architectures with 174 slower eigenvalue decay: 175

$$\mu_{\text{eff}}^{(k)} \sim C(\alpha) \cdot \mu^{k+1},\tag{9}$$

 $\mu_{\rm eff}^{(k)} \sim C(\alpha) \cdot \mu^{k+1}, \tag{9}$  where  $C(\alpha)$  is a factor dependent on architectural parameters. This improvement is achieved through two key mechanisms:

1) Reduced Decay Rate: Effective eigenvalue decay improves from  $\mathcal{O}(\mu^{2k+1})$  to  $\mathcal{O}(\mu^{k+1})$ :

$$\lambda_{\text{new}}^{(l)} = \lambda_{\text{orig}}^{(l)} \cdot \underbrace{\left(1 + \alpha \frac{\langle h_i, h_j \rangle}{\|h_i\| \|h_j\|}\right)}_{C(\alpha) \ge 1}$$

$$(10)$$

2) Feature Preservation: The scaling factor  $C(\alpha)$  amplifies components orthogonal to majority-class features, maintaining minority feature magnitudes during propagation.

**Implications.** These theoretical insights underscore BufferGraph's broad applicability. For *highly imbalanced graphs*, where minority features risk being overshadowed by dominant majority signals, BufferGraph's reduced eigenvalue decay and targeted feature amplification are crucial. These mechanisms preserve minority feature integrity and mitigate data skew, leading to more accurate minority node identification. In *moderately imbalanced or relatively balanced graphs*, even with less severe feature suppression, BufferGraph's controlled message passing from majority to minority nodes and favorable spectral properties enhance model robustness and refine feature learning. This nuanced aggregation benefits overall classification, positioning BufferGraph as both a specialized solution for severe imbalance and a generally beneficial graph learning framework.

## 3 Experiments

This section focuses on answering the following research questions (RQs): (RQ1) How does Buffer-Graph's performance in node classification on naturally class-imbalanced graphs compare to that of existing baseline models (§3.2)? (RQ2) How effectively does BufferGraph outperform other baseline models in node classification across graphs with varying class-imbalance ratios (§3.3)? (RQ3) How does each component contribute to the overall performance of BufferGraph (§3.4)? (RQ4) How does BufferGraph's performance change when different hyperparameter settings are used (§3.5)? (RQ5) Does our inserting ratio affect the performance of BufferGraph (§3.6)?

#### 3.1 Experimental Setup

**Datasets.** To comprehensively evaluate BufferGraph, we conduct experiments across five naturally class-imbalanced datasets: Amazon-Photos, Amazon-Computers [25], Coauthor-CS, Coauthor-Physics [16], and WikiCS [26]. Table 5 summarizes the key statistics of these datasets. The Max/Min ratio represents the number of samples in the largest majority class to that in the smallest minority class, highlighting the natural class imbalance in these datasets.

**Baselines.** We compare BufferGraph against three categories of GCN-backboned baselines: (1) **Loss management strategies** including Reweight [29], PC SoftMax [30], Cross Entropy, and Balanced SoftMax [31]; (2) **Class-imbalanced node classification** methods including GraphSMOTE [19], GraphENS [20], TAM [21], and GraphSHA [22]; (3) **Heterophilic GCN models** including MixHop [32] and H2GCN [33], which handle heterophily similar to majority-minority edges. More details can be found in Appendix F.

Evaluation Metrics. We use four widely adopted metrics for class-imbalanced node classification: Accuracy (Acc.) measuring overall classification accuracy across all nodes; Balanced Accuracy (BAcc.) as the average of per-class accuracy giving equal weight to each class; Macro F1 Score (F1) balancing precision and recall across all classes; and Standard Deviation reported as mean±std across five runs with different random seeds.

Implementation Details. For all experiments, we use the following settings: (1) Random Splitting with dataset division into training/validation/testing sets (6:2:2 ratio); (2) Model Architecture featuring three hidden layers (256 hidden dimensions each); (3) Training Parameters including learning rate 0.01, dropout rate 0.4, and up to 5000 epochs with early stopping; (4) BufferGraph-Specific Parameters with mixup coefficient  $\alpha = 0.5$  and difference loss weight  $\lambda = 1.0$ . For the imbalanced setting, we follow [20] by downsampling the last half of classes in the training set to achieve an imbalance ratio of 10, while maintaining the original distribution in validation and testing sets. For hardware, we provide the information in Appendix E.

#### 3.2 RQ1: Performance on Naturally Imbalanced Graphs

**Random Splitting Results.** Tables 1 and 2 demonstrate BufferGraph's consistent superiority over baselines across five datasets. On Amazon-Computers, BufferGraph achieves 90.22% accuracy (+2% vs 88.39%), 91.85% balanced accuracy (+2% vs 89.23%), and 89.30% F1-score (+2% vs 87.55%). The framework shows particular strength on minority classes (evidenced by high balanced accuracy) while maintaining majority class performance, with consistent improvements across datasets of varying characteristics, confirming its robustness. To further validate BufferGraph's effectiveness,

we present a comparative analysis in Figure 1. The results reveal that BufferGraph significantly improves the performance of minority classes, particularly those with a high percentage of majority-class neighbors, while maintaining strong performance across all classes. This balanced improvement demonstrates BufferGraph's ability to effectively address the challenges of imbalanced node classification.

230

231

234

236

240

241

244

246

247

248

Imbalanced Splitting Results. Tables 3 and 4 show BufferGraph's performance under artificially imbalanced settings with an imbalance ratio of 10. BufferGraph maintains its superior performance, achieving approximately 2% improvement in accuracy, 1% in balanced accuracy, and 1.5% in F1-score on Amazon-Computers compared to the runner-up results. While baseline methods like Reweight show good balanced accuracy but lower overall accuracy (indicating they improve minority class performance at the cost of majority class accuracy), BufferGraph achieves strong performance across all metrics, demonstrating its ability to effectively balance minority and majority class prediction.

**Analysis.** BufferGraph's superior performance stems from two key innovations: First, buffer nodes transform direct one-hop neighbors into two-hop neighbors, reducing the negative influence of majority nodes on minority nodes. Second, adaptive message passing modulates information flow based on node similarity, preserving minority class features while maintaining majority class performance.

**Table 1:** Random splitting experiment results of BufferGraph and other baselines on Amazon-Photos, Amazon-Computers, and Coauthor-CS. We report all metrics with the standard deviation errors for five repetitions. The best result is highlighted by bold text. The runner-up result is highlighted by the underline.

	Dataset	Dataset Amazon-Photos		s	An	nazon-Comput	ers	Coauthor-CS		
	Random Splitting	Acc.	BAcc.	F1	Acc.	BAcc.	F1	Acc.	BAcc.	F1
	Vanilla	92.44±0.16	90.41±0.62	91.20±0.27	87.71±0.36	82.34±1.36	83.03±1.86	92.89±0.41	89.97±0.45	90.70±0.63
Methods	Reweight	92.91±0.36	92.51±0.41	91.72±0.25	86.21±0.71	89.23±0.19	85.09±0.68	92.86±0.03	$90.86\pm0.13$	91.09±0.04
	PC Softmax	91.83±0.33	91.83±0.34	90.31±0.48	87.13±1.49	87.60±0.67	85.09±1.71	92.95±0.12	$91.87\pm0.11$	91.21±0.11
	Balanced Softmax	93.46±0.05	91.19±0.01	92.05±0.05	86.28±0.02	87.42±0.29	83.91±0.02	93.46±0.05	$91.19\pm0.01$	92.05±0.05
	Cross Entropy	93.40±0.20	92.13±0.21	92.30±0.23	86.99±0.22	82.26±0.75	84.21±0.50	93.05±0.14	$90.85\pm0.37$	91.49±0.32
Me	MixHop	92.85±0.42	91.77±1.36	91.20±0.47	85.02±0.26	$75.17\pm1.58$	75.09±1.13	91.97±0.58	$88.50\pm0.49$	$89.41\pm0.43$
	H2GCN	93.06±0.53	92.47±0.27	91.79±0.21	88.39±0.55	$88.01\pm0.32$	87.55±1.18	94.24±0.14	$92.77\pm0.34$	$92.75\pm0.60$
	GraphSmote	89.72±0.45	90.69±0.57	88.90±0.48	85.36±0.72	84.79±1.22	85.22±0.98	87.44±0.24	85.08±0.63	84.26±0.52
	GraphENS	93.37±0.42	92.18±0.36	91.63±0.46	86.35±0.71	87.66±0.54	85.81±0.47	91.65±0.23	90.72±0.39	89.53±0.36
	TAM	90.13±0.33	90.98±0.36	89.15±0.49	85.46±0.11	88.51±0.67	84.52±0.26	92.41±0.04	90.84±0.01	91.35±0.02
	GraphSHA	<u>93.63</u> ±0.23	<u>92.61</u> ±0.66	<u>92.60</u> ±0.38	82.98±0.17	77.73±1.90	79.10±2.22	92.68±0.59	91.00±0.37	90.94±0.51
	BufferGraph	<b>94.47</b> ±0.10	<b>94.28</b> ±0.10	<b>93.12</b> ±0.07	<b>90.22</b> ±0.48	<b>91.85</b> ±0.34	<b>89.30</b> ±0.69	<b>94.90</b> ±0.28	<b>93.88</b> ±0.39	<b>93.70</b> ±0.41

**Table 2:** Random splitting experiment results of BufferGraph on Coauthor-Physics and WikiCS.

	Dataset	C	oauthor-Physic	cs		WikiCS	
	Random Splitting	Acc.	BAcc.	F1	Acc.	BAcc.	F1
	Vanilla	96.22±0.24	94.60±0.42	93.49±0.03	83.20±0.23	80.34±0.41	80.63±0.07
Methods	Reweight PC Softmax Balanced Softmax Cross Entropy	$\begin{array}{c c} 95.70 \pm 0.02 \\ 96.14 \pm 0.07 \\ 96.16 \pm 0.03 \\ \underline{96.50} \pm 0.14 \end{array}$	$\begin{array}{c} 95.06 \pm 0.05 \\ 95.36 \pm 0.11 \\ \underline{95.52} \pm 0.06 \\ 95.30 \pm 0.13 \end{array}$	$\begin{array}{c} 94.52 {\pm} 0.01 \\ 95.07 {\pm} 0.11 \\ 95.05 {\pm} 0.02 \\ \underline{95.25} {\pm} 0.18 \end{array}$	$ \begin{vmatrix} 82.66 \pm 0.17 \\ 82.76 \pm 0.32 \\ \underline{83.83} \pm 0.49 \\ 83.15 \pm 0.23 \end{vmatrix} $	$\begin{array}{c} \underline{82.82} {\pm} 0.09 \\ 81.94 {\pm} 0.50 \\ 82.15 {\pm} 0.51 \\ 82.63 {\pm} 0.74 \end{array}$	$80.92\pm0.20$ $80.46\pm0.45$ $\underline{81.74}\pm0.72$ $81.57\pm0.30$
Met	MixHop H2GCN	94.77±0.69 95.17±0.33	$92.32 \pm 0.86$ $92.33 \pm 0.25$	$93.15\pm1.09$ $93.25\pm0.81$	77.07±1.11 79.77±2.23	$67.96\pm2.02$ $69.69\pm2.62$	$68.57{\pm}1.95$ $72.90{\pm}2.67$
	GraphSmote GraphENS TAM GraphSHA BufferGraph	95.09±0.53 95.46±0.09 95.35±0.19 96.27±0.14 <b>96.78</b> ±0.07	93.01±0.66 95.32±0.04 95.04±0.08 95.51±0.17 <b>96.18</b> ±0.07	93.42±0.76 94.27±0.06 94.04±0.20 95.05±0.11 <b>95.34</b> ±0.06	82.94±0.70 81.78±0.06 80.73±0.34 81.83±1.22 <b>84.47</b> ±0.22	80.42±1.14 80.87±0.12 79.02±0.21 80.76±0.54 <b>84.34</b> ±0.21	80.65±0.75 79.80±0.08 79.02±0.16 78.95±1.53 <b>81.92</b> ±0.08

#### 3.3 RQ2: Performance Across Varying Imbalance Ratios

**Experimental Setup.** To evaluate BufferGraph's robustness to different levels of class imbalance, we artificially adjust the imbalance ratio to 15, 20, and 25 on the Amazon-Computers dataset.

**Results and Analysis.** Figure 4 shows that BufferGraph maintains consistently superior balanced accuracy across all imbalance ratios, outperforming the second-best baseline by at least 3% in each

**Table 3:** Imbalanced splitting experiment results of BufferGraph and other baselines on Amazon-Photos, Amazon-Computers, and Coauthor-CS.

	Dataset	Dataset Amazon-Photos		An	nazon-Comput	ers	Coauthor-CS			
	$\rho = 10$	Acc.	BAcc.	F1	Acc.	BAcc.	F1	Acc.	BAcc.	F1
	Vanilla	92.20±0.54	89.60±0.05	90.41±0.14	83.40±0.29	69.71±0.28	70.79±0.43	92.54±0.55	89.86±0.68	90.53±0.63
Methods	Reweight PC Softmax Balanced Softmax Cross Entropy	92.65±0.36 84.51±0.86 92.81±0.20 91.67±0.16	$\begin{array}{c} 92.34{\pm}0.17 \\ 88.69{\pm}1.27 \\ \underline{93.33}{\pm}0.04 \\ 87.85{\pm}0.40 \end{array}$	90.79±0.36 84.01±2.57 91.44±0.19 89.93±0.35	86.46±0.20 70.48±1.09 87.55±0.24 87.46±0.18	$\begin{array}{c} 89.26{\pm}0.08 \\ 84.92{\pm}1.21 \\ \underline{89.31}{\pm}0.16 \\ 83.49{\pm}0.53 \end{array}$	$85.33\pm0.14$ $70.50\pm0.46$ $86.95\pm0.02$ $85.19\pm0.53$	$93.23\pm0.12$ $92.78\pm0.02$ $93.99\pm0.01$ $94.04\pm0.07$	$\begin{array}{c} 91.74{\pm}0.07 \\ 93.16{\pm}0.06 \\ \underline{93.24}{\pm}0.03 \\ 92.03{\pm}0.03 \end{array}$	91.86±0.03 91.23±0.14 92.35±0.02 92.38±0.04
Met	MixHop H2GCN	91.57±0.84 92.49±0.70	90.46±0.54 91.04±0.29	89.53±0.72 91.55±0.24	84.56±1.15 85.22±0.46	75.48±2.35 85.19±0.30	75.30±2.52 85.55±0.83	88.65±0.51 93.98±0.36	80.59±1.24 92.61±0.36	83.19±1.03 92.33±0.34
	GraphSmote GraphENS TAM GraphSHA BufferGraph	88.31±0.63 92.55±0.07 91.08±0.03 <u>93.56</u> ±0.04 <b>93.91</b> ±0.18	88.15±1.53 91.66±0.37 91.70±0.07 92.46±0.30 <b>93.40</b> ±0.20	87.27±0.28 91.07±0.02 90.15±0.07 <u>92.59</u> ±0.02 <b>92.90</b> ±0.07	85.30±0.66 85.50±0.58 85.79±0.18 85.24±0.52 <b>89.51</b> ±0.35	84.66±0.27 89.21±0.14 88.21±0.69 83.77±0.55 <b>90.54</b> ±0.57	84.35±0.23 85.05±0.69 85.21±0.42 83.31±0.59 <b>88.14</b> ±0.40	88.95±0.19 92.12±0.03 92.53±0.04 92.42±0.16 <b>94.06</b> ±0.02	83.96±0.99 90.49±0.01 90.45±0.13 90.43±0.46 <b>93.78</b> ±0.03	85.56±0.78 89.21±0.16 90.67±0.13 90.21±0.22 <b>92.72</b> ±0.35

**Table 4:** Imbalanced splitting experiment results of our model BufferGraph and other baselines on two class-imbalanced node classification benchmark datasets.

	Dataset	C	oauthor-Physic	cs			
	$\rho$ =10	Acc.	BAcc.	F1	Acc.	BAcc.	F1
	Vanilla	95.65±0.04	93.76±0.12	94.19±0.17	81.30±1.00	75.16±1.53	77.42±1.55
Methods	Reweight PC Softmax Balanced Softmax Cross Entropy	96.35±0.04 95.18±0.09 <b>96.46</b> ±0.05 96.12±0.01	95.12±0.20 95.47±0.08 95.46±0.03 94.53±0.11	95.16±0.09 93.87±0.13 95.32±0.04 94.93±0.03	81.16±0.13 76.01±2.24 82.14±0.04 82.44±0.23	$\begin{array}{c} 81.48 {\pm} 0.28 \\ 80.30 {\pm} 1.42 \\ \underline{82.45} {\pm} 0.21 \\ 78.07 {\pm} 0.51 \end{array}$	79.54±0.34 73.85±2.13 <u>80.10</u> ±0.08 80.06±0.10
Met	MixHop H2GCN	93.88±0.32 93.09±0.17	91.06±0.82 88.60±0.14	91.98±0.27 90.02±0.56	76.91±1.41 77.85±1.83	66.24±0.49 66.30±2.75	66.90±0.94 68.01±0.93
	GraphSmote GraphENS TAM GraphSHA BufferGraph	92.64±0.36 95.35±0.19 95.25±0.06 96.27±0.14 96.26±0.17	92.79±0.11 95.04±0.08 94.11±0.19 95.14±0.04 <b>96.13</b> ±0.14	94.42±0.53 94.04±0.20 93.66±0.16 94.94±0.14 <b>96.15</b> ±0.27	74.96±1.07 80.73±0.30 81.12±0.04 <u>82.60</u> ±0.30 <b>83.42</b> ±0.52	69.43±2.17 79.94±0.27 80.29±0.03 80.34±0.31 <b>84.21</b> ±0.35	$70.82\pm1.93$ $77.83\pm0.45$ $79.32\pm0.12$ $80.00\pm0.52$ $81.58\pm0.45$

case. While all methods show some performance degradation as the imbalance ratio increases, Buffer-Graph exhibits the most stable performance, demonstrating its robustness to severe class imbalance. This stability can be attributed to BufferGraph's buffer node mechanism, which effectively modulates message passing between majority and minority classes regardless of their relative proportions.

### 3.4 RQ3: Ablation Study

**Experimental Setup.** To assess the contribution of each component in BufferGraph, we conduct ablation studies on Amazon-Computers and WikiCS datasets using the GCN backbone under the random splitting setting. BufferGraph uniquely combines buffer node insertion with a specialized message passing scheme and a difference loss. We examine four variants: the **Complete BufferGraph** (the full model); - **DL**, which ablates the difference loss component; - **CMP** (Concrete Message Passing), which removes the adaptively weighted direct message path between the original connected nodes, thereby routing all interaction through the buffer node pathway; and - **UMP** (Update Message Passing via Buffer), which ablates the message pathway through the inserted buffer node, relying only on the adaptively weighted direct path between original nodes.

Results and Analysis. Figure 5 shows that removing the update message passing mechanism (- UMP) leads to the most significant performance drop across all metrics, particularly in F1 scores. This confirms that adaptive message passing is the core component of BufferGraph, as it directly controls how information flows between majority and minority nodes. The removal of other components also notably impacts model performance, demonstrating that each component plays an essential role in BufferGraph's effectiveness. The complete BufferGraph consistently outperforms all ablated variants, confirming the synergistic effect of combining all components.

#### 3.5 RQ4: Parameter Sensitivity Analysis

**Experimental Setup.** We investigate the sensitivity of BufferGraph to two critical hyperparameters: (1) **Mixup coefficient**  $\alpha$  controlling feature interpolation for buffer nodes (tested values: 0.1, 0.3, 0.5, 0.7, 0.9), and (2) **Difference loss weight**  $\lambda$  balancing prediction loss and difference loss (tested values: 0.1, 0.5, 1, 2, 5).

**Results and Analysis.** Figure 6 reveals that: (1) For  $\alpha$ , optimal performance occurs at 0.5 across both datasets, suggesting equal feature mixing is ideal for buffer nodes; (2) For  $\lambda$ , Amazon-Computers shows stable performance from 0.1 to 1 while WikiCS improves over this range, but both degrade significantly at  $\lambda \geq 2$ . These results demonstrate BufferGraph's robustness within moderate parameter ranges ( $\alpha = 0.5$ ,  $\lambda = 1.0$ ), while extreme values harm performance. The optimal settings balance feature mixing and loss weighting effectively across diverse datasets.

### 3.6 RQ5: Buffer Node Ratio Study

**Experimental Setup.** We investigate how the ratio of inserted buffer nodes affects model performance under the random splitting setting on Amazon-Computers and WikiCS datasets. We vary the number of buffer nodes in the graph to evaluate the sensitivity of our approach. Specifically, we randomly selected a certain percentage of all eligible edges to add buffer nodes varying from 0% to 100%.

**Results and Analysis.** Our experiments consistently show that higher insertion ratios lead to improved performance, with the best results achieved when buffer nodes are inserted for all potentially misleading edges (100% insertion ratio). This confirms that our strategy of comprehensive buffer node insertion is effective and that the benefits of modulating message passing outweigh any potential drawbacks from increased graph complexity. Detailed results are provided in Appendix G.5.

#### 4 Related Work

Existing methods to tackle class imbalance in graphs can be categorized into two main approaches. The first approach focuses on generating synthetic nodes and edges [18–20, 22, 34, 35], where methods like DRGCN [34] and GraphSMOTE [19] generate minority nodes to balance classes, while GraphSHA [22] focuses on generating 'hard' nodes of minority classes to enlarge the margin between majority and minority classes. The second approach modifies the learning process itself, such as TAM [21] which introduces connectivity-aware margins. Recently, BAT [36] addresses class imbalance from a topological perspective by theoretically identifying two fundamental phenomena that amplify class-imbalance bias, serving as a plug-and-play module for existing class rebalancing methods.

However, existing methods either overlook the direct influence of majority classes on minority classes through majority-minority edges (empirically demonstrated in Figure 1), or function merely as auxiliary modules without providing a complete solution. In contrast, BufferGraph offers a comprehensive end-to-end model that fundamentally redesigns the message passing mechanism using buffer nodes to effectively address the class imbalance problem. We provide a comparison between previous works and BufferGraph to highlight our novelty in Table 4.

	Method	Handle Majority-Minority Edges	Label-Free Node Generation	Message Propagation Adjustment	No Complex Parameter Tuning
	GraphSMOTE [19]	Х	Х	Х	Х
	GraphENS [20]	Х	X	Х	Х
308	TAM [21]	Х	✓	✓	Х
300	GraphSHA [22]	Х	X	Х	X
	DRGCN [34]	Х	Х	Х	Х
	BAT [36]	X	✓	X	✓
	BufferGraph	✓	✓	✓	✓

#### 5 Conclusion

We introduce BufferGraph, a lightweight framework that addresses class imbalance in graph neural networks through strategic buffer node insertion. Unlike previous approaches focusing on minority node generation, BufferGraph modulates message passing through potentially misleading edges, effectively reducing majority nodes' negative influence on minority nodes. Our experiments on five real-world datasets demonstrate BufferGraph's superior and robust performance, particularly for minority classes with high percentages of majority-class neighbors.

#### References

316

- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 1, 5
- [2] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [3] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017. 1
- [4] Fenyu Hu, Liping Wang, Shu Wu, Liang Wang, and Tieniu Tan. Graph classification by mixture of diverse experts. *arXiv preprint arXiv:2103.15622*, 2021. 1
- Dingyi Zeng, Wanlong Liu, Wenyu Chen, Li Zhou, Malu Zhang, and Hong Qu. Substructure
   aware graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
   volume 37, pages 11129–11137, 2023.
- [6] Zemin Liu, Xingtong Yu, Yuan Fang, and Xinming Zhang. Graphprompt: Unifying pre-training and downstream tasks for graph neural networks. In *Proceedings of the ACM Web Conference* 2023, pages 417–428, 2023. 1
- [7] Xinyu Huang, Dongming Chen, Dongqi Wang, and Tao Ren. Identifying influencers in social networks. *Entropy*, 22(4):450, 2020. 1
- [8] Jayati Bhadra, Amandeep Singh Khanna, and Alexei Beuno. A graph neural network approach for identification of influencers and micro-influencers in a social network:\* classifying influencers from non-influencers using gnn and gcn. In 2023 International Conference on Advances in Electronics, Communication, Computing and Intelligent Information Systems (ICAECIS), pages 66–71. IEEE, 2023.
- [9] Alexander K Taylor, Nuan Wen, Po-Nien Kung, Jiaao Chen, Violet Peng, and Wei Wang. Where does your news come from? predicting information pathways in social media. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2511–2515, 2023. 1
- [10] Dan Lin, Jiajing Wu, Qi Xuan, and K Tse Chi. Ethereum transaction tracking: Inferring
   evolution of transaction networks via link prediction. *Physica A: Statistical Mechanics and its* Applications, 600:127504, 2022. 1
- Kiarash Shamsi, Friedhelm Victor, Murat Kantarcioglu, Yulia Gel, and Cuneyt G Akcora.
   Chartalist: Labeled graph datasets for utxo and account-based blockchains. Advances in Neural
   Information Processing Systems, 35:34926–34939, 2022.
- Qian Wang, Zhen Zhang, Zemin Liu, Shengliang Lu, Bingqiao Luo, and Bingsheng He. Etgraph:
  A pioneering dataset bridging ethereum and twitter. *arXiv preprint arXiv:2310.01015*, 2023. 1
- Sunil Kumar Maurya, Xin Liu, and Tsuyoshi Murata. Not all neighbors are friendly: Learning to choose hop features to improve node classification. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pages 4334–4338, 2022. 1
- Jianke Yu, Hanchen Wang, Xiaoyang Wang, Zhao Li, Lu Qin, Wenjie Zhang, Jian Liao, and Ying Zhang. Group-based fraud detection network on e-commerce platforms. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5463–5475, 2023. 1
- Yuan Li, Bingqiao Luo, Qian Wang, Nuo Chen, Xu Liu, and Bingsheng He. Cryptotrade: A reflective llm-based agent to guide zero-shot cryptocurrency trading. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1094–1106, 2024. 1
- [16] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008. 1, 2, 6
- Min Liu, Siwen Jin, Luo Jin, Shuohan Wang, Yu Fang, and Yuliang Shi. Imbalanced nodes classification for graph neural networks based on valuable sample mining. In *Proceedings* of the 2022 6th international conference on electronic information technology and computer engineering, pages 1957–1962, 2022. 1
- Mengting Zhou and Zhiguo Gong. Graphsr: A data augmentation algorithm for imbalanced node classification. *arXiv preprint arXiv:2302.12814*, 2023. 1, 9

- Tianxiang Zhao, Xiang Zhang, and Suhang Wang. Graphsmote: Imbalanced node classification on graphs with graph neural networks. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 833–841, 2021. 1, 2, 6, 9, 14
- [20] Joonhyung Park, Jaeyun Song, and Eunho Yang. Graphens: Neighbor-aware ego network synthesis for class-imbalanced node classification. In *The Tenth International Conference on Learning Representations*, *ICLR* 2022. International Conference on Learning Representations (ICLR), 2022. 1, 6, 9, 14
- Jaeyun Song, Joonhyung Park, and Eunho Yang. Tam: topology-aware margin loss for class imbalanced node classification. In *International Conference on Machine Learning*, pages
   20369–20383. PMLR, 2022. 1, 6, 9, 14
- Wen-Zhi Li, Chang-Dong Wang, Hui Xiong, and Jian-Huang Lai. Graphsha: Synthesizing harder samples for class-imbalanced node classification. *arXiv preprint arXiv:2306.09612*, 2023. 1, 2, 4, 6, 9, 14
- Meilun Shi, Enguang Zuo, Hanwen Qu, and Xiaoyi Lv. Graphuc: Predictive probability-based graph-structured data augmentation model for class imbalance node classification. In 2024
   4th International Conference on Electronic Information Engineering and Computer Science (EIECS), pages 933–936. IEEE, 2024.
- Nan Chen, Zemin Liu, Bryan Hooi, Bingsheng He, Jun Hu, and Jia Chen. Nodeimport:
   Imbalanced node classification with node importance assessment. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 94–105, 2025.
- Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018. 2, 6
- [26] Péter Mernyei and Cătălina Cangea. Wiki-cs: A wikipedia-based benchmark for graph neural
   networks. arXiv preprint arXiv:2007.02901, 2020. 2, 6
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 4
- [28] Nicolas Keriven. Not too little, not too much: a theoretical analysis of graph (over) smoothing.
   Advances in Neural Information Processing Systems, 35:2268–2281, 2022. 5
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based
   on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, pages 9268–9277, 2019. 6, 14
- Youngkyu Hong, Seungju Han, Kwanghee Choi, Seokjun Seo, Beomsu Kim, and Buru Chang.
   Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6626–6636, 2021. 6,
   14
- Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems*, 33:4175–4186, 2020. 6, 14
- [32] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr
   Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional
   architectures via sparsified neighborhood mixing. In *international conference on machine learning*, pages 21–29. PMLR, 2019. 6, 14
- [33] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems*, 33:7793–7804, 2020. 6, 14
- [34] Min Shi, Yufei Tang, Xingquan Zhu, David Wilson, and Jianxun Liu. Multi-class imbalanced
   graph convolutional network learning. In *Proceedings of the Twenty-Ninth International Joint* Conference on Artificial Intelligence (IJCAI-20), 2020. 9
- Liang Qu, Huaisheng Zhu, Ruiqi Zheng, Yuhui Shi, and Hongzhi Yin. Imgagn: Imbalanced network embedding via generative adversarial graph networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1390–1398, 2021. 9

- Zhining Liu, Ruizhong Qiu, Zhichen Zeng, Hyunsik Yoo, David Zhou, Zhe Xu, Yada Zhu,
   Kommy Weldemariam, Jingrui He, and Hanghang Tong. Class-imbalanced graph learning
   without class rebalancing. In Forty-first International Conference on Machine Learning, 2024.
   9
- Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen
   Chang, and Doina Precup. Revisiting heterophily for graph neural networks. *Advances in neural information processing systems*, 35:1362–1375, 2022. 14

#### 7 A Dataset Details

428

429

430

431

432

437

We provide the statistics of datasets used in the paper in Table 5.

Dataset	#Nodes	#Edges	#Features	#Classes	#Max / Min
Amazon-Photos	7,650	119,081	745	8	5.86
Amazon-Computers	13,752	245,861	767	10	17.72
Coauthor-CS	18,333	163,778	6,805	15	35.05
Coauthor-Physics	34,493	247,962	8,415	5	6.33
WikiCS	11,701	216,123	300	10	9.08

**Table 5:** Statistics of datasets used in the paper.

And these datasets' licenses are as follows:

- Amazon-Photos/Computers: Sourced from Amazon product data under MIT license.
- Coauthor-CS/Physics: Academic collaboration networks from arXiv with CC-BY 4.0 license.
- WikiCS: Wikipedia article network dataset (v1.0) under CC-BY-SA 3.0 license.

#### **B** Notations in BufferGraph

We provide all notations we use in BufferGraph paper in Table 6.

Table 6: Key Notations Used in BufferGraph

Symbol	Description
$\mathcal{G}=(\mathcal{V},\mathcal{E})$	Graph with nodes $\mathcal V$ and edges $\mathcal E$
$\mathbf{A} \in \{0,1\}^{N \times N}$	Adjacency matrix
$oldsymbol{X} \in \mathbb{R}^{N  imes d}$	Node feature matrix
$v_{ m buf}$	Buffer node
$s_{(v_i,v_j)}$	Difference score between nodes $v_i$ and $v_j$
$w_{(v_i,v_j)}$	Edge weight between nodes $v_i$ and $v_j$

### 35 C Algorithm

We provide the full algorithm in Algorithm 1.

### D Complexity Analysis

Given N as the total number of nodes and E as the total number of edges within the graph, the creation of buffer nodes correlates with the subset of the node and edge sets, leading to a computational complexity of O(E). Generating augmented features for buffer nodes requires  $O(E \cdot d)$  time, with d representing the dimensionality of the node features. Additionally, forming augmented edges for the buffer nodes incurs a time complexity of O(E), engaging the whole graph's edge set. Hence,

### Algorithm 1: BufferGraph: Buffer Node Synthesis and Adaptive Message Passing

```
Input: Graph \mathcal{G} = (X, A), training set nodes \mathcal{V}^L and their labels Y^L, number of classes C,
               mixup coefficient \alpha, distance loss coefficient \lambda
    Output: Trained GNN model f_{\theta}
    /* Buffer Node Generation
                                                                                                                                        */
 1 for each edge (v_{src}, v_{tar}) in \mathcal{G} do
         \boldsymbol{X}_{\text{buf}} \leftarrow \alpha \boldsymbol{X}_{\text{src}} + (1 - \alpha) \boldsymbol{X}_{\text{tar}}
         Insert buffer node v_{\text{buf}} with feature X_{\text{buf}} into \mathcal{G} between v_{\text{src}} and v_{\text{tar}}
 4 end
    /* Pre-training Stage
5 Initialize vanilla GCN f_{\theta}
6 \hat{m{Y}} \leftarrow \texttt{PreTrainGCN}(\mathcal{G}, \mathcal{V}^L, m{Y}^L)
7 Identify minority nodes based on \hat{Y}
 8 Decide edges to insert buffer nodes in \mathcal{G} (exclude edges connected to predicted minority nodes)
    /* Training Stage
9 best\_val\_loss \leftarrow \infty
10 Define warmup_epochs
11 for epoch \leftarrow 1 to N do
         Compute node prediction loss \mathcal{L}_{pred}
12
13
         Compute difference loss \mathcal{L}_{\text{diff}}
         \mathcal{L}_{\text{total}} \leftarrow \mathcal{L}_{\text{pred}} + \lambda \cdot \mathcal{L}_{\text{diff}}
14
         Update f_{\theta} by minimizing \mathcal{L}_{\text{total}} with standard backpropagation
15
         if epoch > warmup_epochs then
16
              if val\_loss < best\_val\_loss then
17
                    best\_val\_loss \leftarrow val\_loss
18
19
                    for each edge e_{(v_i,v_j)} \in \mathcal{G} do
20
                         s_{(v_i,v_j)} \leftarrow \sum_k |\hat{y}_{i,k} - \hat{y}_{j,k}|
                         if edge e_{(v_i,v_j)} has a buffer node then
21
                              w_{(v_i,v_i)} \leftarrow 1 - s_{(v_i,v_i)}
22
                         end
23
                    end
24
25
               end
         end
26
27 end
   return trained model f_{\theta}
```

the total additional complexity introduced by our model is  $O(E+E\cdot d+E)$ . As demonstrated by successful experiments on extensive datasets such as Coauthor-Physics and WikiCS, BufferGraph can work well on the large datasets.

### **E** Experimental Environment

All models in our experiments were implemented using Pytorch 2.0.0 in Python 3.9.16, and run on a robust Linux workstation. This system is equipped with two Intel(R) Xeon(R) Gold 6226R CPUs, each operating at a base frequency of 2.90 GHz and a max turbo frequency of 3.90 GHz. With 16 cores each, capable of supporting 32 threads, these CPUs offer a total of 64 logical CPUs for efficient multitasking and parallel computing. The workstation is further complemented by a potent GPU setup, comprising eight NVIDIA GeForce RTX 3090 GPUs, each providing 24.576 GB of memory. The operation of these GPUs is managed by the NVIDIA-SMI 525.60.13 driver and CUDA 12.0, ensuring optimal computational performance for our tasks.

#### F Baselines

446

455

Loss Management Strategies:

- 1. Reweight [29]: This adjusts the weights of classes based on their proposion to the class samples in the dataset, aiming to counteract imbalance.
- 2. PC SoftMax [30]: This improves the model's probability calibration for multi-class classification, ensuring more precise probabilistic predictions for minority classes.
- Cross Entropy: It acts as a core loss function by quantifying the divergence between predicted and actual distributions.
- Balanced Softmax [31]: It is designed to minimize the generalization bound in multi-class softmax regression.

### • Heterophilious GCNs

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

474

476

478

495

496

501

- 1. MixHop [32]: It extracts the features from multi-hop neighbors to get more information.
- 2. H2GCN [33] [37]: It combines three key designs to address heterophily: (1) ego- and neighbor-embedding separation; (2) higher-order neighborhoods; (3) combination of intermediate representations.

#### • Class-imbalanced Node Classification Approaches:

- GraphSMOTE [19]: It synthesizes new instances in minority classes to directly tackle class imbalance.
- GraphENS [20]: It utilizes ensemble strategies to generate new instances in minority classes.
- TAM [21]: It introduces connectivity and distribution-aware margin to guide minority classes' classification.
- 4. GraphSHA [22]: It focus on generating hard minority samples to enlarge the margin between minority and majority classes.
- For general hyper-parameters such as learning rates and the number of layers, we explore layers of 2,3 and explore hidden dimensions of 64, 128, and 256. Among these, we select the layer and dimension that provide the best performance of each baseline on each dataset. For the learning rate, we explore the optimal settings by evaluating the performance at values of 0.001, 0.005, and 0.01. For the dropout rate, we explore the optimal settings for evaluating the performance of 0.1, 0.2, 0.3, 0.4, 0.5.
- For GraphSmote, we opt for the GraphSmote<sub>O</sub> variant, which is tailored to predict discrete values without necessitating pretraining, showcasing superior performance among its various versions [19].
- In the implementation of GraphENS, we set the feature masking rate k to 0.01 and the temperature  $\tau$  to 1 according to the paper's configuration [20].
- For TAM, we select the GraphENS-based iteration, which is identified as the most performant according to the findings reported in the corresponding paper. The default settings from the released code are utilized, with the coefficients for ACM  $\alpha$ , ADM  $\beta$ , and classwise temperature  $\phi$  set to 2.5, 0.5, and 1.2, respectively [21].
- Specifically for GraphSHA, we follow the parameter configurations detailed in the original study, employing the PPR version with a setting of  $\alpha=0.05$  and K=128 [22].

#### **G** Experiment Results

#### G.1 Imbalanced Splitting Experiment

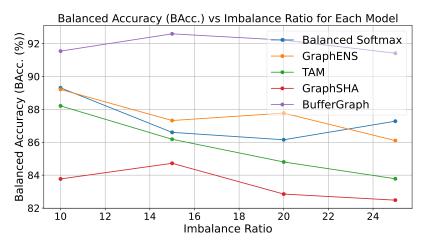
In this section, we provide the results of the imbalanced splitting experiment in Table 3 and Table 4.

#### 498 G.2 Imbalance Ratio Experiment Results

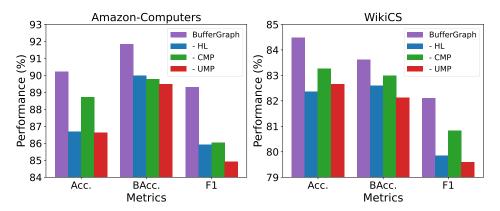
In this section, we provide the results of the imbalance ratio experiment on Amazon-Computers dataset in Figure 4.

#### **G.3** Ablation Study Results

In this section, we provide the results of the ablation study on Amazon-Computers and WikiCS in Figure 5.



**Figure 4:** Balanced accuracy (BAcc.) with increasing imbalance ratios on Amazon-Computers. BufferGraph maintains superior performance even as imbalance becomes more severe.



**Figure 5:** Ablation study on Amazon-Computers and WikiCS.

### **G.4** Parameter Sensitivity Experiment Results

In this section, we provide the results of the parameter sensitivity experiment on Amazon-Computers and WikiCS in Figure 6.

### **G.5** Insertion Ratio Experiment Results

504

505

506

507

508

509

We use the pre-training predicted results to identify the edges most likely to connect majority and minority nodes. We apply various ratios: 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0 to select the top potential edges for inserting buffer nodes and analyze their performance. For instance, a 0.1 ratio means that only the top 10% of potential majority-minority edges, excluding edges neighboring minority nodes, would be inserted with buffer nodes. To quantify the likelihood of an edge connecting majority-minority nodes, we calculate a score for each edge using the following formula:

$$\begin{split} \text{score}(e) &= \frac{1}{2} \bigg( \text{prob\_minority}(v_i) \cdot \text{prob\_majority}(v_j) \\ &+ \text{prob\_minority}(v_j) \cdot \text{prob\_majority}(v_i) \bigg) \end{split}$$

where  $e = (v_i, v_j) \in \mathcal{E}$  is an edge, prob\_minority(v) is the sum of the probabilities of node v belonging to minority classes, and prob\_majority(v) is the sum of the probabilities of node v belonging to majority classes. The scores are then normalized between 0 and 1:

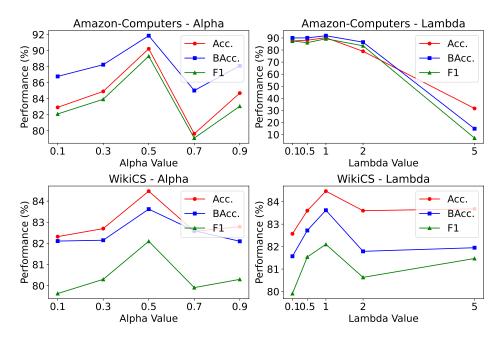


Figure 6: Parameter sensitivity analysis on Amazon-Computers and WikiCS.

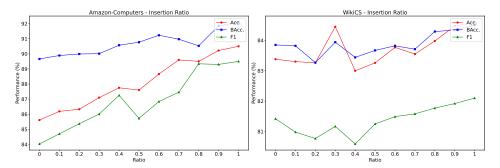


Figure 7: Insertion ratio experiments on Amazon-Computers and WikiCS.

$$\operatorname{normalized\_score}(e) = \frac{\operatorname{score}(e) - \min(\operatorname{score}(e))}{\max(\operatorname{score}(e)) - \min(\operatorname{score}(e))}$$

This edge scoring mechanism effectively identifies potential majority-minority connections, guiding the strategic insertion of buffer nodes. The experimental results show that higher insertion ratios consistently lead to better performance, validating the effectiveness of BufferGraph's design.

### **H** Limitations and Broader Impacts

518

519

520

521

522

523

524

527

528

**Limitations.** BufferGraph has two main limitations: (1) using a fixed mixup coefficient for all buffer nodes may not be optimal across different graph structures, and (2) its effectiveness is currently only validated for node classification, leaving other graph tasks (e.g., edge prediction) for future work.

**Broader Impact.** By improving minority class prediction in imbalanced graphs, BufferGraph has potential applications in domains where fairness and equity are crucial, such as social network analysis, recommendation systems, and fraud detection. Addressing class imbalance can help reduce algorithmic bias against underrepresented groups in graph-based applications. Future work could explore dynamic buffer node generation strategies, adaptive mixup coefficients based on node characteristics, and extensions to other graph learning tasks beyond node classification.