GAMC: An Unsupervised Method for Fake News Detection Using Graph Autoencoder with Masking

Shu Yin[1, 2], **Peican Zhu**[2*], **Lianwei Wu**[1], **Chao Gao**[2], **Zhen Wang**[3, 1, 2†]

[1]School of Computer Science, Northwestern Polytechnical University
[2]School of Artificial Intelligence, OPtics and ElectroNics (iOPEN), Northwestern Polytechnical University
[3]School of Cybersecurity, Northwestern Polytechnical University

Abstract

With the rise of social media, the spread of fake news has become a significant concern, potentially misleading public perceptions and impacting social stability. Although deep learning methods like CNNs, RNNs, and Transformer-based models like BERT have enhanced fake news detection. However, they primarily focus on content and do not consider social context during news propagation. Graph-based techniques have incorporated the social context but are limited by the need for large labeled datasets. To address these challenges, this paper introduces GAMC, an unsupervised fake news detection technique using the Graph Autoencoder with Masking and Contrastive learning. By leveraging both the context and content of news propagation as self-supervised signals, our method reduces the dependency on labeled datasets. Specifically, GAMC begins by applying data augmentation to the original news propagation graphs. Subsequently, these augmented graphs are encoded using a graph encoder and subsequently reconstructed via a graph decoder. Finally, a composite loss function that encompasses both reconstruction error and contrastive loss is designed. Firstly, it ensures the model can effectively capture the latent features, based on minimizing the discrepancy between reconstructed and original graph representations. Secondly, it aligns the representations of augmented graphs that originate from the same source. Experiments on the real-world dataset validate the effectiveness of our method.

Introduction

The rapid development of social media has brought convenience to people's lives (Hou et al. 2023). However, it has also served as a breeding ground for the widespread dissemination of fake news (Wang et al. 2023). The proliferation of fake news has become a major issue in the digital media era, as it could mislead public perception, affect social stability, and even threaten national political security (Zhou and Zafarani 2020). Therefore, the detection of fake news has become a pressing issue that requires effective solutions (Hua et al. 2023; Hou and Yin 2024).

To automatically identify the ever-growing fake news, various approaches have been proposed. Traditional fake news detection methods primarily involve manually designing rules to extract news features and then employing classifiers to categorize these features. However, manually designing rules to extract valuable fake news features could be labor-intensive, and may not always capture complex deceptive features (Zhou and Zafarani 2020).

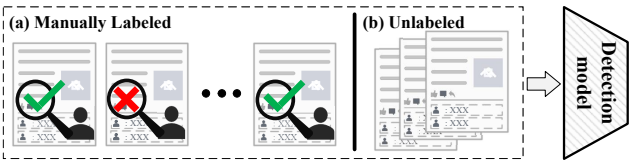

Figure 1: Difference between supervised and unsupervised methods for fake news detection. (a) Existing methods largely rely on manually labeled datasets. The process of manually annotating data is time-consuming, expensive, and often requires expert knowledge to ensure accurate labeling. (b) In contrast, our proposed method, GAMC, is based on unsupervised learning and can directly use unlabeled data for fake news detection. This eliminates the need for costly and time-consuming manual data annotation.

In recent years, deep learning has played an increasingly important role in fake news detection, as it can automatically learn and extract features from large amounts of data, thereby improving the accuracy and efficiency of detection mechanisms (Shu et al. 2017). For example, Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) are utilized to learn local and temporal dependencies in text data respectively (Li et al. 2021b). Furthermore, Transformer-based models, such as BERT, have been employed to understand the context and semantic relationships in news articles better (Devlin et al. 2019). These models, pre-trained on large corpora, have shown remarkable success in capturing the complex linguistic characteristics of fake news (Wu et al. 2023). However, these methods primarily analyze the content of the news, and do not consider the social context information in the process of news propagation (Yuan et al. 2019).

Recognizing this oversight, researchers have proposed graph-based methods that incorporate social context into the detection process (Min et al. 2022). These methods model the spread of news as a graph, capturing the intricate in-

*Corresponding author: ericcan@nwpu.edu.cn
†Corresponding author: w-zhen@nwpu.edu.cn

teractions and relationships among various entities involved in news propagation. However, these supervised methods necessitate large labeled datasets for training as Figure 1 (Fang et al. 2023). The collection and labeling of extensive datasets can be a laborious and resource-intensive task, posing a significant challenge for real-world applications (He et al. 2021).

To address these issues, this paper proposes an unsupervised fake news detection method, GAMC, that employs a Graph Autoencoder with Masking and Contrastive learning. By employing the context and content of the news propagation process as the self-supervised signal, along with a feature reconstruction and contrasting task, this method circumvents the need for labeled datasets. In addition, the application of data augmentations and contrastive ideas can improve the robustness of the model. Specifically, first, data augmentation is employed on the original news propagation graph, which includes random node feature masking and edge dropping, as facilitates the subsequent proxy tasks of feature reconstruction and contrastive. Then, the graph encoder is employed to encode enhanced graphs, yielding the latent representation vector. The latent vectors not only capture the global characteristics of the graph but also contain information about the news propagation process, both its context and content. Once the model has been trained, these latent representation vectors can be used for the task of fake news detection. Next, the graph decoder is used to map the latent representation vector back to the original input, resulting in reconstructed graph vectors. This step is designed to teach the model how to reconstruct the original input from the latent representation vector, thus helping the model to better learn the latent patterns of news propagation. Finally, a composite loss function is designed, composed of reconstruction error loss and contrast loss. On the one hand, the reconstruction error loss aims to minimize the discrepancy between the reconstructed graph representation and the original graph representation, enabling the graph autoencoder to better learn the latent features of the propagation graph. The contrast loss, on the other hand, ensures that the representations of the two augmented graphs generated from the same propagation graph are as similar as possible after reconstruction. Experiments on real datasets show that our method GAMC outperforms existing unsupervised methods. The code is available at https://github.com/cgao-comp/GAMC.

The contribution of this paper can be summarized as follows:

- Self-supervised learning is introduced into the domain of fake news detection, eliminating the dependence on labeled data, which makes the method more applicable to real-world scenarios.
- We proposed a graph autoencoder with reconstruction error loss and contrast loss. The reconstruction error loss aims to minimize the discrepancy between the reconstructed and the original graph representations, while the contrast loss ensures that the representations of two augmented graphs, both derived from the same propagation graph, are as similar as possible after reconstruction.
- We conducted a series of experiments on real-world datasets, demonstrating the effectiveness of the proposed method.

Related Work

To provide a comprehensive understanding of the current landscape of our work, we review two primary areas: fake news detection methods and generative self-supervised graph learning techniques.

Fake News Detection

The task of fake news detection can be viewed as a classification problem. The classification process relies on various factors such as the content of the news, the spread pattern, user reactions, and other related data.

Recently, deep learning has been taking an increasingly prominent role in fake news detection. Ma et al. developed a novel recurrent neural network (RNN) based method for rumor detection on microblogging platforms, which outperformed traditional models using hand-crafted features (Ma et al. 2016). Considering the different events, EANN is proposed to effectively extract event-invariant features from multimedia content, thereby enhancing the detection of fake news on newly arrived events (Wang et al. 2018). To introduce extra knowledge for detecting fake news, Wang et al. proposed a unified framework named KMGCN, using a graph convolutional network to extract textual information, knowledge concepts, and visual information (Wang et al. 2020). However, these methods primarily focus on the content of news, which may fall short in distinguishing ambiguous fake news that is crafted to resemble real news.

Building on this, researchers begin to explore the potential of leveraging social context information in the process of news propagation (Yang et al. 2022). Generally, the social context includes information such as forwarding relationships, comment content, and user preferences, which can provide additional insight into how news spreads in social networks. Bian et al. introduced the bi-directional graph convolutional network that simultaneously captures top-down propagation and bottom-up dispersion features on social media, which enhances traditional deep learning approaches (Bian et al. 2020). To capture rich structural information, GLAN models relationships among source tweets, retweets, and users as a heterogeneous graph, then effectively encodes both local semantic and global structural information for rumor detection (Yuan et al. 2019). Considering the influence of user preferences in news propagation, UPFD employs users' historical posts as an endogenous preference, and the news propagation graph as an exogenous context, integrating internal and external information to better identify disinformation (Dou et al. 2021).

However, these supervised methods depend on large labeled datasets. The acquisition of these labeled datasets often requires considerable time, effort, and domain expertise.

Generative Self-Supervised Graph Learning

Generative self-supervised graph learning leverages the richly structured data in graphs to learn meaningful representations without the need for explicit labels (Wu et al.

2021). To generate diverse and realistic graphs, Li et al. introduced GraphRNN, a deep autoregressive model that addresses the challenges of graph generation and representation learning (You et al. 2018). Kipf et al. developed a novel Graph Auto-Encoder (GAE) based method that learns to encode a graph into a lower-dimensional space and then decode it back into its original structure, outperforming traditional models using hand-crafted features (Kipf and Welling 2016). Considering that most GAEs lack the ability to reconstruct node features, some work has been dedicated to reconstructing masked features, thereby enhancing the efficiency of self-supervised GAEs in graph representation learning for classification tasks (Hou et al. 2022). Our work is inspired by the graph mask autoencoder, and we developed a self-supervised graph autoencoder to obtain representations of news for the task of fake news detection.

Problem Definition

In this paper, the task of fake news detection is to design an automatic discriminator that can learn latent features from a set of unlabeled news. The learned features can then be used to predict the reality of news instances. Specifically, the news dataset can be defined as $D = \{D_1, D_2, \ldots, D_n\}$, where each D_i represents a single news instance in the dataset. Each news instance D_i can be modeled as a graph based on its corresponding propagation process. We aim to learn an unsupervised function, f, as defined below:

$$f : D \rightarrow Y, \tag{1}$$

where D represents the set of news instances with propagation process and $Y \in \{F, R\}$ (i.e. fake news or real news) denotes the set of possible outcomes.

Method

In this section, we introduce the GAMC method for fake news detection tasks, designed to capitalize on the inherent context and content of the news propagation process to function as a self-supervised signal, thereby bypassing the need for labeled datasets. As illustrated in Figure 2, the following parts will detail the procedure of employing GAMC for fake news detection, including data augmentation, graph encoding, graph decoding, and the composite loss function.

Data Augmentation

For the GAMC method, the first step is to model news propagation as graphs. After that, to serve the feature reconstruction task and contrastive task, the data augmentation method is employed to generate different enhanced graphs by transforming the original graph for each piece of news.

To begin with, each piece of news is modeled into a graph $G = (V, A, X)$, based on the forwarding relationship. $V = \{v_n, v_u\}$ represents the set of nodes, where v_n is the news node, and v_u signifies the user nodes that forward the news. A represents the adjacency matrix, which embodies the forwarding relationships. X is the feature matrix. The node feature for the news nodes v_n is the news content embedding, encoded by a pre-trained BERT model, and the node feature for the user nodes v_u is derived from their historical posts, as described in (Dou et al. 2021). The node feature of node i is denoted as x_i, and the feature matrix X is composed by x_i.

Following the construction of the graph, the data augmentation process employs two strategies: random node feature masking and edge dropping. Random node feature masking is the random feature elimination of the nodes when the training begins. Let $V_m \subset V$ be a subset of nodes randomly selected for masking. For each node in V_m, we replace its feature vector with a special mask token, denoted as $x_{[MASK]} \in \mathbb{R}^D$. Then, the masked feature matrix $\widehat{X}$ can be defined as:

$$\widehat{x}_i = \begin{cases} x_{[MASK]}, & \text{if } v_i \in V_m \\ x_i, & \text{if } v_i \notin V_m, \end{cases} \tag{2}$$

where $\widehat{x}_i$ is the augmentation feature of node i, the augmented feature matrix $\widehat{X}$ is constructed by $\widehat{x}_i$.

Edge dropping is the second strategy used for data augmentation. This method disrupts the connectivity of the graph by randomly dropping some edges before training. E_{Drop} is the edge set obtained by randomly sampling from the original edge set E, and A_{Drop} denotes the adjacency matrix of E_{Drop}. Then the augmented adjacency matrix $\widehat{A}$ could be calculated as $\widehat{A} = A - A_{Drop}$. Through the above node feature masking and edge dropping operations, the augmented graph transformed by G can be represented as $\widehat{G} = (V, \widehat{A}, \widehat{X})$. Each of these strategies is applied twice to the original graph, resulting in two distinct augmented graphs, namely $\widehat{G}_1 = (V, \widehat{A}_1, \widehat{X}_1)$ and $\widehat{G}_2 = (V, \widehat{A}_2, \widehat{X}_2)$.

The data augmentation strategy not only ensures the model's effectiveness in the face of complex news propagation patterns but also facilitates the model to reconstruct missing features, thereby promoting the model to capture deeper feature dependencies.

Graph Encoding

Following the data augmentation process, the second integral component of the GAMC method is graph encoding. The purpose of this step is to transform the augmented graphs into a compact and meaningful latent space representation.

The graph encoder in our method is a two-layer Graph Isomorphism Network (GIN) (Xu et al. 2019), which is designed to process the augmented graph $\widehat{G} = (V, \widehat{A}, \widehat{X})$ and generate the graph-level representation vectors. The GIN is selected due to its capacity to capture the topological structure and node features of a graph, making it suitable for the task of fake news detection.

Given an augmented graph $\widehat{G}$, the GIN encoder operates as follows. At the l-th layer, the hidden feature vector $h_i^{(l)}$ for node i is updated using the aggregation function:

$$h_i^{(l)} = MLP\left((1 + \epsilon^{(l)}) \cdot h_i^{(l-1)} + \sum_{j \in \mathcal{N}(i)} \left(h_j^{(l-1)}\right)\right), \tag{3}$$

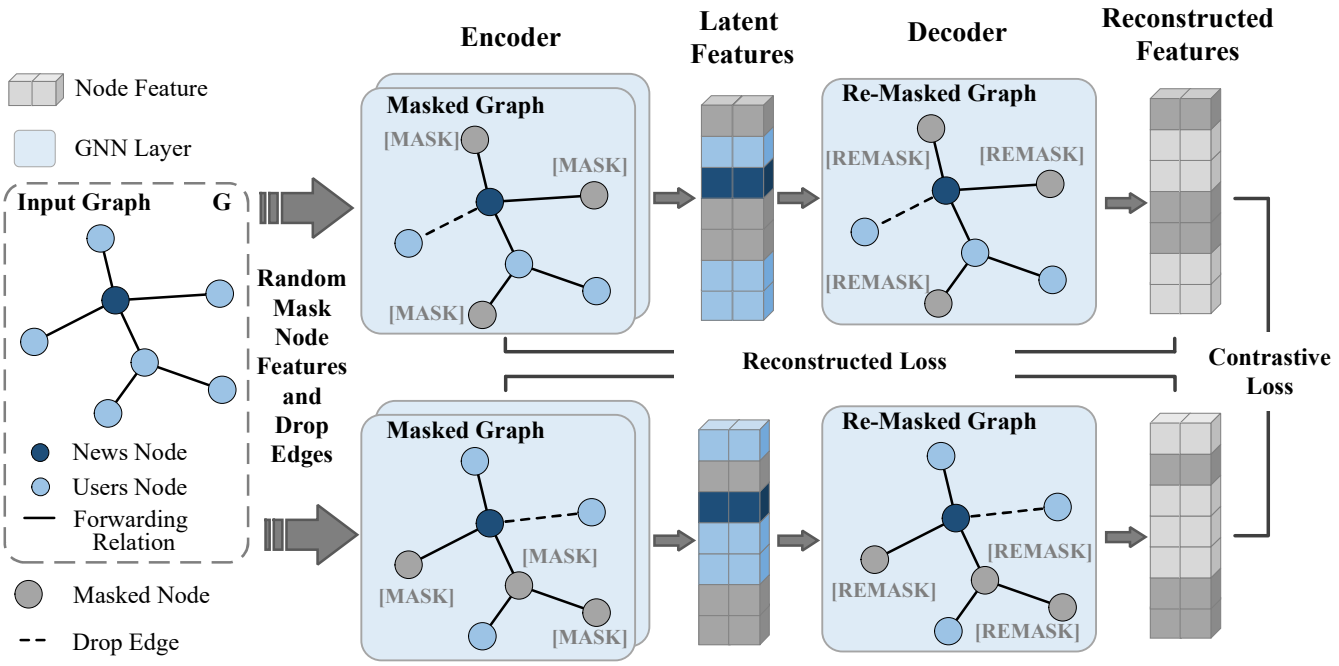

Figure 2: Overview of the GAMC Fake News Detection Method. The original news propagation graph is first augmented through random feature masking and edge dropping, generating two perturbed graphs. These graphs are then processed through a graph encoder to generate latent representation vectors. A decoder subsequently maps these vectors back to their input space, producing reconstructed feature vectors. Minimizing the composite loss function facilitates the effective learning of the propagation graph's latent features by the autoencoder and ensures the similarity of the reconstructed features across the augmented graphs.

where $\mathcal{N}(i)$ is the set of neighboring nodes of i, and $h_i^{(0)} = \widehat{x}_i$ is the input feature vector of node i. This process is iteratively conducted for all nodes until the l-th layer.

After two layers of information propagation, the GIN encoder outputs a set of node embeddings H for all nodes in the graph. Through the above encoding step, for the two augmented graphs $\widehat{G}_1$ and $\widehat{G}_2$, we can obtain the corresponding nodes' latent representations as:

$$H_1 = GIN_{Eecoder}(\widehat{A}_1, \widehat{X}_1), H_2 = GIN_{Eecoder}(\widehat{A}_2, \widehat{X}_2), \tag{4}$$

where node representations H_1 and H_2 will then be pushed into the graph decoder.

While the model has been trained, the latent representation vectors obtained from the GIN encoder can be used directly for news classification tasks. Node embeddings from the encoder are pooled together to generate a graph-level representation vector F for the entire graph:

$$F = \sum_{i=1}^{n} h_i. \tag{5}$$

This graph-level representation vector F captures the overall information of the graph, including both the structural and content information, which are essential for the downstream task of fake news detection.

Graph Decoding

The third core component of the GAMC method is graph decoding. The goal of this step is to map the latent graph-level representation vectors back to the input, namely to obtain reconstructed feature matrices. Before decoding, we perform a re-mask operation on the masked nodes. The re-mask option forces the masked nodes to aggregate from their neighbors to reconstruct their initial features. For each node in V_m, we replace its latent representation with a special mask token, denoted as $h_{[REMASK]} \in \mathbb{R}^D$. The re-mask representation $\widehat{h}_i$ of v_i can be described as:

$$\widehat{h}_i = \begin{cases} h_{[MASK]}, & \text{if } h_i \in V_m \\ h_i, & \text{if } h_i \notin V_m \end{cases}. \tag{6}$$

The re-masked latent representation $\widehat{H}$ of the whole graph is constructed by $\widehat{h}_i$. Correspondingly, after the re-mask operation, the hidden representations of graphs $\widehat{G}_1$ and $\widehat{G}_2$ can be denoted as $\widehat{H}_1$ and $\widehat{H}_2$. Then, we feed the hidden representations of the two graphs into the graph decoder, obtaining the reconstructed features X_1' and X_2'. The detailed aggregation process of decoding is shown in Eq. (3), and we can simplify the process as:

$$X_1' = GIN_{Decoder}(\widehat{A}_1, \widehat{H}_1), X_2' = GIN_{Decoder}(\widehat{A}_2, \widehat{H}_2). \tag{7}$$

Loss Function

The loss function in GAMC is to guide the learning process in a way that the difference between the original and reconstructed graphs is minimized, and the contrast between the two reconstructed graphs derived from the same propagation graph is minimized. We define the loss function in two parts: the reconstruction loss and the contrastive loss.

The reconstruction loss aims to ensure the fidelity of the reconstructed feature matrices X_1' and X_2' to the original feature matrices X_1 and X_2. We use the Mean Squared Error (MSE) between the original and reconstructed feature matrices as reconstruction loss:

$$\mathcal{L}rec = \frac{1}{n}\sum_{i=1}^{n}\left(|X_1 - X_1'|_2^2 + |X_2 - X_2'|_2^2\right), \quad (8)$$

where n is the number of samples. By minimizing this loss in the training phase, the model could produce reconstructed graphs that closely match the originals, which encourages the graph encoder to learn better graph-level representation vectors.

On the other hand, the contrastive loss is designed to minimize the difference between the two reconstructed graphs derived from the same propagation graph. This is achieved by minimizing the cosine similarity between reconstructed features X_1' and X_2':

$$\mathcal{L}con = \frac{X_1' \cdot X_2'}{\|X_1'\|\|X_2'\|}. \quad (9)$$

By minimizing this loss, the model is encouraged to generate similar representations for one augmented graph.

The overall loss function is then a weighted sum of the reconstruction loss and the contrastive loss:

$$\mathcal{L} = \mathcal{L}rec - \alpha\mathcal{L}con, \quad (10)$$

where α is the hyperparameter that controls the balance between the two loss components.

By minimizing this overall loss, our model is trained to generate robust and discriminative graph-level representations that can be effectively used for the task of fake news detection.

Experiments

In this section, we validate the effectiveness of the proposed GAMC method by comparing it with some benchmark models on public datasets. Following this, to analyze and validate the necessity of each component in GAMC, we conduct ablation studies. Finally, we investigate the impact of different parameter values within GAMC on the experimental results.

Datasets and Settings

Datasets To validate the efficiency of GAMC, we carried out evaluations on the FakeNewsNet, a published data source for fake news detection (Shu et al. 2020). This repository is divided into two sub-datasets, PolitiFact and GossipCop. The PolitiFact dataset primarily consists of news related to U.S. politics, while GossipCop is primarily focused on news about Hollywood celebrities. The social context in these two datasets includes the propagation network of news and the history of user posts. Table 1 provides comprehensive statistics of the PolitiFact and GossipCop datasets.

Dataset	PolitiFact	GossipCop
#News	314	5464
#True News	157	2732
#Fake News	157	2732
#Nodes	41054	314262
#Edges	40740	308798

Table 1: Statistics of the datasets. In the two datasets, each graph denotes a piece of news.

Baselines We have conducted a comparison of the proposed method GAMC with the following unsupervised methods:

- **TruthFinder** (Yin, Han, and Yu 2008) is one of the earliest methods for detecting fake news using an unsupervised approach. This method employs an iterative process to determine the veracity of news by assessing the credibility of the source websites of the news.
- **UFD** (Yang et al. 2019) employs a Bayesian network model and a collapsed Gibbs sampling technique. This method leverages users' engagements on social media to understand their opinions regarding news authenticity.
- **GTUT** (Gangireddy et al. 2020) is a graph-based method for fake news detection that identifies a seed set of articles, and then progressively labels all articles in the dataset.
- **UFNDA** (Li et al. 2021a) is an unsupervised fake news detection approach. Utilizing a combination of a Bidirectional GRU (Bi-GRU) layer and self-attention within an autoencoder, the method uncovers hidden relationships between features to detect fake news.
- $(\textbf{UMD})^2$ (Silva et al. 2023) is an unsupervised fake news detection framework that encodes multi-modal knowledge into low-dimensional vectors. This method leverages a teacher-student architecture to determine the truthfulness of news by aligning various modalities, then uses them as guiding signals for veracity assessment.

Additionally, we also conducted comparisons with the following classical supervised methods:

- **SAFE** (Zhou, Wu, and Zafarani 2020) is a multimodal method for fake news detection. It converts images in the news into text, learns the latent representation of text and visual information, then measures the similarity between them to detect fake news.
- **EANN** (Wang et al. 2018) is a multi-modal approach for detecting fake news. It extracts text and image features from news content, and then incorporates an event discriminator using adversarial learning to obtain the event-invariant features of fake news.
- **dEFEND** (Shu et al. 2019) designs a co-attention network to capture the noticeable sentences between news

and comments, which improves the explanatoryness of fake news detection.

- **GACL** (Sun et al. 2022) leverages contrastive learning within the loss function to learn the difference between positive and negative samples. This approach enhances the model's robustness, contributing to efficient fake news detection.
- **FinerFact** (Jin et al. 2022) models the article keywords and related evidence, and then uses a dual-channel kernel graph network to perform fine-grained reasoning on the news, improving the accuracy of fake news detection.

Parameter Settings The experiments were conducted on a server equipped with an Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz and a GeForce RTX 3090Ti graphics card. We implement the proposed GAMC model using PyTorch. We use accuracy, F1 score, precision, and recall as our evaluation metrics across both datasets. During the data augmentation process, we mask 50% of the node features and apply dropout to 20% of the edges. The node features are represented in 768-dimensional space, while the intermediate layer vectors produced by the encoder have a dimensionality of 512. The training procedure consists of 80 epochs, with the Adam optimization algorithm employed to optimize the model. The hyperparameter that controls the balance between the two loss components is set to 0.1. Upon completion of training, we classify the graph-level vectors learned by the graph encoder to predict labels. By conducting experiments using multiple classifiers, we selected Multilayer Perceptron (MLP) as the final classifier. For the experimental results, we run ten times and take the average values.

Overall Performance

Table 2 and Table 3 respectively display the performance of the proposed GAMC method and the unsupervised methods. From the results, compared to existing unsupervised methods, GAMC demonstrates noticeable improvements across all four metrics on the two datasets. Specifically, the accuracy on the PolitiFact dataset increased by 4.49%, and on the Gossipcop dataset, it rose by 19.44%. This improvement can be attributed to GAMC's unique design, which leverages a graph autoencoder with masking and contrast. The approach harnesses both the context and content of news propagation, thereby providing a more holistic and accurate representation. Additionally, the composite loss function, combining reconstruction error loss and contrast loss, ensures not only that the latent features of the propagation graph are accurately captured but also that the representations of the augmented graphs are closely aligned.

Table 4 shows the performance of the proposed GAMC and the supervised methods. Compared with Tables 2 and 3, we can observe that supervised methods tend to outperform unsupervised methods. This is primarily attributable to supervised methods taking advantage of the specific label information provided in training datasets, enabling these models to learn more distinctive and discriminative patterns associated with fake news.

As can be seen in Table 4, on the PolitiFact and Gossipcop datasets, the methods based on news propagation context (deFEND, GACL, and FinerFact) perform better than those based on news content (SAFE and EANN). This is due to social context-based methods effectively capturing the complex interconnections and behavioral patterns involved in news propagation. In contrast, methods exclusively focused on news content could potentially overlook these significant contextual signals.

Methods	**ACC.**	**Prec.**	**Rec.**	**F1.**
TruthFinder	0.581	0.572	0.576	0.573
UFNDA	0.685	0.667	0.659	0.670
UFD	0.697	0.652	0.641	0.647
GTUT	0.776	0.782	0.758	0.767
$(\text{UMD})^2$	0.802	0.795	0.748	0.761
GAMC	**0.838**	**0.836**	**0.827**	**0.831**
variance	*±0.013*	*±0.008*	*±0.011*	*±0.010*

Table 2: Results of GAMC, compared with unsupervised methods on the PolitiFact dataset. We bold the top results.

Methods	**ACC.**	**Prec.**	**Rec.**	**F1.**
TruthFinder	0.668	0.669	0.672	0.669
UFNDA	0.692	0.687	0.662	0.673
UFD	0.662	0.687	0.654	0.667
GTUT	0.771	0.770	0.731	0.744
$(\text{UMD})^2$	0.792	0.779	0.788	0.783
GAMC	**0.946**	**0.941**	**0.946**	**0.943**
variance	*±0.004*	*±0.005*	*±0.003*	*±0.004*

Table 3: Results of GAMC, compared with unsupervised methods on the GossipCop dataset.

Additionally, our unsupervised method GAMC demonstrates superior performance over the classic content-based supervised algorithms, achieving an accuracy improvement of 4.22% and 13.16%, respectively. Compared to classic social context-based supervised algorithms, our method shows a decrease in accuracy by 7.81% on PolitiFact and an increase by 4.29% on Gossipcop. However, in real-world scenarios where labeled data may be scarce or costly to obtain, our GAMC method offers an effective alternative. Furthermore, this unsupervised model opens up new possibilities for continual, on-the-fly fake news detection as it can easily adapt to changing data landscapes. As such, GAMC not only competes with supervised methods but also provides additional flexibility and cost-effectiveness, making it a robust solution for the challenge of fake news detection.

Ablation Study

To further elucidate the importance of each component in our proposed GAMC model, we conduct an ablation study in this section. This analysis aims to evaluate the contribution of individual modules by iteratively removing them and

Dataset	PolitiFact		GossipCop	
	Acc.	F1.	Acc.	F1.
SAFE	0.793	0.775	0.832	0.811
EANN	0.804	0.798	0.836	0.813
dEFEND	0.904	**0.928**	0.808	0.755
GACL	0.867	0.866	0.907	0.905
FinerFact	**0.909**	0.917	0.832	0.868
GAMC	0.838	0.831	**0.946**	**0.943**

Table 4: Results of GAMC, compared with supervised methods on PolitiFact and GossipCop datasets.

observing the effect on the model's performance. We compare GAMC with its various sub-models:

- GMAC-Aug removes the data augmentation, including node feature masking and edge dropping.
- GAMC-L_{rec} removes the reconstruction loss, and only depends on the contrastive loss to optimize the model.
- GAMC-L_{con} removes the contrastive loss while generating only one augmented graph from the propagation graph.

The comparative results of these various sub-models are visually summarized in Table 5. From the results, it can be observed that removing data augmentation (GAMC-Aug) led to a decrease in accuracy, indicating the importance of this feature in capturing the intricacies of news propagation. Data augmentation in GAMC is instrumental in increasing the autoencoder's feature reconstruction ability. The absence of reconstruction loss (GAMC-L_{rec}) made a noticeable difference in performance, weakening the model's ability to accurately regenerate the original graph structure. The contrastive loss helps the model to recognize similarities and differences between different instances, enhancing its discrimination power. By generating only one augmented graph (GAMC-L_{con}), the model loses the capability to contrast between various augmented views of the data. In conclusion, each component of the GAMC model plays a critical role in ensuring optimal performance.

Dataset	PolitiFact		GossipCop	
	Acc.	F1.	Acc.	F1.
GAMC-Aug	0.804	0.801	0.907	0.893
GAMC-L_{rec}	0.771	0.767	0.891	0.878
GAMC-L_{con}	0.816	0.812	0.926	0.917
GAMC	**0.838**	**0.836**	**0.946**	**0.943**

Table 5: Results of sub-models of GAMC on PolitiFact and GossipCop datasets.

Parameter Discussion

To ensure that our proposed GAMC model achieves optimal performance, an investigation and analysis of the parameters was conducted in this section. The mask rate λ and edge drop rate γ stand out as vital tunable parameters, impacting the model's capacity to understand and process the underlying data structure. For comprehensive insights into their influence, we conduct a series of experiments using both rates varying between 0.1 to 0.9 to encapsulate their entire effective range. Figure 3 shows that when the mask rate is

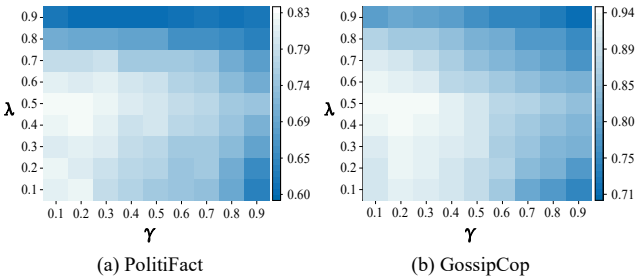

(a) PolitiFact (b) GossipCop

Figure 3: Parameter analysis for λ and γ in GAMC. (a) shows the results on PolitiFact, and (b) shows the results on Gossipcop.

set to 0.5 and the edge drop rate set to 0.2 the result is the best. At a high mask rate, a substantial portion of node information becomes occluded. This leads to the model losing critical information, making it harder to discern patterns and structures essential for its tasks. Samely, a high edge dropping disrupts the inherent structure and connectivity of the original graph. It makes the graph too sparse, thereby losing significant relational data between nodes. On the other hand, the low mask and edge drop rate might not provide enough reconstruction clues.

Conclusion

In this study, we introduced GAMC, a novel unsupervised approach to fake news detection. By executing data augmentations like node feature masking and edge dropping, we engender enhanced graphs. Subsequently, we implemented a graph encoding and decoding strategy. Furthermore, the composite loss function, including both the reconstruction error loss and the contrast loss, optimally synergizes these components. The reconstruction error loss facilitates the reconstruction of the original graph from its representation vector. The contrast loss facilitates the aligning representations of augmented graphs from the same original graph. Experiments validate that our method manifests effectively in fake news detection, eliminating the need for extensive labeled datasets. However, the method requires a certain level of propagation to detect fake news. In the future, we will integrate knowledge bases to facilitate efficient detection at the initial stages of fake news emergence.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (nos. U22B2036, 62025602, 62261136549, 11931015, 62073263 and 62202381); the Fok Ying-Tong Education Foundation, China (Grant No. 171105); and sponsored by CAAI-Huawei MindSpore Open Fund and the Tencent Foundation and XPLORER PRIZE.

References

Bian, T.; Xiao, X.; Xu, T.; Zhao, P.; Huang, W.; Rong, Y.; and Huang, J. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 549–556.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.

Dou, Y.; Shu, K.; Xia, C.; Yu, P. S.; and Sun, L. 2021. User Preference-Aware Fake News Detection. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2051–2055.

Fang, L.; Feng, K.; Zhao, K.; Hu, A.; and Li, T. 2023. Unsupervised Rumor Detection Based on Propagation Tree VAE. *IEEE Transactions on Knowledge and Data Engineering*, 1–16.

Gangireddy, S. C. R.; P, D.; Long, C.; and Chakraborty, T. 2020. Unsupervised Fake News Detection: A Graph-Based Approach. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, 75–83.

He, Z.; Li, C.; Zhou, F.; and Yang, Y. 2021. Rumor Detection on Social Media with Event Augmentations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020–2024.

Hou, D.; Wang, Z.; Gao, C.; and Li, X. 2023. Sequential attention source identification based on feature representation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 4794–4802.

Hou, D.; Yin, S.; Gao, C.; Li, X.; and Wang, Z. 2024. Propagation Dynamics of Rumor vs. Non-rumor across Multiple Social Media Platforms Driven by User Characteristics. arXiv:2401.17840.

Hou, Z.; Liu, X.; Cen, Y.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. GraphMAE: Self-Supervised Masked Graph Autoencoders. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 594–604.

Hua, J.; Cui, X.; Li, X.; Tang, K.; and Zhu, P. 2023. Multimodal fake news detection through data augmentation-based contrastive learning. *Applied Soft Computing*, 136: 110125.

Jin, Y.; Wang, X.; Yang, R.; Sun, Y.; Wang, W.; Liao, H.; and Xie, X. 2022. Towards fine-grained reasoning for fake news detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5746–5754.

Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*.

Li, D.; Guo, H.; Wang, Z.; and Zheng, Z. 2021a. Unsupervised Fake News Detection Based on Autoencoder. *IEEE Access*, 9: 29356–29365.

Li, Z.; Liu, F.; Yang, W.; Peng, S.; and Zhou, J. 2021b. A survey of convolutional neural networks: analysis, applications, and prospects. *IEEE Transactions on Neural Networks and Learning Systems*.

Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B. J.; Wong, K.-F.; and Cha, M. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 3818–3824.

Min, E.; Rong, Y.; Bian, Y.; Xu, T.; Zhao, P.; Huang, J.; and Ananiadou, S. 2022. Divide-and-Conquer: Post-User Interaction Network for Fake News Detection on Social Media. In *Proceedings of the ACM Web Conference 2022*, 1148–1158.

Shu, K.; Cui, L.; Wang, S.; Lee, D.; and Liu, H. 2019. DEFEND: Explainable Fake News Detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 395–405.

Shu, K.; Mahudeswaran, D.; Wang, S.; Lee, D.; and Liu, H. 2020. FakeNewsNet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data*, 8(3): 171–188.

Shu, K.; Sliva, A.; Wang, S.; Tang, J.; and Liu, H. 2017. Fake News Detection on Social Media: A Data Mining Perspective. *SIGKDD Explor. Newsl.*, 19(1): 22–36.

Silva, A.; Luo, L.; Karunasekera, S.; and Leckie, C. 2023. Unsupervised Domain-agnostic Fake News Detection using Multi-modal Weak Signals. arXiv:2305.11349.

Sun, T.; Qian, Z.; Dong, S.; Li, P.; and Zhu, Q. 2022. Rumor Detection on Social Media with Graph Adversarial Contrastive Learning. In *Proceedings of the ACM Web Conference 2022*, 2789–2797.

Wang, Y.; Ma, F.; Jin, Z.; Yuan, Y.; Xun, G.; Jha, K.; Su, L.; and Gao, J. 2018. EANN: Event Adversarial Neural Networks for Multi-Modal Fake News Detection. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 849–857.

Wang, Y.; Qian, S.; Hu, J.; Fang, Q.; and Xu, C. 2020. Fake News Detection via Knowledge-Driven Multimodal Graph Convolutional Networks. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*, 540–547.

Wang, Z.; Hou, D.; Gao, C.; Li, X.; and Li, X. 2023. Lightweight Source Localization for Large-Scale Social Networks. In *Proceedings of the ACM Web Conference 2023*, 286–294.

Wu, L.; Lin, H.; Tan, C.; Gao, Z.; and Li, S. Z. 2021. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE Transactions on Knowledge and Data Engineering*.

Wu, L.; Rao, Y.; Zhang, C.; Zhao, Y.; and Nazir, A. 2023. Category-Controlled Encoder-Decoder for Fake News Detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(2): 1242–1257.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.

Yang, R.; Wang, X.; Jin, Y.; Li, C.; Lian, J.; and Xie, X. 2022. Reinforcement subgraph reasoning for fake news detection. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2253–2262.

Yang, S.; Shu, K.; Wang, S.; Gu, R.; Wu, F.; and Liu, H. 2019. Unsupervised Fake News Detection on Social Media: A Generative Approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 5644–5651.

Yin, X.; Han, J.; and Yu, P. S. 2008. Truth Discovery with Multiple Conflicting Information Providers on the Web. *IEEE Transactions on Knowledge and Data Engineering*, 20(6): 796–808.

You, J.; Ying, R.; Ren, X.; Hamilton, W.; and Leskovec, J. 2018. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *International conference on machine learning*, 5708–5717.

Yuan, C.; Ma, Q.; Zhou, W.; Han, J.; and Hu, S. 2019. Jointly embedding the local and global relations of heterogeneous graph for rumor detection. In *2019 IEEE international conference on data mining (ICDM)*, 796–805. IEEE.

Zhou, X.; Wu, J.; and Zafarani, R. 2020. SAFE: Similarity-Aware Multi-modal Fake News Detection. In *Pacific-Asia Conference on knowledge discovery and data mining*, 354–367.

Zhou, X.; and Zafarani, R. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. *ACM Computing Surveys*, 53(5): 1–40.