

DDGCN: Dual Dynamic Graph Convolutional Networks for Rumor Detection on Social Media

Mengzhu Sun¹, Xi Zhang^{1*}, Jiaqi Zheng¹, Guixiang Ma²

¹Key Laboratory of Trustworthy Distributed Computing and Service (MoE),
Beijing University of Posts and Telecommunications, China

²University of Illinois at Chicago
{2019110945, zhangx, zjq1}@bupt.edu.cn, guixiang.ma@intel.com

Abstract

Detecting rumors on social media has become particular important due to the rapid dissemination and adverse impacts on our lives. Though a set of rumor detection models have exploited the message propagation structural or temporal information, they seldom model them altogether to enjoy the best of both worlds. Moreover, the dynamics of knowledge information associated with the comments are not involved, either. To this end, we propose a novel Dual-Dynamic Graph Convolutional Networks, termed as DDGCN, which can model the dynamics of messages in propagation as well as the dynamics of the background knowledge from Knowledge graphs in one unified framework. Specifically, two Graph Convolutional Networks are adopted to capture the above two types of structure information at different time stages, which are then combined with a temporal fusing unit. This allows for learning the dynamic event representations in a more fine-grained manner, and incrementally aggregating them to capture the cascading effect for better rumor detection. Extensive experiments on two public real-world datasets demonstrate that our proposal yields significant improvements compared to strong baselines and can detect rumors at early stages.

Introduction

With the proliferation of social media, it is convenient for users to obtain information, express opinions and communicate with each others. However, social media also enables the widespread of false rumor information at a high rate, bringing threats to society and causing disorders, especially during big events such as the U.S. presidential election election (Allcott and Gentzkow 2017) and the COVID-19 pandemic (Islam et al. 2020). Therefore, it is of prominent importance to detect rumors on social media at very early stage.

In order to scale with the increasing amount of misinformation, massive efforts have been devoted to automatic rumor detection. Conventional detection methods mainly focus on text mining with textual contents, either based on feature engineering (Castillo, Mendoza, and Poblete 2011; Kwon et al. 2013; Liu et al. 2015; Ma et al. 2015) or deep learning models (Ma et al. 2016; Ruchansky, Seo, and Liu 2017). Recent studies have shown that the diffusion patterns modeled as propagation tree or graph structures can

provide useful clues for distinguishing rumors from non-rumors. They thus propose kernel-based models (Wu, Yang, and Zhu 2015; Ma, Gao, and Wong 2017), recursive neural network (Ma, Gao, and Wong 2018) or graph-based models (Huang et al. 2019; Bian et al. 2020) to learn high-level spatial structural representations. However, these spatial structure based approaches largely overlooked or oversimplified the temporal structure associated with the message propagation.

The temporal structure refers to the sequence and interval of the (replied or forwarded) messages along the timeline, which can be used to further differentiate the diffusion patterns (Huang et al. 2020; Xia, Xuan, and Yu 2020). For example, considering two comments r_a and r_b which corresponds to the same post. r_a is taking place before r_b or after r_b can lead to the same spatial structure but different temporal structures. Recent studies have highlighted the importance of incorporating the temporality for rumor detection as it can model more fine-grained dynamics of the message stream, achieving better detection performance, especially in early detection situations. Therefore, it is desirable to consider both spatial and temporal structures in message propagation, which is rarely covered by previous studies.

Meanwhile, another parallel line of works focus on introducing the extra knowledge (e.g., knowledge graph) for rumor identification (Zhang et al. 2019; Wang et al. 2020; Hu et al. 2021). They commonly utilize the knowledge information associated with the post contents to complement semantic information and thus improve the high-level post representations. The intuition behind is that different from conventional classification tasks, the rumor detection models have to deal with many new and unseen events. Due to the limited words posted on social media, the extra background knowledge can greatly assist in judging the credibility of a post, which is also a common paradigm used by human judgement. However, most of these knowledge based methods only use the background knowledge information in the source post, ignoring that from the user comments. Additionally, the knowledge information would also evolve as the message spreads. However, such knowledge dynamics is not exploited by existing models.

In this paper, we aim to simultaneously model the dynamics of messages in propagation as well as the dynamics of the associated background knowledge in one unified frame-

*Corresponding author

work for detecting rumors in a timely manner. Intuitively, to effectively encode the structure information among propagating messages and among knowledge entities, we would like to use graph convolutional networks (GCNs) to learn their high-level graph representations. However, conventional GCN models are not able to track the evolving graph representations. To this end, we propose a novel Dual-Dynamic Graph Neural Network (DDGCN) framework, which includes two coupled dynamic GCNs, one of which operates on the evolving propagation graphs to capture both spatial and temporal structure of the messages as dynamic propagation representations, and the other operates on the evolving knowledge graphs associated with the messages to learn dynamic knowledge representations. Furthermore, we propose a sequential fusing method to combine the above two representations. Specifically, it consists of a sequence of temporal fusing units, and each unit would combine the two intermediate representations at the end of each time stage, and then passes the fused information to the successor unit in a sequential manner. This framework can therefore incrementally learn better representations of an event as time goes on and also enables rumor detection at early stages. The main contributions are summarized as follows:

- To the best of knowledge, we are the first to consider the dynamic characteristics of the knowledge information for the rumor detection task.
- We present a novel dual-dynamic GCN by modeling both the dynamics of propagation structure and the dynamics of knowledge entity structure, and incrementally fuse them at each time stage with temporal fusing units.
- We empirically show our proposed method can not only outperform the strong baselines on two real world datasets but also have a good ability on detecting rumors at an early stage.

Related Work

Spatial structure based rumor detection. Diffusion patterns modeled as propagation trees or graph structures can provide useful clues for distinguishing rumors from non-rumors. Early methods rely on hand-crafted feature engineering to extract spatial structure features (Wu, Yang, and Zhu 2015; Ma, Gao, and Wong 2017). Recently, a line of deep neural networks are proposed to capture the propagation patterns. Ma et al. (Ma, Gao, and Wong 2018) presents a tree based recursive neural network to capture the content semantics and propagation cues. Huang et al. (Huang et al. 2019) adopts GCNs to capture the spatial structure of message propagation. Bian et al. (Bian et al. 2020) proposes a novel bi-directional graph model by operating on both top-down and bottom-up propagation of rumors. Song et al. (Song et al. 2021) proposes an adversary-aware model to generate the adversarial responses with the consideration of the response position in the propagation structure.

Temporal structure based rumor detection. The temporal structure in the information diffusion process can also provide useful features for rumor detection. A line of methods extract handcrafted temporal features (Kwon et al. 2013; Ma et al. 2015). Deep neural networks have also been

utilized to capture more effective temporal structure. Ma et al. (Ma et al. 2016) exploits a recurrent neural network based model to capture the variation of semantics in propagation. Liu et al. (Liu and Wu 2018) models the temporal structure by combining the recurrent and convolutional networks. Xia et al. (Xia, Xuan, and Yu 2020) proposes a state-independent and time-evolving network for rumor detection based on fine-grained event state detection and segmentation. Huang et al. (Huang et al. 2020) proposes a spatial-temporal structure neural network for rumor detection. However, most of spatial structure based or temporal structure based methods only consider single structure information, and none of them has involved the external knowledge to facilitate the task.

Knowledge-enhanced detection. The knowledge-based rumor detection models can be classified into two categories. One line of work (Ciampaglia et al. 2015; Fionda and Pirrò 2018; Pan et al. 2018; Shi and Weninger 2016) focuses on fact checking, which commonly uses structure triples (head, relation, tail) extracted from the post contents to compare with the faithful triples from KG. However, such fact-checking methods can not satisfy many real-world scenarios since they lack ground-truth knowledge for reference. The other line of work uses external knowledge to supplement the post contents to produce better representations for rumor detection. Zhang et al. (Zhang et al. 2019) utilizes the multimodal knowledge-aware representation and the event-invariant features to form the event representations. Cui et al. (Cui et al. 2020) incorporates an article-entity bipartite graph and a medical knowledge graph to better model the embedding of news. KMGCN (Wang et al. 2020) is proposed to model the global structure among texts, images, and knowledge concepts to obtain comprehensive semantic representations. CompareNet (Hu et al. 2021) compares the news articles to the knowledge base (KB) through entities for fake news detection. However, these methods neither consider the knowledge information in the comments nor the dynamics of knowledge information, which has been fully utilized in this work.

Dynamic Graph Convolutional Networks. GCNs have been used in a variety of graph-based tasks, and also have proved its effectiveness for rumor detection by modeling the structure among posts, comments and users (Yang et al. 2020; Bian et al. 2020). However, these models don't consider the temporal information in diffusion. One related work is the dynamic GCN proposed for predicting social events (Deng, Rangwala, and Ning 2019), but it doesn't consider the knowledge information and is only equipped with one dynamic GCN.

To summarize, we distinguish ourselves from these approaches by considering the dynamics of knowledge information as well as the dynamics of the message propagation information in one unified framework. A comparison of different models is provided in Table 3.

Methodology

Problem Definition

The rumor detection task can be defined as a binary classification problem. Let $\mathcal{E} = \{\mathcal{E}_1, \dots, \mathcal{E}_n\}$ be a set of event

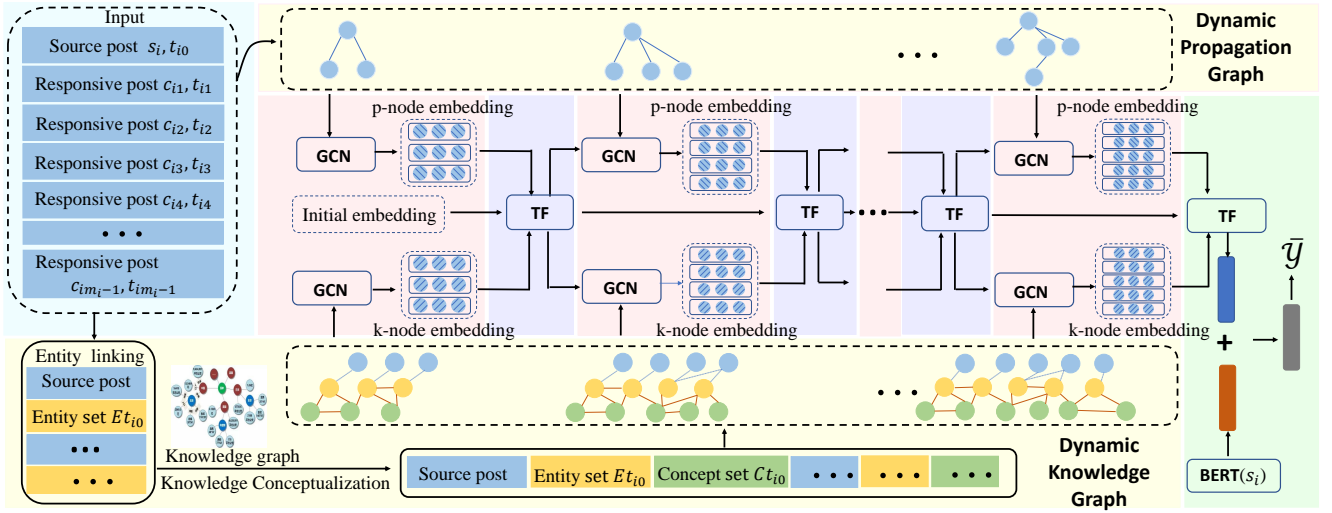


Figure 1: The framework of the proposed dual-dynamic graph convolutional network. It consists of three components: (1) top and bottom yellow block: the dynamic graph construction module to build the temporal rumor propagation graph and the temporal event-entity-concept tripartite knowledge graph; (2) middle block: the dual-dynamic GCN module composed of dual-static GCN units (pink block) and temporal fusion units (purple block); (3) right green block: a rumor classification module. The left blue block is the inputs of the model.

Notations	Descriptions
\mathcal{E}_i	the i -th event
\mathcal{E}_i^c	post and comment sequence of the i -th event
\mathcal{E}_i^t	release time sequence of the i -th event
γ	the number of time stages
Δt_i	equal time interval of event \mathcal{E}_i
\mathcal{E}_{ir}	the r -th stage subset of \mathcal{E}_i
$\mathcal{G}_{ir}^p = \langle V_{ir}^p, E_{ir}^p \rangle$	propagation graph at stage r
$\mathcal{G}_{ir}^k = \langle V_{ir}^k, E_{ir}^k \rangle$	knowledge graph at stage r

Table 1: Important notations and descriptions

instances for rumor detection, where \mathcal{E}_i is the i -th event and n is the number of events. We denote the set of the post and comment contents as $\mathcal{E}_i^c = \{s_i, c_{i1}, \dots, c_{im_i-1}\}$, where s_i is the source post text, and c_{ij} is the j -th comment text and m_i refers to the number of post and comments in \mathcal{E}_i . Note that, s_i can also be regarded as c_{i0} . We can also obtain the relative release time sequence \mathcal{E}_i^t of all the posts for event \mathcal{E}_i , that is, $\mathcal{E}_i^t = \{t_{i0}, t_{i1}, \dots, t_{im_i-1}\}$, where $t_{i0} = 0$ ($i = 1, 2, \dots, n$). Then \mathcal{E}_i^c and \mathcal{E}_i^t are combined to get $\mathcal{E}_i = \{(c_{i0}, t_{i0}), (c_{i1}, t_{i1}), \dots, (c_{im_i-1}, t_{im_i-1})\}$.

Then we divide \mathcal{E}_i into γ stages along its time span (γ is thus a hyperparameter). For every stage $r \in \{1, 2, \dots, \gamma\}$, it has an equal time interval $\Delta t_i = \frac{t_{im_i-1}}{\gamma}$. Consequently, the r -th sub-event of \mathcal{E}_i is $\mathcal{E}_{ir} = \{(c_{i\pi}, t_{i\pi}) \mid t_{i\pi} \leq r\Delta t_i\}$.

We need to learn a model $f : \mathcal{E} \rightarrow \mathcal{Y}$ to classify each event \mathcal{E}_i into the predefined categories $\mathcal{Y} = \{0, 1\}$, which is the ground-truth label of the event (0 denotes non-rumor while 1 denotes rumor).

For ease of understanding, we list the important mathematical notations used throughout the paper in Table 1.

Overview

Figure 1 presents our framework, which mainly consists of three components: (1) a Dynamic Graph Construction Module to build the temporal rumor propagation graph and the temporal event-entity-concept tripartite knowledge graph, respectively; (2) a Graph Convolution Networks Module composed of dual-static GCN units and temporal fusion units to obtain the structural semantic features of an event and to fuse propagation and knowledge information at each time stage; (3) a Classification Module that aggregates the final propagation, knowledge and textual information and produces classification labels.

Dynamic Graph Construction Module

We use this module to build two dynamic graphs: one is the dynamic propagation graph and the other is the dynamic knowledge graph, which are illustrated separately.

Dynamic Propagation Graph Construction. For event \mathcal{E}_i , based on its source post and comments, we construct a sequence of propagation graphs $\{G_{i1}^p, \dots, G_{i\gamma}^p\}$, where $G_{ir}^p = \langle V_{ir}^p, E_{ir}^p \rangle$ is the propagation graph at the r -th time stage. V_{ir}^p denotes the set of vertices in G_{ir}^p , thus $V_{ir}^p = \{c_{i\pi} \mid t_{i\pi} \leq r\Delta t_i\}$. E_{ir}^p represents the set of edges between the source post and its comments or between two comments. For example, if c_{i2} is a comment on c_{i1} , there would be an edge between them in graph G_{ir}^p . For simplicity, we don't consider the directions on edges and treat G_{ir}^p as an undirected graph, whose adjacency matrix is $A_{ir}^p \in \{0, 1\}^{m_i \times m_i}$. Besides, we initialize the representation of each node l with its word embedding vector $h_l^p \in \mathbb{R}^F$.

Dynamic Knowledge Graph Construction. Different from existing studies that merely exploit the knowledge in-

formation in the source post for rumor detection, we additionally introduce the knowledge information from comments. Thus, the knowledge information will also evolve as the event propagates. The purpose of this module is to capture such knowledge dynamics. More specifically, we not only extract the knowledge entities from the source post and comment texts, but also model the structural information between entities to obtain richer knowledge semantics.

We first introduce the knowledge extraction and conceptualization procedure. Given the source post and comments, the entity linking solution TAGME¹ and shuyantech² (Chen et al. 2018) are used to link the ambiguous entity mentions to the corresponding entities in KG for English and Chinese texts, respectively. Then, for each identified entity, we acquire its conceptual information from an existing KG, such as YAGO and Probase by conceptualization. We take “isA” relation to get the concepts. For instance, given a short text “Welcome to Ferguson where Americans started waking up to the militarization of their police force.”, we obtain the entity set $T = \{\text{Ferguson unrest, Militarization, Police force}\}$ by entity linking. Then, we conceptualize the entities in T and acquire its concept set $Concept_{\text{Ferguson unrest}} = \{\text{Event, Case, City}\}$, $Concept_{\text{Militarization}} = \{\text{Social movement, Government policy}\}$, and $Concept_{\text{Police force}} = \{\text{Organization, Institution}\}$ from external knowledge graphs. Following this procedure, given an event \mathcal{E}_i , we can obtain its entity set Et_i and concept set Ct_i .

We then illustrate how to construct the dynamic knowledge graphs for each time stage, i.e., $\{G_{i1}^k, \dots, G_{i\gamma}^k\}$, where $G_{ir}^k = \langle V_{ir}^k, E_{ir}^k \rangle$ is the knowledge graph at the r -th stage for event \mathcal{E}_i , whose vertex set is the combination of the vertex set of the corresponding propagation graph, the entity set and the concept set, that is, $V_{ir}^k = V_{ir}^p \cup Et_{ir} \cup Ct_{ir}$. Note that we have built the edges between posts and comments in the propagation graph, and here we do not involve such edges in the knowledge graph but introduce other types of edges with the following rules:

(1) We build a post-entity edge between a post (or comment) node from V_{ir}^p and a entity node from Et_{ir} if the post (or comment) node contains a word that can be linked to the entity. Their edge weight is set based on the term frequency-inverse document frequency (TF-IDF) of the entity in the post, where term frequency is the number of times the entity appears in the post, and inverse document frequency is the logarithmically scaled inverse fraction of the number of posts that can be connected to the entities;

(2) We build entity-entity edges, entity-concept edges, concept-concept edges according to their point-wise mutual information (PMI), which is a widely used measure for associations. We use PMI to calculate the weights between two nodes, i.e., either entity nodes or concept nodes. Specifically, we employ a fixed-size window on all posts for gathering node co-occurrence statistics. Then, we calculate the PMI of node pairs following the procedure in (Wang et al. 2020). Note that the statistics are based on the global corpus rather than a specific post.

¹<https://sobigdata.d4science.org/group/tagme/>

²<http://shuyantech.com/entitylinking>

Therefore, for adjacency matrix A_{ir}^k of the dynamic knowledge graph, we only preserve the entity-entity, entity-concept or concept-concept edges with positive PMI scores and post-entity edges with TF-IDF scores.

We initialize the representations of each node v_l^k in the dynamic knowledge graph with its word embedding vector $h_l^k \in \mathbb{R}^F$. Note that, if a node h_l^k exists in both V_{ir}^p and V_{ir}^k , it has the same initial embedding in both graphs.

Dual-Dynamic GCN Module

A dual-dynamic GCN module is composed of dual-static GCN units and temporal fusion units. The numbers of both units are equal to the number of time stages, i.e., γ . We take the dynamic propagation graph and the dynamic knowledge graph as inputs and encode both of them at each time stage in a sequential manner.

A single Dual-Static GCN unit. We adopt two GCNs to encode dynamic propagation graph and dynamic knowledge graph simultaneously to obtain propagation-level features and knowledge-level features at the same time.

We first introduce how to encode the dynamic propagation graph, and the procedure for dynamic knowledge graph is similar. Let $H_{ir}^p \in \mathbb{R}^{m_i \times F}$ be a matrix containing features of all m_i nodes in propagation graph G_{ir}^p , where F is the dimension of the feature vectors. Let $h_{iru}^p \in \mathbb{R}^F$ denote the feature vector for node u in the graph. Its adjacency matrix is $A_{ir}^p \in \mathbb{R}^{m_i \times m_i}$. Then a two-layer GCN is defined as

$$H_{ir}^{p(1)} = g \left(\hat{A}_{ir}^p H_{ir}^{p(0)} W_r^{p(0)} + b_r^{p(0)} \right) \quad (1)$$

$$H_{ir}^{p(2)} = g \left(\hat{A}_{ir}^p H_{ir}^{p(1)} W_r^{p(1)} + b_r^{p(1)} \right) \quad (2)$$

where $H_{ir}^{p(0)} = H_{ir}^p$, $W_r^{p(0)} \in \mathbb{R}^{F \times F^1}$, $b_r^{p(0)} \in \mathbb{R}^{F^1}$, $W_r^{p(1)} \in \mathbb{R}^{F^1 \times F^2}$, $b_r^{p(1)} \in \mathbb{R}^{F^2}$ are trainable parameters. F^1 and F^2 are the output node feature dimension sizes for the first and second layer respectively. g is a non-linear activation function. \hat{A}_{ir}^p is the normalized symmetric adjacency matrix defined as:

$$\hat{A}_{ir}^p = (\tilde{D}_{ir}^p)^{-\frac{1}{2}} \tilde{A}_{ir}^p (\tilde{D}_{ir}^p)^{-\frac{1}{2}} \quad (3)$$

Here, $\tilde{A}_{ir}^p = A_{ir}^p + I_{m_i}$ and \tilde{D}_{ir}^p is the degree matrix. I_{m_i} is an identity matrix with dimensions of m_i and $\tilde{D}_{ir}^p = \sum_v \tilde{A}_{ir}^p$.

The source post of an event plays an indispensable role in rumor detection. To fully utilize such information, enlightened by the idea of root feature enhancement in (Bian et al. 2020), we concatenate the hidden feature vector of each node with the node vector of the source post learned from the previous layer, and obtain an enhanced feature matrix for the l -th ($l = \{1, 2\}$) GCN layer as

$$\tilde{H}_{ir}^{p(l)} = \text{Concat} \left(H_{ir}^{p(l)}, (H_{ir}^{p(l-1)})_{\text{Source}} \right) \quad (4)$$

By substituting $H_{ir}^{p(l)}$ with the source enhanced $\tilde{H}_{ir}^{p(l)}$ in Eq.(2), we obtain

$$H_{ir}^{p(2)} = g \left(\hat{A}_{ir}^p \tilde{H}_{ir}^{p(1)} W_r^{p(1)} + b_r^{p(1)} \right) \quad (5)$$

$$\tilde{H}_{ir}^{p(2)} = \text{Concat} \left(H_i^{p(2)}, (H_i^{p(1)})^{\text{Source}} \right) \quad (6)$$

$$\tilde{H}_{ir}^{p(2)} = \text{ReLU}(W_r^p \tilde{H}_{ir}^{p(2)} + b_r^p) \quad (7)$$

where $W_r^{p(1)} \in \mathbb{R}^{2F^1 \times F^2}$, $b_r^{p(1)} \in \mathbb{R}^{F^2}$, the new $\tilde{H}_{ir}^{p(2)} \in \mathbb{R}^{m_i \times F}$, $W_r^p \in \mathbb{R}^{2F^2 \times F}$, $b_r^p \in \mathbb{R}^F$. And for simplicity, we set $F = F^1 = F^2$.

Similarly, the hidden feature matrices of the dynamic knowledge graph with source enhancement, $\tilde{H}_{ir}^{k(1)}$ and $\tilde{H}_{ir}^{k(2)}$ are obtained in the same manner as Eq.(4)-(7).

A Temporal Fusion Unit. Based on a single Dual-Static GCN unit, we obtain node embeddings at each time stage from the two GCNs. After that, inspired by the gating mechanism, a Temporal Fusion Unit is proposed to combine and project node embeddings, and the fused embedding will serve as the initial node embedding for the next-stage Dual-Static GCN unit. Specifically, we combine three node embeddings, that is, two node embeddings learned with GCNs on the propagation graph and knowledge graph, as well as the initial node embeddings at the first time stage. For this purpose, we employ three linear transformations to $\tilde{H}_{ir}^{p(2)}$, $\tilde{H}_{ir}^{k(2)}$ and H_{i0}^p , respectively, with three trainable weight matrices $W_{tf}^0, W_{tf}^p, W_{tf}^k \in \mathbb{R}^{F \times F}$.

$$\begin{aligned} *H_{i0}^p &= W_{tf}^0 H_{i0}^p \\ *H_{ir}^{p(2)} &= W_{tf}^p \tilde{H}_{ir}^{p(2)} \\ *H_{ir}^{k(2)} &= W_{tf}^k \tilde{H}_{ir}^{k(2)} \end{aligned} \quad (8)$$

Then we apply concatenations on the three transformed features. Note that we only concatenate the source post and comment node embeddings, as they exist in both graphs. Their initial node embeddings at the first time are both pre-trained word embeddings and thus identical. The concatenation matrix goes through a linear function and a tanh activation function, and produces \hat{H}_{ir} , that is,

$$\begin{aligned} \text{Concat}_H &= \text{Concat}(*H_{i0}^p, *H_{ir}^{p(2)}, *H_{ir}^{k(2)}[V_{ir}^p]) \\ \hat{H}_{ir} &= \text{tanh}(W_{tf} \text{Concat}_H + b_{tf}) \end{aligned} \quad (9)$$

where $W_{tf} \in \mathbb{R}^{3F \times F}$, $b_{tf} \in \mathbb{R}^F$. \hat{H}_{ir} would be used by the Dual-Dynamic GCN module in Eq.(11) as the initial node embeddings for the two GCNs at time stage $r + 1$, i.e.,

$$\begin{aligned} \hat{H}_{ir}^p &= \hat{H}_{ir} \\ \hat{H}_{ir}^k[V_{ir}^p] &= \hat{H}_{ir} \\ \hat{H}_{ir}^k[V_{ir}^k - V_{ir}^p] &= \tilde{H}_{ir}^{k(2)}[V_{ir}^k - V_{ir}^p] \end{aligned} \quad (10)$$

In Eq. (8)(9), at each time stage, we concatenate the initial node embeddings $*H_{i0}^p$ in order to emphasize the importance of the initial semantic information. Note that this is different from the source enhancement method in Eq.(4), which only involves the source post information in the form of its node embedding learned at the previous stage. In contrast, we further utilize the semantic embeddings of comments, and use their initial embeddings at the first stage rather than the previous stage. So the two methods can complement to each other and cooperate to learn better node embeddings.

Connecting the Above Two Units. As the propagation graph and the knowledge graph differ at different time stages, it is necessary to capture such dynamics to allow for better modeling the temporal semantics. To this end, we propose to connect a dual-static GCN unit and a temporal fusion unit to encode and fuse the graph structure information for each time stage in the form of node embeddings, and the embeddings will be propagated to the next-stage dual-static GCN as initial node embeddings. In this way, we can capture and aggregate the dynamic graph structure information in a sequential manner.

Specifically, we design the following cross-stage graph convolutional layer, that is

$$\begin{aligned} H_{i(r+1)}^p &= g \left(\hat{A}_{ir+1}^p \hat{H}_{ir}^p W_{r+1}^{p(0)} + b_{r+1}^{p(0)} \right) \\ H_{i(r+1)}^k &= g \left(\hat{A}_{ir+1}^k \hat{H}_{ir}^k W_{r+1}^{k(0)} + b_{r+1}^{k(0)} \right) \end{aligned} \quad (11)$$

where $\hat{A}_{ir+1}^p \in \mathbb{R}^{m_i \times m_i}$ and $\hat{A}_{ir+1}^k \in \mathbb{R}^{m_i \times m_i}$ are the normalized symmetric adjacency matrices at time stage $r + 1$ as defined in Eq.(3). $W_{r+1}^{p(0)}$ and $W_{r+1}^{k(0)} \in \mathbb{R}^{F \times F^1}$, and $\{b_{r+1}^{p(0)}, b_{r+1}^{k(0)}\} \in \mathbb{R}^{F^1}$ are trainable parameters of the dual dynamic GCN first layer at time stage $r + 1$. Note that \hat{H}_{ir}^p and \hat{H}_{ir}^k come from the temporal fusion unit.

Rumor Classification Module

The output of the last TF module are the post node and comment node representations. Here we employ mean-pooling operators to aggregate information from the node representations. It is formulated as

$$H_i = \text{MEAN}(\hat{H}_{i\gamma}^p) \quad (12)$$

Then, we concatenate the representations of propagation and knowledge H_i and representation of source post node s_i to merge the information as

$$\tilde{S} = \text{Concat}(H_i, \text{BERT}(s_i)) \quad (13)$$

Finally, the label of the event \hat{y} is calculated via several full connection layers and a sigmoid layer:

$$\hat{y}_i = \sigma(w_f \tilde{S} + b_f) \quad (14)$$

where w_f and b_f are the weight and bias parameters. We then use cross entropy loss as the rumor classification loss:

$$\mathcal{L}_c = - \sum_i y_i \log \hat{y}_i \quad (15)$$

where y_i is the ground truth label of the i -th instance.

Experiments

Datasets

We conduct experiments on two public real-world datasets, i.e., Pheme (Zubiaga, Liakata, and Procter 2017) and Weibo (Ma, Gao, and Wong 2018), which are collected from Twitter and Weibo respectively. Each event has a label, rumor or non-rumor. The datasets contain rich information including the post texts, user comments, and their release time. Table 2 shows the statistics of the datasets.

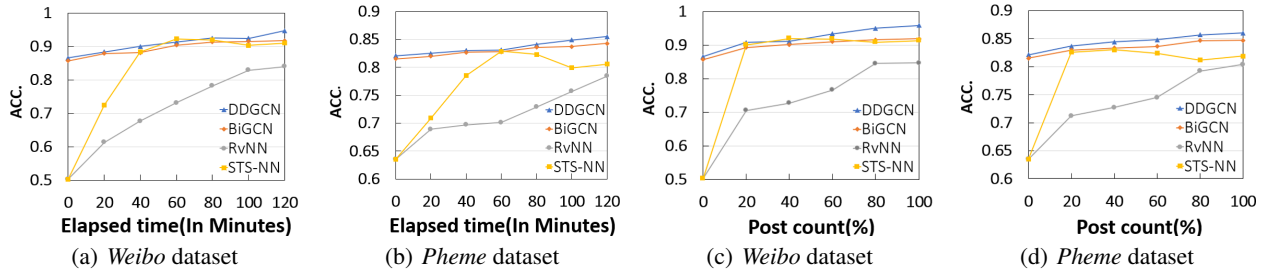


Figure 2: Early rumor detection accuracy with the increase of observation time or percentage of the number of comments.

Statistic	Weibo	PHEME
# of events	4657	5748
# of Non-rumors	2312	3654
# of Rumors	2345	2094
Avg. time length/ event (min)	108506	347
Avg. # of posts/ event	804	16
Max # of posts/ event	59318	346
Min # of posts/ event	10	1

Table 2: Statistics of Datasets

Comparison Methods

We compare with the following baselines:

BERT (Devlin et al. 2018) is a pre-trained language model based on deep bidirectional transformers. We use it to obtain the representation of the source post for classification.

Transformer (Vaswani et al. 2017) uses the self-attention mechanism and position encoding to extract textual features for Seq2seq learning. We only use its encoder to learn text representations for classification.

EANN (Wang et al. 2018) is a GAN-based model to extract event invariant features to facilitate detecting newly arrived events. Note that different from the original EANN, we don't use pictures due to the lack of pictures in our dataset.

QSAN (Tian et al. 2020) integrates the quantum-driven text encoding and a signed attention mechanism to model complex semantics between source post and responsive posts.

RumorGAN (Ma, Gao, and Wong 2019) generates uncertain or conflicting voices to enhance the discriminator to learn stronger rumor representations.

RVNN (Ma, Gao, and Wong 2018) learns discriminative features from contents by following their non-sequential propagation structure. RVNN includes a bottom-up and a top-down tree-structured network, denoted as $RVNN_{BU}$ and $RVNN_{TD}$ respectively.

BiGCN (Bian et al. 2020) is a GCN-based model that can embed both propagation and dispersion structures and enhance node representations with root node features.

KMGCN (Wang et al. 2020) uses a graph convolution network to incorporate visual information and KG to enhance the semantic representation. We don't use visual information due to the lack of pictures in our dataset.

STS-NN (Huang et al. 2020) jointly models the spatial structure and temporal structure in message propagation.

Table 3 compares these approaches from the perspective of feature modeling.

Experiment Setup

We adopt the default optimization settings reported in corresponding papers for all comparison methods. We implement our method with Pytorch framework (Paszke et al. 2017). The parameters are optimized using Adam algorithm (Kingma and Ba 2014). We split the PHEME dataset and Weibo dataset into training, validation, and testing set with a split ratio of 7:1:2 without overlapping. We select the best parameter settings based on the performance on the validation set. In our model, we use the pre-trained BERT (Devlin et al. 2018): bert-base-uncased for English, and bert-base-chinese for Chinese. We employ Accuracy, Precision, Recall, and F1 as evaluation metrics. More reproducibility details are listed in Appendix. We randomly split the datasets into five parts, and conduct 5-fold cross-validation to obtain the final results. We set the number of time stages $\gamma = 3$.

Rumor Classification Performance

Table 3 shows the performance of the compared models. On both datasets, our model DDGCN significantly outperforms all the other approaches in all the metrics, which confirms that considering the dual-dynamic information would benefit the rumor detection task.

BERT outperforms the other models that work with only text information, demonstrating its superior capability in capturing the textual semantics for rumor detection. We can observe that the performance of RumorGAN is comparable to the models that work only with text information. We can also observe that the methods with the spatial structure and text information perform better than those methods that work only with text and temporal structure. For example, compared to RumorGAN, BiGCN yields an improvement of 5.2% and 6.4% in terms of accuracy on Weibo and PHEME, respectively. The results demonstrate that the spatial structure is more effective than temporal structure. Moreover, it can be observed that KMGCN shows better performance on PHEME dataset than most of text-based models and temporal-based models, indicating that knowledge features can provide complementary information and thus improve performance. STS-NN achieves comparable performance compared to BiGCN on Weibo dataset, but on PHEME dataset it is not as good as BiGCN. The possible reason is as

Method	Features				Weibo				PHEME			
	Text	Temporal	Spatial	KG	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
BERT	✓				0.866	0.867	0.866	0.866	0.821	0.818	0.788	0.803
Transformer	✓				0.772	0.779	0.772	0.775	0.774	0.755	0.648	0.697
EANN	✓				0.866	0.867	0.859	0.862	0.794	0.796	0.791	0.792
QSAN	✓				0.871	0.850	0.870	0.860	0.723	0.631	0.580	0.604
RumorGAN	✓	✓			0.867	0.866	0.866	0.866	0.783	0.785	0.783	0.782
RVNN _{TD}	✓		✓		0.847	0.843	0.847	0.845	0.804	0.803	0.803	0.803
RVNN _{BU}	✓		✓		0.903	0.901	0.897	0.900	0.789	0.788	0.788	0.789
BiGCN	✓		✓		0.919	0.918	0.916	0.913	0.847	0.840	0.834	0.835
KMGCN	✓			✓	0.861	0.864	0.856	0.860	0.812	0.775	0.753	0.764
STS-NN	✓	✓	✓		0.913	0.902	0.898	0.900	0.819	0.816	0.791	0.800
DDGCN	✓	✓	✓	✓	0.948	0.953	0.948	0.950	0.855	0.846	0.841	0.844

Table 3: Results of comparison among different models on Weibo and PHEME Datasets. We run the models five times, and report average results here.

Method	Weibo		PHEME	
	ACC	F1	ACC	F1
Our Model	0.948	0.950	0.855	0.844
-w/o TF	0.929	0.930	0.842	0.833
-w/o dynamic	0.909	0.910	0.841	0.835
-w/o knowledge	0.922	0.925	0.833	0.819
-w/o propagation	0.904	0.905	0.832	0.821

Table 4: Results of comparison among different variants on Weibo and PHEME datasets.

the message propagates, the impacts of the source post information would decay in STS-NN model due to a sequence of gating units. In contrast, we emphasize the impacts of the initial semantic embeddings in the dual-static GCN unit and temporal fusion unit, as the source posts are quite important for the rumor detection task. Additionally, compared to the baselines, we capture more types of semantics involving both temporal and spatial structure, as well as the extra knowledge in source posts and comments, and thus can achieve the best performance.

Ablation Study

We investigate the effects of our proposed components by defining the following variations: (1) **w/o Knowledge**: removing dynamic knowledge GCNs; (2) **w/o Propagation**: removing dynamic propagation GCNs; (3) **w/o TF**: replacing the temporal fusion unit with the *Concatenate* operation. (4) **w/o dynamic**: utilizing static GCNs instead of dynamic GCNs. Specifically, we only use the propagation graph and knowledge graph representations in the final time stage as input, and the model only contains one dual-static GCN unit and one temporal fusion unit.

From Table 4, we can observe that all ablation variants perform worse than the complete DDGCN model on both datasets. Specifically, when removing dynamic propagation GCN layer, the accuracy drops by 4.4% on Weibo dataset and 2.3% on PHEME dataset. It indicates the importance of spatial structure information. When removing dynamic knowledge GCNs, the decrease in accuracy on Weibo

dataset is 2.6% and on PHEME dataset is 2.2%. The replacement of the dynamic graph leads to a accuracy decrease by 3.9% on Weibo and 1.4% on PHEME. And the replacement of the temporal fusion unit also degrades the accuracy by 1.9% on Weibo and 1.3% on PHEME. The results demonstrate the necessity of the temporal information for better performance.

Early Rumor Detection Performance

To construct the early detection task, we set up a series of detection delays and only use the post and comments posted before that delay for detection. We adopt two types of delays: one regards the time elapsed since the source post is released, and the other is set according to the count of received comments. For comparisons, we choose RVNN_{TD}, BiGCN and STS-NN as early detection baselines since they consider temporal information and show good performance in Table 3. Figure 2 shows the performance results at different delays on Weibo and PHEME. It can be observed that DDGCN is able to reach a high accuracy at every early stage after the source post is posted, and almost consistently outperforms others at each time stage.

Conclusion

In this paper, we propose a dual-dynamic graph convolutional networks (DDGCN) to model the spatial structure, temporal structure, external knowledge and text information in one unified framework. DDGCN includes two coupled dynamic GCNs to capture multi-view information in propagation, one of which operates on the evolving propagation graphs and the other operates on the evolving knowledge graphs. Furthermore, we propose the temporal fusing unit to combine the two intermediate graph representations at each time stage, and then passes the fused information to the successor unit in a sequential manner, in order to incrementally aggregate the time-evolving information for classification. Experiments on two public datasets show that DDGCN performs better than a set of strong baselines and supports rumor early detection.

Acknowledgements

This work was supported by the Natural Science Foundation of China (No.61976026, No.61902394) and 111 Project (B18008).

References

- Allcott, H.; and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2): 211–36.
- 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.
- Castillo, C.; Mendoza, M.; and Poblete, B. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, 675–684.
- Chen, L.; Liang, J.; Xie, C.; and Xiao, Y. 2018. Short text entity linking with fine-grained topics. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 457–466.
- Ciampaglia, G. L.; Shiralkar, P.; Rocha, L. M.; Bollen, J.; Menczer, F.; and Flammini, A. 2015. Computational fact checking from knowledge networks. *PloS one*, 10(6): e0128193.
- Cui, L.; Seo, H.; Tabar, M.; Ma, F.; Wang, S.; and Lee, D. 2020. DETERRENT: Knowledge Guided Graph Attention Network for Detecting Healthcare Misinformation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 492–502.
- Deng, S.; Rangwala, H.; and Ning, Y. 2019. Learning dynamic context graphs for predicting social events. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1007–1016.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fionda, V.; and Pirrò, G. 2018. Fact checking via evidence patterns. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 3755–3761.
- Hu, L.; Yang, T.; Zhang, L.; Zhong, W.; Tang, D.; Shi, C.; Duan, N.; and Zhou, M. 2021. Compare to The Knowledge: Graph Neural Fake News Detection with External Knowledge. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 754–763. Association for Computational Linguistics.
- Huang, Q.; Zhou, C.; Wu, J.; Liu, L.; and Wang, B. 2020. Deep spatial-temporal structure learning for rumor detection on Twitter. *Neural Computing and Applications*, 1–11.
- Huang, Q.; Zhou, C.; Wu, J.; Wang, M.; and Wang, B. 2019. Deep structure learning for rumor detection on twitter. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Islam, M. S.; Sarkar, T.; Khan, S. H.; Kamal, A.-H. M.; Hasan, S. M.; Kabir, A.; Yeasmin, D.; Islam, M. A.; Chowdhury, K. I. A.; Anwar, K. S.; et al. 2020. COVID-19-related infodemic and its impact on public health: A global social media analysis. *The American journal of tropical medicine and hygiene*, 103(4): 1621.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kwon, S.; Cha, M.; Jung, K.; Chen, W.; and Wang, Y. 2013. Prominent features of rumor propagation in online social media. In *2013 IEEE 13th international conference on data mining*, 1103–1108. IEEE.
- Liu, X.; Nourbakhsh, A.; Li, Q.; Fang, R.; and Shah, S. 2015. Real-time rumor debunking on twitter. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, 1867–1870.
- Liu, Y.; and Wu, Y.-F. B. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Thirty-second AAAI conference on artificial intelligence*.
- 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, IJCAI'16*, 3818–3824. AAAI Press. ISBN 9781577357704.
- Ma, J.; Gao, W.; Wei, Z.; Lu, Y.; and Wong, K.-F. 2015. 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*, 1751–1754.
- Ma, J.; Gao, W.; and Wong, K.-F. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. Association for Computational Linguistics.
- Ma, J.; Gao, W.; and Wong, K.-F. 2018. Rumor detection on twitter with tree-structured recursive neural networks. Association for Computational Linguistics.
- Ma, J.; Gao, W.; and Wong, K.-F. 2019. Detect rumors on twitter by promoting information campaigns with generative adversarial learning. In *The World Wide Web Conference*, 3049–3055.
- Pan, J. Z.; Pavlova, S.; Li, C.; Li, N.; Li, Y.; and Liu, J. 2018. Content based fake news detection using knowledge graphs. In *International semantic web conference*, 669–683. Springer.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch.
- Ruchansky, N.; Seo, S.; and Liu, Y. 2017. *CSI: A Hybrid Deep Model for Fake News Detection*, 797–806. New York, NY, USA: Association for Computing Machinery. ISBN 9781450349185.
- Shi, B.; and Weninger, T. 2016. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowledge-based systems*, 104: 123–133.

- Song, Y.-Z.; Chen, Y.-S.; Chang, Y.-T.; Weng, S.-Y.; and Shuai, H.-H. 2021. Adversary-Aware Rumor Detection. In *ACL-IJCNLP: Findings*.
- Tian, T.; Liu, Y.; Yang, X.; Lyu, Y.; Zhang, X.; and Fang, B. 2020. QSAN: A Quantum-probability based Signed Attention Network for Explainable False Information Detection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 1445–1454.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- 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.
- Wu, K.; Yang, S.; and Zhu, K. Q. 2015. False rumors detection on sina weibo by propagation structures. In *2015 IEEE 31st international conference on data engineering*, 651–662. IEEE.
- Xia, R.; Xuan, K.; and Yu, J. 2020. A State-independent and Time-evolving Network with Applications to Early Rumor Detection. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 9042–9051.
- Yang, X.; Lyu, Y.; Tian, T.; Liu, Y.; Liu, Y.; and Zhang, X. 2020. Rumor Detection on Social Media with Graph Structured Adversarial Learning. In *IJCAI*, 1417–1423.
- Zhang, H.; Fang, Q.; Qian, S.; and Xu, C. 2019. Multi-modal knowledge-aware event memory network for social media rumor detection. In *Proceedings of the 27th ACM International Conference on Multimedia*, 1942–1951.
- Zubiaga, A.; Liakata, M.; and Procter, R. 2017. Exploiting context for rumour detection in social media. In *International Conference on Social Informatics*, 109–123. Springer.