

DropMessage: Unifying Random Dropping for Graph Neural Networks

Taoran Fang¹, Zhiqing Xiao¹, Chunping Wang², Jiarong Xu³, Xuan Yang¹, Yang Yang^{1*}

¹Zhejiang University

²FinVolution Group

³Fudan University

{fangtr, yangya}@zju.edu.cn, wangchunping02@xinye.com, jiarongxu@fudan.edu.cn

Abstract

Graph Neural Networks (GNNs) are powerful tools for graph representation learning. Despite their rapid development, GNNs also face some challenges, such as *over-fitting*, *over-smoothing*, and *non-robustness*. Previous works indicate that these problems can be alleviated by *random dropping* methods, which integrate augmented data into models by randomly masking parts of the input. However, some open problems of random dropping on GNNs remain to be solved. First, it is challenging to find a universal method that are suitable for all cases considering the divergence of different datasets and models. Second, augmented data introduced to GNNs causes the incomplete coverage of parameters and unstable training process. Third, there is no theoretical analysis on the effectiveness of random dropping methods on GNNs. In this paper, we propose a novel random dropping method called *DropMessage*, which performs dropping operations directly on the propagated messages during the message-passing process. More importantly, we find that DropMessage provides a unified framework for most existing random dropping methods, based on which we give theoretical analysis of their effectiveness. Furthermore, we elaborate the superiority of DropMessage: it stabilizes the training process by reducing sample variance; it keeps information diversity from the perspective of information theory, enabling it become a theoretical upper bound of other methods. To evaluate our proposed method, we conduct experiments that aims for multiple tasks on five public datasets and two industrial datasets with various backbone models. The experimental results show that DropMessage has the advantages of both effectiveness and generalization, and can significantly alleviate the problems mentioned above. A detailed version with full appendix can be found on arXiv: <https://arxiv.org/abs/2204.10037>.

Introduction

Graphs, ubiquitous in the real world, are used to present complex relationships among various objects in numerous domains such as social media (social networks), finance (trading networks), and biology (biological networks). As powerful tools for representation learning on graphs, graph neural networks (GNNs) have attracted considerable attention recently (Defferrard, Bresson, and Vandergheynst 2016; Kipf

and Welling 2017; Velickovic et al. 2018; Ding, Tang, and Zhang 2018). In particular, GNNs adopt a message-passing schema (Gilmer et al. 2017), in which each node aggregates information from its neighbors in each convolutional layer, and have been widely applied in various downstream tasks such as node classification (Kipf and Welling 2017), link prediction (Kipf and Welling 2016), vertex clustering (Ramaswamy, Gedik, and Liu 2005), and recommendation systems (Ying et al. 2018).

Yet, despite their rapid development, training GNNs on large-scale graphs is facing several challenges such as *over-fitting*, *over-smoothing*, and *non-robustness*. Indeed, compared to other data forms, gathering labels for graph data is expensive and inherently biased, which limits the generalization ability of GNNs due to over-fitting. Besides, representations of different nodes in a GNN tend to become indistinguishable as a result of aggregating information from neighbors recursively. This phenomenon of over-smoothing prevents GNNs from effectively modeling the higher-order dependencies from multi-hop neighbors (Li, Han, and Wu 2018; Xu et al. 2018; Chen et al. 2020; Zhao and Akoglu 2020; Oono and Suzuki 2020, 2019). Recursively aggregating schema makes GNNs vulnerable to the quality of input graphs (Zhu et al. 2019; Zügner, Akbarnejad, and Günnemann 2018). In other words, noisy graphs or adversarial attacks can easily influence a GNN’s performance.

The aforementioned problems can be helped by *random dropping* methods (Hinton et al. 2012; Rong et al. 2019; Feng et al. 2020), which integrate augmented data into models by randomly masking parts of the input. These methods (Maaten et al. 2013; Matsuoka 1992; Bishop 1995; Cohen, Rosenfeld, and Kolter 2019) focus on randomly dropping or sampling existing information, and can also be considered as a data augmentation technique. Benefiting from the advantages of being unbiased, adaptive, and free of parameters, random dropping methods have greatly contributed to improving the performance of most GNNs.

However, some open questions related to random dropping methods on GNNs still exist. First, a general and critical issue of existing random dropping methods is that augmented data introduced to GNNs make parameters difficult to converge and the training process unstable. Moreover, it is challenging to find an optimal dropping method suitable to all graphs and models, because different graphs and models are equipped

*Corresponding author.




Figure 1: Illustrations of DropMessage and other existing random dropping methods. Considering the messages propagated by the center node (*i.e.*, Node 1), DropMessage allows to propagate distinct messages to different neighbor nodes, and its induced message matrix can be arbitrary. The induced message matrices of other methods obey some explicit constraints and can be regarded as special forms of DropMessage.

with their own properties and the model performance can be influenced greatly by employing various dropping strategies. Furthermore, the answer to how to choose a proper dropping rate when applying these methods is still unclear, and so far no theoretical guarantee has been provided to explain why random dropping methods can improve the performance of a GNN.

In this paper, we propose a novel random dropping method called DropMessage, which can be applied to all message-passing GNNs. As Figure 1 suggests, existing random dropping methods perform dropping on either the node feature matrix (Hinton et al. 2012; Feng et al. 2020) or the adjacency matrix (Rong et al. 2019), while our DropMessage performs dropping operations on the propagated messages, which allows the same node to propagate different messages to its different neighbors. Besides, we unify existing dropping methods into our framework and demonstrate theoretically that conducting random dropping methods on GNNs is equivalent to introducing additional regularization terms to their loss functions, which makes the models more robust. Furthermore, we also elaborate the superiority of our DropMessage whose sample variance is much smaller and training process is more stable. From the perspective of information theory, DropMessage keeps the property of information diversity, and is theoretically regarded as an upper bound of other random dropping methods. To sum up, the contributions of this paper are as follows:

- We propose a novel random dropping method, called DropMessage, for all message-passing GNNs. Existing random dropping methods on GNNs can be unified into our framework via performing masking in accordance with the certain rule on the message matrix. In other words, these methods can be regarded as one special form of DropMessage.
- We theoretically demonstrate the effectiveness of the random dropping methods, filling the gap in this field.
- We conduct sufficient experiments for different down-

stream tasks, and the experimental results show that DropMessage can better alleviate over-fitting, over-smoothing, and non-robustness compared to other existing random dropping methods.

Related Work

To sort out the key logic of our work, we first review some related work about random dropping methods with a particular focus on GNNs.

In general, random dropping can be regarded as a form of feature-noising schema that alleviates over-fitting by artificially corrupting the training data. As a representative work, Dropout is first introduced by Hinton et al. (Hinton et al. 2012) and has been proved to be effective in many scenarios (Abu-Mostafa 1990; Burges and Schölkopf 1996; Simard et al. 1998; Rifai et al. 2011; Maaten et al. 2013). Besides, Bishop (Bishop 1995) demonstrates the equivalence of corrupted features and L_2 -type regularization. Wager et al. (Wager, Wang, and Liang 2013) show that the dropout regularizer is first-order equivalent to an L_2 regularizer that being applied after scaling the features by an estimate of the inverse diagonal Fisher information matrix.

With the rapid development of GNNs, random dropping has also been generalized to the graph field, thus leading to three most common methods: *Dropout* (Hinton et al. 2012), *DropEdge* (Rong et al. 2019) and *DropNode* (Feng et al. 2020). Dropout performs random dropping operation on the node feature matrix, while DropEdge and DropNode, as the name implies, respectively act on the adjacency matrix (edges) and nodes. These random dropping methods can also be regarded as special forms of data augmentation (Shorten and Khoshgoftaar 2019; Frid-Adar et al. 2018; Buslaev et al. 2020; Ding et al. 2022; Velickovic et al. 2019), with the advantage of not requiring parameter estimation (Papp et al. 2021; Luo et al. 2021; Chen, Ma, and Xiao 2018; Zeng et al. 2020) and easy to apply. All the methods mentioned above can be used to alleviate *over-fitting* and *over-smoothing* on

GNNs. However, they can achieve effective performance only on some specific datasets and GNNs. The question of how to find an optimal dropping method that suitable for most cases still remains to be explored. Moreover, there is no theoretical explanation about the effectiveness of random dropping methods on GNNs, which adds some ambiguity to the function of these methods.

Notations and Preliminaries

Notations. Let $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ represent the graph, where $\mathbf{V} = \{v_1, \dots, v_n\}$ denotes the set of n nodes, and $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$ is the set of edges between nodes. The node features can be denoted as a matrix $\mathbf{X} = \{x_1, \dots, x_n\} \in \mathbb{R}^{n \times c}$, where x_i is the feature vector of the node v_i , and c is the dimensionality of node features. The edges describe the relations between nodes and can be represented as an adjacent matrix $\mathbf{A} = \{a_1, \dots, a_n\} \in \mathbb{R}^{n \times n}$, where a_i denotes the i -th row of the adjacency matrix, and $\mathbf{A}(i, j)$ denotes the relation between nodes v_i and v_j . Also, the node degrees are given by $\mathbf{d} = \{d_1, \dots, d_n\}$, where d_i computes the sum of edge weights connected to node v_i . Meanwhile, the degree of the whole graph is calculated by $\mathbf{d}(\mathbf{G}) = \sum_i^n d_i$. When we apply message-passing GNNs on \mathbf{G} , the message matrix can be represented as $\mathbf{M} = \{m_1, \dots, m_k\} \in \mathbb{R}^{k \times c'}$, where m_i is a message propagated between nodes, k is the total number of messages propagated on the graph, and c' is the dimension number of the messages.

Message-passing GNNs. Most of the existing GNN models adopt the message-passing framework, where each node sends messages to its neighbors and simultaneously receives messages from its neighbors. In the process of the propagation, node representations are updated based on node feature information and messages from neighbors, which can be formulated as

$$h_i^{(l+1)} = \gamma^{(l)}(h_i^{(l)}, \mathcal{AGG}_{j \in \mathcal{N}(i)}(\phi^{(l)}(h_i^{(l)}, h_j^{(l)}, e_{j,i}))) \quad (1)$$

where $h_i^{(l)}$ denotes the hidden representation of node v_i in the l -th layer, and $\mathcal{N}(i)$ is a set of nodes adjacent to node v_i ; $e_{j,i}$ represents the edge from node j to node i ; $\phi^{(l)}$ and $\gamma^{(l)}$ are differentiable functions; and \mathcal{AGG} represents the aggregation operation. From the perspective of the message-passing schema, we can gather all the propagated messages into a message matrix $\mathbf{M} \in \mathbb{R}^{k \times c'}$. Specifically, each row of the message matrix \mathbf{M} corresponds to a message propagated on a directed edge, which can be expressed as below:

$$\mathbf{M}_{(i,j)}^{(l)} = \phi^{(l)}(h_i^{(l)}, h_j^{(l)}, e_{j,i}^{(l)})$$

where ϕ denotes the mapping that generates the messages, c' is the dimension number of the messages, and the row number k of the message matrix \mathbf{M} is equal to the directed edge number in the graph.

Our Approach

In this section, we introduce our proposed *DropMessage*, which can be applied to all message-passing GNNs. We first

describe the details of our approach, and further prove that the most common existing random dropping methods, i.e., Dropout, DropEdge and DropNode, can be unified into our framework. Based on that, we give a theoretical explanation of the effectiveness of these methods. After that, we theoretically analyze the superiority of DropMessage in terms of stabilizing the training process and keeping information diversity. Finally, we derive a theoretical upper bound to guide the selection of dropping rate δ .

DropMessage

Algorithm description. Different from existing random dropping methods, DropMessage performs directly on the message matrix \mathbf{M} instead of the feature matrix or the adjacency matrix. More specifically, DropMessage conducts dropping on the message matrix with the dropping rate δ , which means that $\delta|\mathbf{M}|$ elements of the message matrix will be masked in expectation. Formally, this operation can be regarded as a sampling process. For each element $\mathbf{M}_{i,j}$ in the message matrix, we generate an independent mask $\epsilon_{i,j}$ to determine whether it will be preserved or not, according to a Bernoulli distribution $\epsilon_{i,j} \sim \text{Bernoulli}(1 - \delta)$. Then, we obtain the perturbed message matrix $\widetilde{\mathbf{M}}$ by multiplying each element with its mask. Finally, we scale $\widetilde{\mathbf{M}}$ with the factor of $\frac{1}{1-\delta}$ to guarantee that the perturbed message matrix is equal to the original message matrix in expectation. Thus, the whole process can be expressed as $\widetilde{\mathbf{M}}_{i,j} = \frac{1}{1-\delta} \epsilon_{i,j} \mathbf{M}_{i,j}$, where $\epsilon_{i,j} \sim \text{Bernoulli}(1 - \delta)$. The applied GNN model then propagates information via the perturbed message matrix $\widetilde{\mathbf{M}}$ instead of the original message matrix. It should be noted that DropMessage only affects on the training process.

Practical implementation. In practice, we do NOT need to generate the complete message matrix explicitly, because each row in the message matrix represents a distinct directed edge in the graph, and our proposed DropMessage can be applied to every directed edge independently. This property allows DropMessage to be easily parallelized, e.g., edge-wise, node-wise, or batch-wise, and to be applied to the message-passing backbone model *without increasing time or space complexity*.

Unifying random dropping methods. As we have mentioned above, DropMessage differs from existing methods by directly performing on messages instead of graphs. However, in intuition, the dropping of features, edges, nodes or messages will all eventually act on the message matrix. It inspires us to explore the theoretical connection between different dropping methods. As a start, we demonstrate that Dropout, DropEdge, DropNode, and DropMessage can all be formulated as Bernoulli sampling processes in Table 1. More importantly, we find that existing random dropping methods are actually special cases of DropMessage, and thus can be expressed in a uniform framework.

Lemma 1. *Dropout, DropEdge, DropNode, and DropMessage perform random masking on the message matrices in accordance with certain rules.*

We provide the equivalent operation on the message matrix

Method	Formula
Dropout	$\tilde{\mathbf{X}}_{i,j} = \epsilon \mathbf{X}_{i,j}$
DropEdge	$\tilde{\mathbf{A}}_{i,j} = \epsilon \mathbf{A}_{i,j}$
DropNode	$\tilde{\mathbf{X}}_i = \epsilon \mathbf{X}_i$
DropMessage	$\tilde{\mathbf{M}}_{i,j} = \epsilon \mathbf{M}_{i,j}$

s.t. $\epsilon \sim \text{Bernoulli}(1 - \delta)$

Table 1: Overview of different random dropping methods in a view of Bernoulli sampling process.

of each method below.

Dropout. Dropping the elements $X_{drop} = \{\mathbf{X}_{i,j} | \epsilon_{i,j} = 0\}$ in the feature matrix \mathbf{X} is equivalent to masking elements $M_{drop} = \{\mathbf{M}_{i,j} | \text{source}(\mathbf{M}_{i,j}) \in X_{drop}\}$ in the message matrix \mathbf{M} , where $\text{source}(\mathbf{M}_{i,j})$ indicates which element in the feature matrix that $\mathbf{M}_{i,j}$ corresponds to.

DropEdge. Dropping the elements $E_{drop} = \{\mathbf{E}_{i,j} | \mathbf{A}_{i,j} = 1 \text{ and } \epsilon_{i,j} = 0\}$ in the adjacency matrix \mathbf{A} is equivalent to masking elements $M_{drop} = \{\mathbf{M}_i | \text{edge}(\mathbf{M}_i) \in E_{drop}\}$ in the message matrix \mathbf{M} , where $\text{edge}(\mathbf{M}_i)$ indicates which edge that \mathbf{M}_i corresponds to.

DropNode. Dropping the elements $V_{drop} = \{\mathbf{X}_i | \epsilon_i = 0\}$ in the feature matrix \mathbf{X} is equivalent to masking elements $M_{drop} = \{\mathbf{M}_i | \text{node}(\mathbf{M}_i) \in V_{drop}\}$ in the message matrix \mathbf{M} , where $\text{node}(\mathbf{M}_i)$ indicates which row in the feature matrix that \mathbf{M}_i corresponds to.

DropMessage. This method directly performs random masking on the message matrix \mathbf{M} .

According to above descriptions, we find DropMessage conduct finest-grained masking on the message matrix, which makes it the most flexible dropping method, and other methods can be regarded as a special form of DropMessage.

Theoretical explanation of effectiveness. Previous studies have explored and explained why random dropping works in the filed of computer vision (Wager, Wang, and Liang 2013; Wan et al. 2013). However, to the best of our knowledge, the effectiveness of random dropping on GNNs has not been studied yet. To fill this gap, based on the unified framework of existing methods, we next provide a theoretical analysis.

Theorem 1. *Unbiased random dropping on GNNs methods introduce an additional regularization term into the objective functions, which makes the models more robust.*

Proof. For analytical simplicity, we assume that the downstream task is a binary classification and we apply a single layer GCN (Kipf and Welling 2017) as the backbone model, which can be formulated as $\mathbf{H} = \bar{\mathbf{B}}\mathbf{M}\mathbf{W}$, where \mathbf{M} denotes the message matrix, \mathbf{W} denotes the transformation matrix, $\mathbf{B} \in \mathbb{R}^{n \times k}$ indicates which messages should be aggregated by each node and $\bar{\mathbf{B}}$ is its normalized form. Also, we adopt sigmoid as non-linear function and present the result as $\mathbf{Z} = \text{sigmoid}(\mathbf{H})$. When we use cross-entropy as loss

function, the objective function can be expressed as follows:

$$L_{CE} = \sum_{j,y_j=1} \log(1 + e^{-h_j}) + \sum_{k,y_k=0} \log(1 + e^{h_k}) \quad (2)$$

When performing random dropping on graphs, we use the perturbed message matrix $\tilde{\mathbf{M}}$ instead of the original message matrix \mathbf{M} . Thus, the objective function in expectation can be expressed as follows:

$$E(\tilde{L}_{CE}) = L_{CE} + \sum_i \frac{1}{2} z_i (1 - z_i) \text{Var}(\tilde{h}_i) \quad (3)$$

More details of the derivation can be found in Appendix. As shown in Equation 3, random dropping methods on graphs introduce an extra regularization to the objective function. For binary classification tasks, this regularization enforces the classification probability approach to 0 or 1, thus a clearer judgment can be obtained. By reducing the variance of \tilde{h}_i , random dropping methods motivate the model to extract more essential high-level representations. Therefore, the robustness of the models is enhanced. It is noted that Equation 3 can be well generalized to multi-classification tasks by extending dimension of the model output. Formally, when dealing with the multi-classification task, the final objective function can be expressed as $E(\tilde{L}_{CE}) = L_{CE} + \sum_i \frac{1}{2} z_i^{c_i} (1 - z_i^{c_i}) \text{Var}(\tilde{h}_i^{c_i})$, where c_i is the label of node v_i , and the superscript indicates which dimension of the vector is selected. \square

Advantages of DropMessage

We give two additional analysis to demonstrate the advantages of DropMessage on two aspects: stabilizing the training process and keeping diverse information.

Reducing sample variance. All random dropping methods are challenged by the problem of unstable training process. As existing works suggest, it is caused by the random noises introduced into each training epoch. These noises then add the difficulty of parameter coverage and the instability of training process. Generally, *sample variance* can be used to measure the degree of stability. According to Table 1, the input of each training epoch can be regarded as a random sample of the whole graph, and the sample variance is calculated by the average difference of every two independent samples. Compared with other random dropping methods, DropMessage effectively alleviates the aforementioned problem by reducing the sample variance.

Theorem 2. *DropMessage presents the smallest sample variance among existing random dropping methods on message-passing GNNs with the same dropping rate δ .*

We leave the proof in Appendix. Intuitively, DropMessage independently determines whether an element in the message matrix is masked or not, which is exactly the smallest Bernoulli trail for random dropping on the message matrix. By reducing the sample variance, DropMessage diminishes the difference of message matrices among distinct training epochs, which stabilizes the training process and expedites the convergence. The reason why DropMessage has the minimum sample variance is that it is the finest-grained random

dropping method for GNN models. When applying DropMessage, each element $M_{i,j}$ will be independently judged that whether it should be masked.

Keeping diverse information. In the following, we compare different random dropping methods with their degree of losing *information diversity*, from the perspective of information theory.

Definition 1. *The information diversity consists of feature diversity and topology diversity. We define feature diversity as $FD_G = \text{card}(\{\|M_{SN(v_i),l}\|_0 \geq 1\})$, where $v_i \in \mathbf{V}$, $l \in [0, c)$, $SN(v_i)$ indicates the slice of the row numbers corresponding to the edges sourced from v_i ; topology diversity is defined as $TD_G = \text{card}(\{\|M_j\|_0 \geq 1\})$, where $j \in [0, k)$. $M \in \mathbb{R}^{k \times c}$ represents the message matrix, $\|\cdot\|_0$ calculates the zero norm of the input vector, and $\text{card}(\cdot)$ counts the number of elements in the set.*

In other words, feature diversity is defined as the total number of preserved feature dimensions from distinct source nodes; topology diversity is defined as the total number of directed edges propagating at least one dimension message. With the above definition, we claim that a method possesses the ability of *keeping information diversity* only under the condition where neither the feature diversity nor the topology diversity decreases after random dropping.

Lemma 2. *None of Dropout, DropEdge, and DropNode is able to keep information diversity.*

According to Definition 1, when we drop an element of the feature matrix \mathbf{X} , all corresponding elements in the message matrix are masked and the feature diversity is decreased by 1. When we drop an edge in adjacency matrix, the corresponding two rows for undirected graphs in the message matrix are masked and the topology diversity is decreased by 2. Similarly, when we drop a node, *i.e.*, a row in the feature matrix, elements in the corresponding rows of the message matrix are all masked. Both the feature diversity and the topology diversity are therefore decreased. Thus, for all of these methods, their feature and topology information cannot be completely recovered by propagated messages, leading to the loss of information diversity.

Theorem 3. *DropMessage can keep information diversity in expectation when $\delta_i \leq 1 - \min(\frac{1}{d_i}, \frac{1}{c})$, where δ_i is the dropping rate for node v_i , d_i is the out-degree of v_i , and c is the feature dimension.*

Proof. DropMessage conducts random dropping directly on message matrix \mathbf{M} . To keep the diversity of the topology information, we expect that at least one element of each row in message matrix \mathbf{M} can be preserved in expectation:

$$E(|\mathbf{M}_f|) \geq 1 \Rightarrow (1 - \delta)c \geq 1 \Rightarrow \delta \leq 1 - \frac{1}{c} \quad (4)$$

To keep the diversity of the feature information, we expect that for every element in the feature matrix \mathbf{X} , at least one of its corresponding elements in the message matrix \mathbf{M} is preserved in expectation:

$$E(|\mathbf{M}_e|) \geq 1 \Rightarrow (1 - \delta_i)d_i \geq 1 \Rightarrow \delta_i \leq 1 - \frac{1}{d_i} \quad (5)$$

Therefore, to keep the information diversity, the dropping rate δ_i should satisfy both Equation 4 and Equation 5 as

$$\delta_i \leq 1 - \min(\frac{1}{d_i}, \frac{1}{c}) \quad (6)$$

□

From the perspective of information theory, a random dropping method with the capability of keeping information diversity can preserve more information and theoretically perform better than those without such capability. Thus, it can explain why our method performs better than those existing dropping methods. Actually, we may only set one dropping rate δ for the whole graph rather than for each node in practice. Consequently, both DropMessage and other methods may lose some information. However, DropMessage still preserves more information than other methods with the same dropping rate even under this circumstance. It is demonstrated that DropMessage remains its advantage in real-world scenarios.

Experiments

Experimental Setup

We empirically validate the effectiveness and adaptability of our proposed DropMessage in this section. In particular, we explore the following questions: 1) Does DropMessage outperform other random dropping methods on GNNs? 2) Could DropMessage further improve the robustness and training efficiency of GNNs? 3) Does information diversity (described in Definition 1) matter in GNNs?

Datasets. We employ 7 graph datasets in our experiments, including 5 public datasets *Cora*, *CiteSeer*, *PubMed*, *ogbn-arxiv*, *Flickr* and 2 industrial datasets *FinV*, *Telecom*.

- *Cora*, *CiteSeer*, *PubMed*, *ogbn-arxiv*: These 4 different citation networks are widely used as graph benchmarks (Sen et al. 2008; Hu et al. 2020). We conduct node classification tasks on each dataset to determine the research area of papers/researchers. We also consider link prediction on the first three graphs to predict whether one paper cites another.
- *Flickr*: It is provided by Flickr, the largest photo-sharing website (Zeng et al. 2020). One node in the graph represents one image uploaded to Flickr. If two images share some common properties (*e.g.*, same geographic location, same gallery, or comments by the same user), an edge between the nodes of these two images will appear. We conduct the node classification task that aims to categorize these images into 7 classes determined by their tags.
- *FinV*, *Telecom*: These are two real-world mobile communication networks provided by FinVolution Group (Yang et al. 2019) and China Telecom (Yang et al. 2021), respectively. In the two datasets, nodes represent users, and edges indicate the situation where two users have communicated with each other at a certain frequency. The task is to identify whether a user is a default borrower or a telecom fraudster.

Baseline methods. We compare our proposed DropMessage with other existing random dropping methods, including

Model \ Task & Dataset	Node classification							Link prediction		
	Cora	CiteSeer	PubMed	ogbn-arxiv	Flickr	Telecom	FinV	Cora	CiteSeer	PubMed
GCN	80.68	70.83	78.97	70.08	0.5188	0.6080	0.4220	0.9198	0.8959	0.9712
GCN-Dropout	<u>83.16</u>	71.48	<u>79.13</u>	<u>71.16</u>	<u>0.5222</u>	0.6601	0.4526	0.9278	0.9107	<u>0.9766</u>
GCN-DropEdge	81.69	71.43	79.06	70.88	0.5214	<u>0.6650</u>	<u>0.4729</u>	<u>0.9295</u>	0.9067	0.9762
GCN-DropNode	83.04	72.12	79.00	70.98	0.5213	0.6243	0.4571	0.9238	0.9052	0.9748
GCN-DropMessage	83.33	<u>71.83</u>	79.20	71.27	0.5223	0.6710	0.4876	0.9305	<u>0.9071</u>	0.9772
GAT	81.35	70.14	77.20	70.32	0.4988	0.7050	0.4467	0.9118	0.8895	0.9464
GAT-Dropout	82.41	71.31	78.31	71.28	0.4998	0.7382	0.4539	0.9182	0.9055	0.9536
GAT-DropEdge	81.82	71.17	77.70	70.67	<u>0.5004</u>	<u>0.7568</u>	0.4896	0.9206	0.9037	0.9493
GAT-DropNode	82.08	<u>71.44</u>	77.98	70.96	0.4992	0.7214	0.4647	0.9224	0.9104	0.9566
GAT-DropMessage	<u>82.20</u>	71.48	<u>78.14</u>	<u>71.13</u>	0.5013	0.7574	<u>0.4861</u>	<u>0.9216</u>	<u>0.9076</u>	<u>0.9553</u>
APPNP	81.45	70.62	79.79	69.11	0.5047	0.6217	0.3952	0.9058	0.8844	0.9531
APPNP-Dropout	82.23	71.93	<u>79.92</u>	<u>69.36</u>	0.5055	0.6578	0.4023	0.9119	0.9071	0.9611
APPNP-DropEdge	82.75	<u>72.10</u>	79.83	69.15	<u>0.5061</u>	<u>0.6591</u>	0.4149	<u>0.9139</u>	<u>0.9131</u>	<u>0.9626</u>
APPNP-DropNode	81.79	71.50	79.81	69.27	0.5053	0.6412	<u>0.4182</u>	0.9068	0.8979	0.9561
APPNP-DropMessage	<u>82.37</u>	72.65	80.04	69.72	0.5072	0.6619	0.4378	0.9165	0.9141	0.9634

Table 2: Comparison results of different random dropping methods. The best results are in bold, while the second-best ones are underlined.

Dropout (Hinton et al. 2012), DropEdge (Rong et al. 2019), and DropNode (Feng et al. 2020). We adopt these dropping methods on various GNNs as the backbone model, and compare their performances on different datasets.

Backbone models. In this paper, we mainly consider three mainstream GNNs as our backbone models: GCN (Kipf and Welling 2017), GAT (Velickovic et al. 2018), and APPNP (Klicpera, Bojchevski, and Günnemann 2019). We take the official practice of these methods while make some minor modifications. All these backbone models have random dropping modules for different steps in their model implementation. For instance, GAT models perform random dropping after self-attention calculation, while APPNP models perform random dropping at the beginning of each iteration. For a fair comparison, we unify the implementation of random dropping modules in the same step for different backbone models. We fix Dropout, DropEdge, and DropNode on the initial input and fix DropMessage at the start point of the message propagation process.

Comparison Results

Table 2 summarizes the overall results. For the node classification task, the performance is measured by accuracy on four public datasets (Cora, CiteSeer, PubMed, ogbn-arxiv). As for Flickr and two imbalanced industrial datasets, we employ F1 scores. When it comes to the link prediction task, we calculate the AUC values for comparisons. Considering the space limitation, the std values of the experimental results are presented in the Appendix.

Effect of random dropping methods. It is observed that random dropping methods consistently outperform GNNs without random dropping in both node classification and link prediction. Besides, we see that the effects of random dropping methods vary over different datasets, backbone models, and downstream tasks. For example, random dropping methods on APPNP obtain an average accuracy improvement of

1.4% on CiteSeer, while 0.1% on PubMed. Meanwhile, random dropping methods achieve 2.1% accuracy improvement for GCN on Cora, while only 0.8% for GAT.

Comparison of different dropping methods. Our proposed DropMessage works well in all settings, exhibiting its strong adaptability to various scenarios. Overall, we have 21 settings under the node classification task, each of which is a combination of different backbone models and datasets (e.g., GCN-Cora). It is showed that DropMessage achieves the optimal results in 15 settings, and gets sub-optimal results in the rest. As to 9 settings under the link prediction task, DropMessage achieves the optimal results in 5 settings, and sub-optimal results in the rest. Moreover, the stable performance of DropMessage over all datasets compared to other methods is clearly presented. Taking DropEdge as the counterexample, it appears strong performance on industrial datasets but demonstrates a clear drop on public ones. A reasonable explanation is that the message matrix patterns reserved by distinct mask methods vary from each other as presented in Table 1. With the favor of its finest-grained dropping strategy, DropMessage obtains smaller inductive bias. Thus, compared with other methods, DropMessage is more applicable in most scenarios.

Additional Results

Robustness analysis. We study the robustness of dropping methods through measuring their ability of handling perturbed graphs. To guarantee that the initial data is comparatively clean, we conduct experiments on three citation networks: Cora, CiteSeer, and PubMed. We randomly add a certain ratio of edges into these datasets and perform the node classification. We find that all the random dropping methods have positive effects when the perturbation rate increases from 0% to 30%. The average improvement in the case of 30% perturbation reached 37% compared to that without perturbation, which indicates that the random dropping methods




Figure 2: Over-Smoothing and Training Process Analysis.

strengthen the robustness of GNN models. Besides, our proposed DropMessage shows its versatility and outperforms other dropping methods in noisy situations. Detailed results are exhibited in Appendix.

Over-smoothing analysis. Over-smoothing is a common issue on GNNs (Li, Han, and Wu 2018), which implies that the node representations become indistinguishable as the network depth increases. In this part, we evaluate the effects that various random dropping methods exert on this issue, and measure the degree of over-smoothing by MADGap (Chen et al. 2020). It should be noted that here a smaller value indicates the more indistinguishable node representations and vice versa. Experiments are conducted on Cora with GCNs serving as backbone models. Figure 2a and Figure 2b show the relative increase of MADGap values and test accuracies of the final node representations compared to the original model without any random dropping techniques. The results indicate that all these random dropping methods can alleviate over-smoothing by increasing the MADGap values and test accuracies when the depth of the model increases. Among all random dropping methods, our proposed DropMessage exhibits a superiority of consistency. It obtains an average improvement of 3.3% on MADGap values and an average improvement of 4.9% on test accuracies compared to other random dropping methods when the layer number $l \geq 3$. This result can be explained by the fact that DropMessage can generate more various messages than other methods, which prevents the nodes from converging to the same representations to some extent. A more detailed theoretical explanation can be found in Appendix.

Training process analysis. We conduct experiments to analyze the loss during the training process when employing different random dropping methods. Figure 2c shows the change of loss in GCN training processes when employing different random dropping methods on Cora. Furthermore, the similar training loss curves can be drawn under other experimental settings. The experimental results suggest that DropMessage presents the smallest sample variance among all methods, thus achieving the fastest convergence and the most stable performance. This is consistent with the theoretical results in Section .

Information diversity analysis. We conduct experiments to

Model	GCN		GAT		APPNP	
	AVG	NW	AVG	NW	AVG	NW
Accuracy (%)	81.62	82.67	80.81	81.61	80.71	81.56

Table 3: Classification accuracy (%) for information diversity analysis (where AVG denotes average, and NW denotes nodewise).

evaluate the importance of information diversity for message-passing GNN models. We set Cora as our experimental dataset, which contains 2708 nodes and 5429 edges. The average node degree of Cora is close to 4. According to Equation 6, the upper bound of dropping rate is calculated from the node degree and the feature dimension. The feature dimension number of Cora is 1433, which is much larger than the number of node degree. Therefore, the upper bound is only determined by the degree of the node. In (backbone)-nodewise settings, we set the dropping rate to be equal to its upper bound $\delta_i = 1 - \frac{1}{d_i}$ for each node. In (backbone)-average settings, we set the dropping rate $\delta_i = 0.75 + \epsilon_i$, where $\epsilon_i \sim Uniform(-0.15, 0.15)$. Both of these settings employ DropMessage. The average random dropping rate of all nodes is almost identical under these two settings, but only the former one can keep the information diversity in expectation. Table 3 presents the results. The (backbone)-nodewise settings outperform (backbone)-average settings regardless of which backbone model is selected.

Conclusion

In this paper, we propose DropMessage, a general random dropping method for message-passing GNN models. We first unify all random dropping methods to our framework via performing dropping on the message matrix and analyzing their effects. Then we illustrate the superiority of DropMessage theoretically in stabilizing the training process and keeping information diversity. Due to its fine-grained dropping operations on the message matrix, DropMessage shows greater applicability in most cases. By conducting experiments for multiple tasks on five public datasets and two industrial datasets, we demonstrate the effectiveness and generalization of our proposed method.

Acknowledgments

This work was partially supported by NSFC (62176233), the National Key Research and Development Project of China (2018AAA0101900), and the Fundamental Research Funds for the Central Universities.

References

- Abu-Mostafa, Y. S. 1990. Learning from hints in neural networks. *Journal of complexity*, 6(2): 192–198.
- Bishop, C. M. 1995. Training with Noise is Equivalent to Tikhonov Regularization. *Neural Computation*, 7: 108–116.
- Burges, C.; and Schölkopf, B. 1996. Improving the Accuracy and Speed of Support Vector Machines. In *NIPS*.
- Buslaev, A. V.; Parinov, A.; Khvedchenya, E.; Iglovikov, V. I.; and Kalinin, A. A. 2020. Alumentations: fast and flexible image augmentations. *ArXiv*, abs/1809.06839.
- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and Relieving the Over-smoothing Problem for Graph Neural Networks from the Topological View. In *AAAI*.
- Chen, J.; Ma, T.; and Xiao, C. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. *ICLR*.
- Cohen, J. M.; Rosenfeld, E.; and Kolter, J. Z. 2019. Certified Adversarial Robustness via Randomized Smoothing. In *ICML*.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NIPS*.
- Ding, K.; Xu, Z.; Tong, H.; and Liu, H. 2022. Data Augmentation for Deep Graph Learning: A Survey. *ArXiv*, abs/2202.08235.
- Ding, M.; Tang, J.; and Zhang, J. 2018. Semi-supervised learning on graphs with generative adversarial nets. In *CIKM*, 913–922.
- Feng, W.; Zhang, J.; Dong, Y.; Han, Y.; Luan, H.; Xu, Q.; Yang, Q.; Kharlamov, E.; and Tang, J. 2020. Graph Random Neural Networks for Semi-Supervised Learning on Graphs. *NeurIPS*, 33.
- Frid-Adar, M.; Diamant, I.; Klang, E.; Amitai, M. M.; Goldberger, J.; and Greenspan, H. 2018. GAN-based Synthetic Medical Image Augmentation for increased CNN Performance in Liver Lesion Classification. *Neurocomputing*, 321: 321–331.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *ICML*, 1263–1272.
- Hinton, G. E.; Srivastava, N.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *ArXiv*, abs/1207.0580.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *arXiv preprint arXiv:2005.00687*.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. In *ArXiv*, volume abs/1611.07308.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Klicpera, J.; Bojchevski, A.; and Günnemann, S. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *ICLR*.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning. In *AAAI*.
- Luo, D.; Cheng, W.; Yu, W.; Zong, B.; Ni, J.; Chen, H.; and Zhang, X. 2021. Learning to Drop: Robust Graph Neural Network via Topological Denoising. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*.
- Maaten, L. V. D.; Chen, M.; Tyree, S.; and Weinberger, K. Q. 2013. Learning with Marginalized Corrupted Features. In *ICML*.
- Matsuoka, K. 1992. Noise injection into inputs in back-propagation learning. *IEEE Trans. Syst. Man Cybern.*, 22: 436–440.
- Oono, K.; and Suzuki, T. 2019. On Asymptotic Behaviors of Graph CNNs from Dynamical Systems Perspective. *ArXiv*, abs/1905.10947.
- Oono, K.; and Suzuki, T. 2020. Graph Neural Networks Exponentially Lose Expressive Power for Node Classification. *arXiv: Learning*.
- Papp, P. A.; Martinkus, K.; Faber, L.; and Wattenhofer, R. 2021. DropGNN: Random Dropouts Increase the Expressiveness of Graph Neural Networks. In *NeurIPS*.
- Ramaswamy, L.; Gedik, B.; and Liu, L. 2005. A distributed approach to node clustering in decentralized peer-to-peer networks. In *TPDS*, volume 16, 814–829.
- Rifai, S.; Dauphin, Y.; Vincent, P.; Bengio, Y.; and Muller, X. 2011. The Manifold Tangent Classifier. In *NIPS*.
- Rong, Y.; Huang, W.; Xu, T.; and Huang, J. 2019. Dropedge: Towards deep graph convolutional networks on node classification. In *ICLR*.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. volume 29, 93–93.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6: 1–48.
- Simard, P. Y.; LeCun, Y. A.; Denker, J. S.; and Victorri, B. 1998. Transformation invariance in pattern recognition—tangent distance and tangent propagation. In *Neural networks: tricks of the trade*, 239–274. Springer.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Velickovic, P.; Fedus, W.; Hamilton, W. L.; Lio, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. *ICLR*.

- Wager, S.; Wang, S. I.; and Liang, P. 2013. Dropout Training as Adaptive Regularization. In *NIPS*.
- Wan, L.; Zeiler, M. D.; Zhang, S.; LeCun, Y.; and Fergus, R. 2013. Regularization of Neural Networks using DropConnect. In *ICML*.
- Xu, K.; Li, C.; Tian, Y.; Sonobe, T.; Kawarabayashi, K.; and Jegelka, S. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In *ICML*.
- Yang, Y.; Xu, Y.; Sun, Y.; Dong, Y.; Wu, F.; and Zhuang, Y. 2021. Mining Fraudsters and Fraudulent Strategies in Large-Scale Mobile Social Networks. *IEEE Transactions on Knowledge and Data Engineering*, 33: 169–179.
- Yang, Y.; Xu, Y.; Wang, C.; Sun, Y.; Wu, F.; Zhuang, Y.; and Gu, M. 2019. Understanding Default Behavior in On-line Lending. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*.
- Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *SIGKDD*, 974–983.
- Zeng, H.; Zhou, H.; Srivastava, A.; Kannan, R.; and Prasanna, V. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. *ArXiv*, abs/1907.04931.
- Zhao, L.; and Akoglu, L. 2020. PairNorm: Tackling Over-smoothing in GNNs. *ArXiv*, abs/1909.12223.
- Zhu, D.; Zhang, Z.; Cui, P.; and Zhu, W. 2019. Robust graph convolutional networks against adversarial attacks. In *SIGKDD*, 1399–1407.
- Zügner, D.; Akbarnejad, A.; and Günnemann, S. 2018. Adversarial Attacks on Neural Networks for Graph Data. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.