

Exploring Graph-aware Multi-View Fusion for Rumor Detection on Social Media

Yang Wu^{1,2}, Jing Yang¹, Xiaojun Zhou¹, Liming Wang¹, and Zhen Xu¹

¹ Institute of Information Engineering, Chinese Academy of Sciences

² School of Cyber Security, University of Chinese Academy of Sciences
{wuyang0419, yangjing, zhouxiaojun, wangliming, xuzhen}@iie.ac.cn

Abstract. Automatic detecting rumors on social media has become a challenging task. Previous studies focus on learning indicative clues from conversation threads for identifying rumor information. However, these methods only model rumor conversation threads from various views but fail to fuse multi-view features very well. In this paper, we propose a novel multi-view fusion framework for rumor representation learning and classification. It encodes the multiple views based on Graph Convolutional Networks (GCN), and leverages Convolutional Neural Networks (CNN) to capture the consistent and complementary information among all views and fuse them together. Experimental results on two public datasets demonstrate that our method outperforms state-of-the-art approaches.

Keywords: Graph · Multi-view Fusion · Rumor Detection.

1 Introduction

Social media has become an essential platform for people to obtain and share information. While bringing convenience, it also gives us too many dire challenges, one of which is the proliferation of rumors. A *rumor* is generally defined as a statement that emerges and spreads among people whose truth value is true, unverified, or false [5, 24, 23, 16]. Because social media lacks effective authentication techniques for user-generated content, users publish rumors without scruples, significantly reducing global Internet information credibility. Social psychology research results show that humans are only slightly higher than 50% (55%-58%) in their ability to identify deceptive information [26], which means that people are very easily deceived by false rumors. Facing the massive amount of information on social media, it seems pretty powerless to verify rumors manually. Therefore, it is necessary to develop an automatic and assistant approaches to debunk rumors on social media.

Rumor detection in this paper is the task of automatically determining the veracity value of the source claim in a conversation thread [16], which is the same as *rumor verification* defined in some other studies [13, 32, 39]. False (true) rumor means the veracity of the rumor claim is false (true) [23]. As shown in Figure 1, retweeted posts about the source claim form a conversation thread.




Fig. 1. Conversation threads of four rumorous source posts from dataset, SemEval-2017. Nodes may express stances on their parents or source node, such as supporting (●), denying (○), querying (◐), and commenting (◑). The bright yellow nodes (●) denote source posts. Stance labels come from the dataset SemEval-2017 (refer to Section 5.1 for details)

On the topological structure, false rumors propagate shallowly around the root node, while true rumors present the characteristics of multi-point and multi-branch propagation, and the branches are long. A false rumor usually needs to attract more people’s attention to spread it quickly, so its propagation is more likely to break out at the root node. In comparison, true rumors have no purpose of spreading, so their spreads are relatively scattered. On the textual content, false rumors are more possibly to be denied or queried by most replies, while true rumors tend to be supported by people [23]. Therefore, the conversation threads under different source claims (false and true rumors) tend to exhibit different patterns in both topological and textual information.

Besides textual information, existing studies mine the temporal, sequential, and structural properties of a rumorous conversation thread. Previous works model the conversation thread as a temporal model [13, 16] or a sequential model [19, 9, 39], but ignores the structure information, which also reflects the behavior of rumors. Therefore, some studies model the rumor conversation as an undirected tree-structured graph and learn its structural features by GCN [32] or Graph Attention Networks (GAT) [17]. However, there are various relationships among nodes in a conversation thread, such as the top-down propagation relationship along retweet chains [1, 7] and bottom-up dispersion relationship within communities [1, 28] as shown in Figure 2, which undirected GCN neglects. Ma et al. designed tree-structured Recursive Neural Networks (RvNN) [23] to learn rumor features from top-down view. Nevertheless, their model only learns one pattern of rumors, lacking comprehensiveness. Each type of relationship among nodes is considered as a view. The top-down view indicates rumor propagation from the source post to forwarded posts along different paths. But it cannot represent information aggregation from leaf nodes to the root node across a social community, while the cases in the bottom-up view are on the contrary. So it is desirable to utilize the complementary information from both views. Thus, recent studies use two GCNs to encode the top-down view and bottom-up view synchronously and obtain representations of two views by applying mean pooling over all nodes [1, 31]. However, they just concatenate two-view feature representations without considering node-to-node complementarity over both views. As

shown in Figure 2 left, two views share the same set of nodes (i.e., posts). Although with the initial embedding value, the node embedding will be updated according to view-specific relationships. Intuitively, the way that learning representations of individual views followed by simple concatenation loses the global perspective. Just like identifying the shape of an object according to its three-view drawing, it is easy to make mistakes if observing three views separately. Therefore, we think learning view-wise node representations and fusing them is the right way to make full use of consistent and complementary information in multiple views for rumor detection, which is still a challenge.

To achieve the aforementioned goal, we propose a novel Graph-aware Multi-View Convolutional Neural Network (GMVCN) for rumor detection. We are inspired by computer vision. A color image can be represented by red, green, and blue (RGB) color channels, and each pixel position has three values corresponding to RGB, respectively. The RGB channels are the three-color views of the image. The actual color of the image can only be seen by superimposing the three color views on each other. In computer vision, the information of RGB views is effectively integrated by CNN for tasks such as image classification [15]. Inspired by this, we regard the rumor conversation thread as a color image and each node as a pixel. The top-down propagation view and bottom-up dispersion view are regarded as two channels of the image. Both views share the same set of nodes with initial values. In order to capture view-specific features, we leverage GCN to encode individual views, which updates the node embedding according to view-specific structures. Then a CNN-based sub-module is utilized to capture the consistent and complementary information between the two views and fuse them into a global rumor representation vector for prediction.

The contributions of this work are as follows. (1) We propose a novel graph-aware multi-view convolutional neural network to effectively integrate the multi-view information of a conversation thread for rumor detection. To the best of our knowledge, this is the first study to explore the multi-view fusion of conversation threads for rumor detection on social media. (2) We innovatively treat a rumor conversation thread as an image, nodes as pixels, and multiple views as image channels. We utilize GCN to learn view-specific features and then explore to capture rumor characteristics from the consistent and complementary information of both views with CNN. (3) Experimental results on two public datasets show that our model outperforms several state-of-the-art approaches. Besides, we visualize what our proposed model has captured, and the results effectively verify that our method can learn inherent properties possessed by rumors on social media.

2 Related Work

2.1 Rumor detection

Most previous methods learn rumor characteristics from text contents, user profiles and retweet propagation [2, 6, 3, 37, 14, 27]. Ma et al. [21] developed a time series model to capture the variation of social context information over

time. Some studies model propagation patterns as tree structures based on kernel methods [34, 22], etc. However, these traditional approaches heavily rely on feature engineering, which is labor-intensive and limited.

Recently, researchers explored applying deep neural networks to detect rumors. Ma et al. [20] leveraged Recurrent Neural Networks (RNN) to learn the representations of relevant posts over spread time. Liu et al. [18] utilized both RNN and CNN to extract user profiles along the propagation time to determine the truth of posts. Lu et al. [19] proposed a hybrid-based model which considers source tweets and user profiles. Yu et al. [39] divided each long thread into shorter subthreads and learned local and global interactions among them based on a hierarchical Transformer framework. However, these approaches ignore the structure features of conversation threads, as shown in Figure 1, which can reflect how posts are propagated over social media.

In order to capture indicative patterns from both text contents and propagation structures based on deep learning methods, Ma et al. [23] proposed an RvNN-based model to learn hidden representations of tweets along propagation trees. Lin et al. [17] enhanced the representation learning of each post and the interactions among them based on undirected graph neural Networks with multiple attention mechanisms. The tree-structured graphs encoded in these methods have only one edge type, but in reality, the relationships between nodes are various. Hence, Bian et al. [1] focused on both top-down propagation relationship and bottom-up dispersion relationship among nodes for rumor detection. Wei et al. [31] removed unreliable relationships between nodes in rumor conversation threads based on Bian’s work. However, the shortcoming of these works is that they cannot effectively integrate multiple views in rumor conversations to mine the difference between false rumors and true rumors from a global perspective.

2.2 Multi-view Graph Learning

In recent years, Graph Neural Networks have been adopted for various tasks according to their remarkable performance of representation learning on structured data, e.g., text classification [41], recommendation system [33], etc. Representative graph neural networks include GCN [11], GAT [30], and so on. However, these methods are tailored for single-view network representation [35]. There are amount of multi-view networks in the real world, where each view consists of a type of relationship. Thus, substantial efforts have been dedicated to exploring multi-view graph learning [40, 36, 35]. They mainly investigated how to integrate each view’s node representations into a global node representation. Unlike them, we explore solving the problem of how to fuse features of multi-view graphs into a global graph feature representation vector for rumor detection.

3 Preliminaries

3.1 Problem Statement

We define a rumor detection dataset as a set of conversation threads $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$, where c_i is the i -th conversation thread and m is the num-

ber of conversation threads. Each thread c_i consists of a source claim r_i and a number of reply posts sorted in chronological order, which is denoted as $c_i = \{r_i, x_1^i, x_2^i, \dots, x_{n_i-1}^i, G_i\}$, where x_j^i is the j -th relevant retweet, and G_i is the propagation structure. Specifically, $G_i = \langle V_i, E_i \rangle$ is a tree-structured graph with the root node r_i , where $V_i = \{r_i, x_1^i, \dots, x_{n_i-1}^i\}$, and $E_i = \{e_{st}^i | s, t = 0, \dots, n_i - 1\}$ refers a set of edges that represents reply relationship among nodes. For example, if x_2^i has a response to x_1^i , there will be an directed edge $x_1^i \rightarrow x_2^i$ in E_i , i.e., e_{12}^i .

The rumor detection task in our paper is formulated as a supervised classification problem, aiming to learn a classification function $f : \mathcal{C} \rightarrow \mathcal{Y}$, where \mathcal{C} and \mathcal{Y} are the sets of conversations and labels, respectively.

3.2 Graph Convolutional Networks

In recent years, graph convolutional networks have been demonstrated superior performance for dealing with graph data in a variety of NLP tasks, such as text classification [38], modeling knowledge graph [29], and recommendation system [8], etc. GCNs capture information from a node's direct and indirect neighbors by multiple Graph Convolution Layers (GCL) and then update the node's representation. Given the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ for a graph G . Its layer-wise propagation rule is:

$$\mathbf{H}_{l+1} = \sigma(\hat{\mathbf{A}}\mathbf{H}_l\mathbf{W}_l) \quad (1)$$

where l is the layer number, $\hat{\mathbf{A}} = \tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{A}}\tilde{\mathbf{D}}^{-\frac{1}{2}}$ the normalized symmetric weight matrix ($\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$, $\tilde{\mathbf{D}}_{ii} = \sum_j \tilde{\mathbf{A}}_{ij}$, where \mathbf{I}_N is the identity matrix), $\mathbf{W}_l \in \mathbb{R}^{D \times D}$ is a layer-specific trainable weight matrix, $\sigma(\cdot)$ is an activation function, e.g., the ReLU function, $\mathbf{H}_l \in \mathbb{R}^{N \times D}$ is the hidden feature matrix in the l^{th} layer.

4 Proposed Method

In this section, we introduce our proposed Graph-aware Multi-View Convolutional neural Network model that learns suitable high-level representations from rumor conversation threads. Figure 2 illustrates an overview of our model, which consists of three main components, including multi-view embedding, multi-view fusing, and classification. We will depict how to apply the GMVCN model to determine the veracity value of the source post r_i in the conversation thread c_i . To better present our method, we omit the subscript i in the following content. The other conversations are calculated in the same manner.

4.1 Multi-View Embedding

Given $G = \langle V, E \rangle$, we transform the edge set E to an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, where $\mathbf{A}_{ij} = 1$ if node x_j directly replies the node x_i . Our model focuses

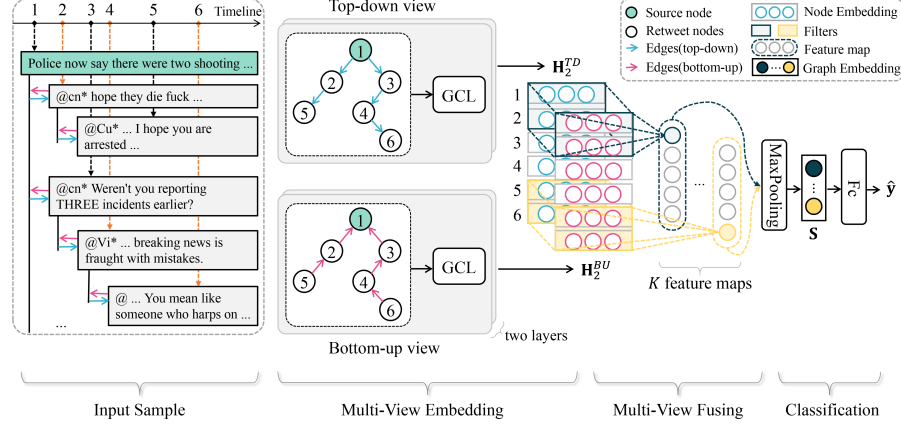


Fig. 2. The architecture of our graph-aware multi-view convolutional neural networks

on two different views of the conversation thread, the top-down propagation view and the bottom-up dispersion view, which have opposite direct edges. \mathbf{A} only contains edges in the top-down view from the parent node to its children nodes. In order to reduce overfitting of our model, we apply the Dropedge technique [25] on G as previous research [1, 31], which randomly drops out p percentage of edges. And then we obtain a new adjacency matrix \mathbf{A}' , so the adjacency matrix of the top-down view is represented as $\mathbf{A}^{TD} = \mathbf{A}'$. For bottom-up view, the adjacency matrix is represented as $\mathbf{A}^{BU} = \mathbf{A}'^T$. The nodes of two views share the same initial feature matrix \mathbf{X} , where the top-5000 words represent features of each node in terms of the TF-IDF values. Then their embeddings are updated by two different GCNs, respectively. The graph convolution operation in one GCN layer is depicted as Eq. (1). For top-down view, we feed \mathbf{A}^{TD} and \mathbf{X} to two-layer GCN as follows:

$$\mathbf{H}_1^{TD} = \sigma(\hat{\mathbf{A}}^{TD} \mathbf{X} \mathbf{W}_1^{TD}) \quad (2)$$

$$\mathbf{H}_2^{TD} = \sigma(\hat{\mathbf{A}}^{TD} \mathbf{H}_1^{TD} \mathbf{W}_2^{TD}) \quad (3)$$

where $\mathbf{H}_1^{TD} \in \mathbb{R}^{n \times d_1}$ and $\mathbf{H}_2^{TD} \in \mathbb{R}^{n \times d_2}$ denote the hidden feature representations of the two-layer GCN. $\mathbf{W}_1^{TD} \in \mathbb{R}^{d \times d_1}$ and $\mathbf{W}_2^{TD} \in \mathbb{R}^{d_1 \times d_2}$ are convolution filter matrices shared by all nodes in the conversation thread. Here $\sigma(\cdot)$ is the ReLU function. To avoid overfitting, we apply Dropout to GCN layers. For bottom-up view, the calculation of two-layer GCN is similar to Eq. (2) and Eq. (3), which obtains \mathbf{H}_1^{BU} and \mathbf{H}_2^{BU} , respectively.

4.2 Multi-View Fusing

After top-to-down and bottom-to-up graph embedding, each node n_i has two representations $\mathbf{h}_i^{TD} \in \mathbb{R}^{d_2}$ and $\mathbf{h}_i^{BU} \in \mathbb{R}^{d_2}$ corresponding to two views. We

treat each conversation thread as a color image, each view as a channel, and each node as a pixel. Therefore, the feature representations of the rumor conversation thread are formulated as follows:

$$\mathbf{H} = \begin{bmatrix} \mathbf{h}_1^{TD}, \mathbf{h}_2^{TD}, \dots, \mathbf{h}_n^{TD} \\ \mathbf{h}_1^{BU}, \mathbf{h}_2^{BU}, \dots, \mathbf{h}_n^{BU} \end{bmatrix} \in \mathbb{R}^{2 \times n \times d_2} \quad (4)$$

Given the multi-channel input \mathbf{H} , a convolution operation, which involves multiple filters $\mathbf{W}_k \in \mathbb{R}^{2 \times l \times d_2}, k \in [1, K]$, is applied to a window of t nodes to produce new features according to the following operation:

$$\mathbf{s}_i^k = \sigma(\mathbf{W}_k * \mathbf{H}_{i:i+t-1} + \mathbf{b}) \quad (5)$$

where \mathbf{W}_k and \mathbf{b} are learnable parameter matrices, $*$ denotes multi-channel convolutional operations, and $\sigma(\cdot)$ is the ReLU activation function. We apply the filters to each possible window of t nodes in \mathbf{H} , and obtain the following feature map:

$$\mathbf{s}^k = [\mathbf{s}_1^k, \mathbf{s}_2^k, \dots, \mathbf{s}_{n-t+1}^k] \quad (6)$$

We then use max-over-time pooling operation on the above feature map and get the most important feature $\hat{\mathbf{s}}^k = \max \{\mathbf{s}^k\}$ corresponding to this particular filter. All maximum values of the feature maps produced by all filters are concatenated as the final representation of the conversation thread:

$$\mathbf{S} = [\hat{\mathbf{s}}^1, \hat{\mathbf{s}}^2, \dots, \hat{\mathbf{s}}^K] \quad (7)$$

where K is the number of filters.

4.3 Classification

Finally, the learned representation of the conversation thread is fed to a full connection layer with softmax normalization for predicting the probability of the source post being a false rumor:

$$\hat{\mathbf{y}} = \text{softmax}(FC(\mathbf{S})) \quad (8)$$

The loss function is devised to minimize the cross-entropy value of the predictions and ground truth distributions:

$$\mathcal{L}_c = - \sum_i^{|Y|} \mathbf{y}^i \log \hat{\mathbf{y}}^i \quad (9)$$

where \mathbf{y}^i denotes the vector representing distribution of ground truth label for source claim in c_i .

5 Experiments and Results

In this section, we first evaluate the performance of GMVCN compared with several baselines on rumor detection. Then we further analyze our proposed model, i.e., the effect of every model component.

Table 1. Statistics of the datasets

Dataset	#Thread	#Tweet	Depth	#True rumors	#False rumors	#Unverified rumors
SemEval	325	5568	3.5	145	74	106
PHEME	2402	105354	2.8	1067	638	697

5.1 Datasets

To evaluate the effectiveness of our GMVCN model, we conduct experiments on two public real-world datasets, SemEval-2017 task 8³ and PHEME⁴. The statistics of two datasets are shown in Table 1.

SemEval-2017 task 8 [4] is composed of 325 rumorous conversation threads, which have been split into training, development, and test sets. These conversation threads are related to 10 events, and the test set covers two events data that the training and development set do not have. In addition to the rumor category label of each conversation thread, each post in the thread is annotated with its stance label (support, deny, query, or comment). Here stance labels only assist in analyzing the performance of our proposed model (refer to Section 1 and Section 5.7 for details) and will not be utilized for rumor detection.

PHEME [42] has 2402 conversation threads related to nine events. Following previous studies, we conduct 9-fold cross-validation to obtain robust results. In each fold, conversation threads related to one event are used for testing, and all the conversations related to the other eight events are used for training.

Both datasets have the problem of class distribution imbalance, so we choose Macro- F_1 as the primary evaluation metric and accuracy as the additional evaluation metric. We run each model five times on both datasets and calculate the average results.

5.2 Experimental Settings

The output sizes of two GCN layers are $d_1=d_2=64$. The window size of filters is $t=1$, and the filter number is $K=64$. We set the rate of Dropedge and that of Dropout are both 0.5. We train our model with a 0.0005 learning rate and 64 batch size. We apply an L2-regularizer to the weights of our model to prevent overfitting, and the weight penalty is 0.001. Our proposed model is trained for 100 epochs. The optimizer we used is Adam [10]. The hyperparameters of baseline methods are referenced from original papers. The source code can be available in Github⁵.

5.3 Baselines

- BranchLSTM [12]: A architecture that models sequential branches in a conversation thread based on Long Short-Term Memory (LSTM).

³ <https://alt.qcri.org/semeval2017/task8/index.php?id=data-and-tools>

⁴ https://figshare.com/articles/dataset/PHEME_dataset_for_Rumour_Detection_and_Veracity_Classification/6392078

⁵ <https://github.com/wuyang45/GMVCN-master>

- TD-RvNN [23]: A rumor detection approach models top-down propagation structure using tree-structured recursive neural networks.
- Hierarchical GCN-RNN [32]: A hierarchical framework leverages GCN and RNN to model the structure and temporal property of the threads, respectively.
- PLAN [9]: A transformer-based model that encodes the conversation thread using a randomly initialized Transformer.
- Hierarchical Transformer [39]: A extension model of BERT that first models the interactions in each subthread and then encodes global interactions of all posts based on a Transformer layer.
- Bi-GCN [1]: A GCN-based model that learns high-level representations from top-down and bottom-up views of conversation threads.
- ClaHi-GAT [17]: A rumor detection model based on GAT that represents the conversation thread as an undirected graph.
- EBGCN [31]: Variants of Bi-GCN that adaptively adjust weights of unreliable relations by a Bayesian method.

5.4 Main Results

Table 2. Results of rumor detection.

Method	SemEval		PHEME	
	Macro- F_1	Accuracy	Macro- F_1	Accuracy
BranchLSTM	0.491	0.500	0.259	0.314
TD-RvNN	0.509	0.536	0.264	0.341
Hierarchical GCN-RNN	0.540	0.536	0.317	0.356
PLAN	0.581	0.571	0.361	0.438
Hierarchical Transformer	0.592	0.607	0.372	0.441
Bi-GCN	0.607	0.617	0.316	0.442
ClaHi-GAT	0.539	0.536	0.369	0.556
EBGCN	0.639	0.643	0.375	0.521
GMVCN (Ours)	0.721	0.721	0.441	0.647

Table 2 shows the performance of our proposed model GMVCN and all baselines on SemEval 2017 and PHEME datasets. Our proposed model GMVCN has significantly superior performance among all the baselines over all metrics across two datasets.

First, we observe that all structure-based methods (RvNN, Hierarchical GCN-RNN, Bi-GCN, ClaHi-GAT, EBGCN, and our GMVCN) perform better than branchLSTM, which indicates the effectiveness of the structural message-passing mechanism. BranchLSTM decomposes one conversation thread along branches of the entire tree, and each branch is encoded by LSTM. Since LSTM can only process data with sequential structure, branchLSTM fails to learn high-level representations of rumors from structure information, resulting in its worse performance. Unlike branchLSTM, PLAN and Hierarchical Transformer are sequential

models with better performance, thanks to the powerful performance of Transformer. Transformer breaks through the spatial limitation of sequential structure and can be regarded as a fully connected graph. Therefore, mining the structural properties of conversation threads is critical for detecting rumors.

Second, due to integrating various information from multiple directed views of conversation threads, GMVCN, Bi-GCN, and EBGCN beat other structure-based models, except for Bi-GCN in PHEME, which performs slightly worse. RvNN, GCN-RNN, and ClaHi-GAT model conversations from a single view, so that they can only learn single patterns of rumors, limiting their performances. ClaHi-GAT outperforms Bi-GCN and EBGCN in the PHEME dataset. The reason might be the attention mechanisms used in ClaHi-GAT, including graph-level attentions, post-level attentions, and event-level attentions. A variety of complex attention mechanisms can capture more helpful information, which is more conducive to the performance of ClaHi-GAT on data of real-world scenarios, i.e., PHEME, where train sets and test sets are neither balanced nor correlated.

Third, GMVCN is significantly superior to Bi-GCN and EBGCN, which demonstrates that indicative features learned in multiple views fusion help improve the performance of rumor detection. Bi-GCN and EBGCN merely concatenate feature representations of multiple views as the final representation of a conversation without deeply mining inter-view complement. However, different views represent the conversation thread from different angles. Just like in spatial geometry, the shape of an object can only be accurately identified by combining three-dimensional views. Hence, the multi-view fusion allows GMVCN to capture clues of rumors from a global perspective, which help improve the performance of our model much more.

All these observations demonstrate that our proposed model GMVCN can effectively detect rumors by integrating the multi-view structural information of the conversation thread.

5.5 Ablation Study




Fig. 3. Ablation study of GMVCN on two datasets.

To analyze the effectiveness of each component of GMVCN, we perform an ablation study in this subsection. “All” denotes the entire mode. We remove each component from the entire model and obtain: (1) “-CNN”: GMVCN without CNN-based submodule. Like the approaches used in Bi-GCN and EBGCN, we apply the mean-pooling operation on top-down GCN and bottom-up GCN to obtain their representations and then concatenate both features for predicting. (2) “-BU”: GMVCN without bottom-up GCN only represents conversations from the top-down view. (3) “-TD”: GMVCN without top-down GCN. (4) “-Dir”: We model the conversation thread as an undirected tree structure, which is encoded by a two-layer GCN followed by a CNN submodule. (5) “GCN”: The vanilla GCN with no direction, i.e., GMVCN without multiple views and inter-view learning.

From Figure 3, we draw the following conclusions. First, GMVCN-BU and GMVCN-TD cannot always achieve better results than GMVCN-Dir, but our GMVCN consistently outperforms GMVCN-BU, GMVCN-TD, and GMVCN-Dir in all evaluation metrics. This indicates the superiority of simultaneously considering the top-down and bottom-up views. Second, the performance of GMVCN-TD drops more significantly than that of GMVCN-BU. It demonstrates that the top-down propagation view can better reflect the characteristics of rumors than the bottom-up dispersion view. Third, after removing the CNN-based component, GMVCN and GMVCN-Dir drop 14.3%, 7.9% in Macro- F_1 on SemEval, and 14.6%, 3.4% in Macro- F_1 on PHEME, respectively. There are similar trends in accuracy on both datasets. This proves the effectiveness of CNN in rumor detection, which not only effectively integrates multi-view information but also captures conducive features to identify rumors based on a single view.

5.6 Parameter Sensitivity

Effect of the different window sizes of filters. We vary the filter size value in CNN to investigate its influence on rumor detection performance. The results of these experiments are exhibited in Figure 4(a). We can observe that our proposed model obtains the best performance on both datasets when the window size of filter is 1. And then, the results are a fall followed by a transitory rise with the increase of the size of filters. This is in line with our intuition. Unlike the correlation between adjacent pixels in an image, there may be no direct correlation between posts in chronological order, e.g., there is no direct connection between adjacent posts (2 and 3, 4 and 5, 5 and 6) of the false rumor conversation thread in Figure 2. Therefore, when the window size is larger than 1, the model learns much noise that affects the performance. However, with the increase in window size, the correlation between users will be slightly enhanced, and the model performance will be improved somewhat. In addition, there are few participating users and data in the early stage of rumor propagation, so the smaller the window, the more conducive the model is to the early detection of rumors. This opens up the possibility of applying our model to early rumor detection.




Fig. 4. Rumor detection results of GMVCN w.r.t. different window size of filters and different rate of Dropedge on two datasets.

Effect of the different rates of Dropedge. As shown in Figure 4(b), we test the performance of GMVCN by setting hyper-parameter p from 0 to 0.6. The performance rises gradually and reaches the highest point at $p=0.5$. After that, it declines. There exist plenty of unreliable relationships in conversation threads, which leads to severe error accumulation, and weakens model robustness [31]. With the rate of Dropedge increases, the number of unreliable edges decreases. Although few sound edges are inevitably cut, the increased performance means the model is more robust and can learn more compelling features. However, with further clipping edges, the performance will decline. Our experimental evidence suggests that this reasonable rate can perform the best.

5.7 Case Study

To further illustrate that our proposed model can capture distinguishing features, we visualize feature maps of several representative samples correctly identified by GMVCN from the test set of SemEval to compare the difference between false rumors and true rumors. Feature maps are visualized using heatmap, where Y-axis represents the node index along the time sequence in the propagation conversation (0 is the root node), and the X-axis is the dimension of the feature vector. The larger the feature value, the darker the color. In Figure 5, each column is the pair of a conversation graph and its feature map, i.e., [(a)(d)], [(b)(e)], [(c)(f)], [(g)(j)], [(h)(k)], and [(i)(l)]. The meanings represented by the colors of nodes are the same as those in Figure 1.

Figure 5(a-f) shows three false rumors' conversation topology graphs and their corresponding feature maps learned by the CNN layer of our model. We can see that the three graphs have simple structures dominated by first-order propagation, and many people disagree with the source claims. Their feature maps are highly similar, where the values of the root node are always larger than those of other nodes. This indicates that the source post of a rumor always plays a vital role in rumor detection, which is consistent with the conclusions of previous work [17, 1]. From figure 5(g-l), we can find that true rumors are usually discussed by a small group of people first and then form some long-lasting branch discussions. In the three feature maps (j-l), the feature values

of early posts are significantly larger than those of later posts, which is utterly different from feature maps of false rumors.

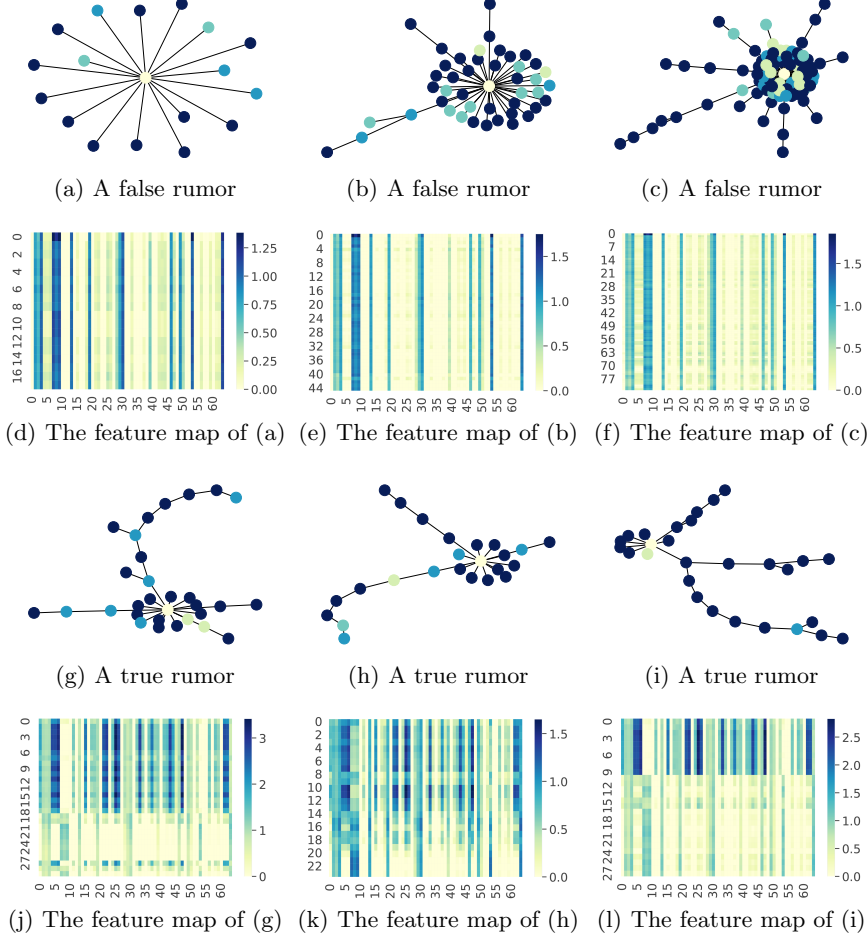


Fig. 5. False and true rumors correctly identified by GMVCN and their corresponding feature maps learned by CNN.

These comparative cases prove that our proposed GMVCN model can notice salience indicators to distinguish false rumors from true rumors by jointly learning textual and structural features. Notably, the feature maps of false (true) rumors are similar, which indicates that our model can learn shared features among all false (true) rumor events, which is beneficial to improve the performance of rumor detection in newly emerging events. Whether it is a false rumor or a true rumor, there are some common features in their feature maps

(i.e., the feature values of all nodes are highlighted in several dimensions). This might be due to the shared features of all event conversations on social media, such as tree topology, words, punctuation marks, etc.

6 Conclusion

In this paper, we proposed a novel graph-aware multi-view convolutional neural network to collaboratively integrate the abundance of information from multiple conversation views for rumor detection on social media. Inspired by computer vision, we regard a rumours conversation thread as a color image, multi-view graphs as multiple channels, and graph nodes as pixels. GCNs in our model could encode the discriminative features of each view. CNNs in our model could capture the consistent and complementary information over all views and fuse these view-level features to generate the global conversation representation. Experimental results on two real-world datasets confirm that our GMVCN outperforms state-of-the-art baselines in very large margins.

Acknowledgments

This research was supported by National Research and Development Program of China (No.2019YFB1005200).

References

1. Bian, T., Xiao, X., Xu, T., Zhao, P., Huang, W., Rong, Y., Huang, J.: Rumor detection on social media with bi-directional graph convolutional networks. In: Proceedings of the AAAI conference on artificial intelligence. pp. 549–556 (2020)
2. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on twitter. In: Proceedings of the 20th international conference on World wide web. pp. 675–684 (2011)
3. Chen, Y., Conroy, N.J., Rubin, V.L.: Misleading online content: recognizing click-bait as” false news”. In: Proceedings of the 2015 ACM on workshop on multimodal deception detection. pp. 15–19 (2015)
4. Derczynski, L., Bontcheva, K., Liakata, M., Procter, R., Hoi, G.W.S., Zubiaga, A.: Semeval-2017 task 8: Rumoureal: Determining rumour veracity and support for rumours. arXiv preprint arXiv:1704.05972 (2017)
5. DiFonzo, N., Bordia, P.: Rumor, gossip and urban legends. *Diogenes* 54(1), 19–35 (2007)
6. Feng, S., Banerjee, R., Choi, Y.: Syntactic stylometry for deception detection. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 171–175 (2012)
7. Han, S., Zhuang, F., He, Q., Shi, Z., Ao, X.: Energy model for rumor propagation on social networks. *Physica A: Statistical Mechanics and its Applications* 394, 99–109 (2014)

8. He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M.: Lightgcn: Simplifying and powering graph convolution network for recommendation. In: Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval. pp. 639–648 (2020)
9. Khoo, L.M.S., Chieu, H.L., Qian, Z., Jiang, J.: Interpretable rumor detection in microblogs by attending to user interactions. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 8783–8790 (2020)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
11. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
12. Kochkina, E., Liakata, M., Augenstein, I.: Turing at semeval-2017 task 8: Sequential approach to rumour stance classification with branch-lstm. arXiv preprint arXiv:1704.07221 (2017)
13. Kochkina, E., Liakata, M., Zubiaga, A.: All-in-one: Multi-task learning for rumour verification. arXiv preprint arXiv:1806.03713 (2018)
14. Kwon, S., Cha, M., Jung, K., Chen, W., Wang, Y.: Prominent features of rumor propagation in online social media. In: 2013 IEEE 13th international conference on data mining. pp. 1103–1108. IEEE (2013)
15. Li, Q., Cai, W., Wang, X., Zhou, Y., Feng, D.D., Chen, M.: Medical image classification with convolutional neural network. In: 2014 13th international conference on control automation robotics & vision (ICARCV). pp. 844–848. IEEE (2014)
16. Li, Q., Zhang, Q., Si, L.: Rumor detection by exploiting user credibility information, attention and multi-task learning. In: Proceedings of the 57th annual meeting of the association for computational linguistics. pp. 1173–1179 (2019)
17. Lin, H., Ma, J., Cheng, M., Yang, Z., Chen, L., Chen, G.: Rumor detection on twitter with claim-guided hierarchical graph attention networks. arXiv preprint arXiv:2110.04522 (2021)
18. Liu, Y., Wu, Y.F.: Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In: Proceedings of the AAAI conference on artificial intelligence (2018)
19. Lu, Y.J., Li, C.T.: Gcan: Graph-aware co-attention networks for explainable fake news detection on social media. arXiv preprint arXiv:2004.11648 (2020)
20. Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B.J., Wong, K.F., Cha, M.: Detecting rumors from microblogs with recurrent neural networks (2016)
21. Ma, J., Gao, W., Wei, Z., Lu, Y., Wong, K.F.: Detect rumors using time series of social context information on microblogging websites. In: Proceedings of the 24th ACM international on conference on information and knowledge management. pp. 1751–1754 (2015)
22. Ma, J., Gao, W., Wong, K.F.: Detect rumors in microblog posts using propagation structure via kernel learning (2017)
23. Ma, J., Gao, W., Wong, K.F.: Rumor detection on twitter with tree-structured recursive neural networks (2018)
24. Qazvinian, V., Rosengren, E., Radev, D., Mei, Q.: Rumor has it: Identifying misinformation in microblogs. In: Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing. pp. 1589–1599 (2011)
25. Rong, Y., Huang, W., Xu, T., Huang, J.: The truly deep graph convolutional networks for node classification. CoRR abs/1907.10903 (2019), <http://arxiv.org/abs/1907.10903>

26. Rubin, V.L.: On deception and deception detection: Content analysis of computer-mediated stated beliefs. *Proceedings of the American Society for Information Science and Technology* 47(1), 1–10 (2010)
27. Sampson, J., Morstatter, F., Wu, L., Liu, H.: Leveraging the implicit structure within social media for emergent rumor detection. In: *Proceedings of the 25th ACM international on conference on information and knowledge management*. pp. 2377–2382 (2016)
28. Thomas, S.A.: Lies, damn lies, and rumors: an analysis of collective efficacy, rumors, and fear in the wake of katrina. *Sociological Spectrum* 27(6), 679–703 (2007)
29. Tian, A., Zhang, C., Rang, M., Yang, X., Zhan, Z.: Ra-gcn: Relational aggregation graph convolutional network for knowledge graph completion. In: *Proceedings of the 2020 12th International Conference on Machine Learning and Computing*. pp. 580–586 (2020)
30. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. *stat* 1050, 20 (2017)
31. Wei, L., Hu, D., Zhou, W., Yue, Z., Hu, S.: Towards propagation uncertainty: Edge-enhanced bayesian graph convolutional networks for rumor detection. *arXiv preprint arXiv:2107.11934* (2021)
32. Wei, P., Xu, N., Mao, W.: Modeling conversation structure and temporal dynamics for jointly predicting rumor stance and veracity. *arXiv preprint arXiv:1909.08211* (2019)
33. Wei, Y., Wang, X., He, X., Nie, L., Rui, Y., Chua, T.S.: Hierarchical user intent graph network for multimedia recommendation. *IEEE Transactions on Multimedia* (2021)
34. Wu, K., Yang, S., Zhu, K.Q.: False rumors detection on sina weibo by propagation structures. In: *2015 IEEE 31st international conference on data engineering*. pp. 651–662. IEEE (2015)
35. Xie, Y., Zhang, Y., Gong, M., Tang, Z., Han, C.: Mgat: Multi-view graph attention networks. *Neural Networks* 132, 180–189 (2020)
36. Xu, L., Wei, X., Cao, J., Yu, P.S.: Multi-task network embedding. *International Journal of Data Science and Analytics* 8(2), 183–198 (2019)
37. Yang, F., Liu, Y., Yu, X., Yang, M.: Automatic detection of rumor on sina weibo. In: *Proceedings of the ACM SIGKDD workshop on mining data semantics*. pp. 1–7 (2012)
38. Yao, L., Mao, C., Luo, Y.: Graph convolutional networks for text classification. In: *Proceedings of the AAAI conference on artificial intelligence*. pp. 7370–7377 (2019)
39. Yu, J., Jiang, J., Khoo, L.M.S., Chieu, H.L., Xia, R.: Coupled hierarchical transformer for stance-aware rumor verification in social media conversations (2020)
40. Zhang, H., Qiu, L., Yi, L., Song, Y.: Scalable multiplex network embedding. In: *IJCAI*. vol. 18, pp. 3082–3088 (2018)
41. Zhang, Y., Yu, X., Cui, Z., Wu, S., Wen, Z., Wang, L.: Every document owns its structure: Inductive text classification via graph neural networks. *arXiv preprint arXiv:2004.13826* (2020)
42. Zubiaga, A., Liakata, M., Procter, R., Wong Sak Hoi, G., Tolmie, P.: Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PloS one* 11(3), e0150989 (2016)