

A Unified Propagation Forest-based Framework for Fake News Detection

Lingwei Wei^{1,2} and Dou Hu^{1,2} and Yantong Lai^{1,2} and Wei Zhou^{1*} and Songlin Hu^{1,2}

¹ Institute of Information Engineering, Chinese Academy of Sciences

² School of Cyber Security, University of Chinese Academy of Sciences

{weilingwei, hudou, laiyantong, zhouwei, husonglin}@iie.ac.cn

Abstract

Fake news's quick propagation on social media brings severe social ramifications and economic damage. Previous fake news detection usually learn semantic and structural patterns within a single target propagation tree. However, they are usually limited in narrow signals since they do not consider latent information cross other propagation trees. Motivated by a common phenomenon that most fake news is published around a specific hot event/topic, this paper develops a new concept of propagation forest to naturally combine propagation trees in a semantic-aware clustering. We propose a novel Unified Propagation Forest-based framework (UniPF) to fully explore latent correlations between propagation trees to improve fake news detection. Besides, we design a root-induced training strategy, which encourages representations of propagation trees to be closer to their prototypical root nodes. Extensive experiments on four benchmarks consistently suggest the effectiveness and scalability of UniPF.

1 Introduction

Recently, social media platforms have facilitated information dissemination greatly. Nevertheless, they quicken the proliferation of fake news as well due to the lack of authoritative regulators (Zhou et al., 2021). Its extensive dissemination would trigger great panic in society and severely impair the public and individuals (Difonzo and Bordia, 2007; Jin et al., 2017; Jankowski et al., 2020). To keep social media a healthy environment, it is desirable and socially beneficial to detect fake news.

Among previous works on automatic fake news detection, textual news material is utilized by almost all extant studies on fake news detection (Ma et al., 2016; Ruchansky et al., 2017; Popat, 2017; Potthast et al., 2018; Zhou et al., 2020). However, because fake news is purposefully designed to deceive readers by imitating actual news, detecting

and distinguishing them solely from news material is challenging (Afroz et al., 2012; Shu et al., 2020a). Hence, an increasing number of works (Ma et al., 2016; Bian et al., 2020; Hu et al., 2021; Song et al., 2021; Wei et al., 2021) have been devoted to learning potential propagation patterns by investigating relationships among tweets for each news article.

However, most approaches usually consider that propagation trees are independent and ignore latent correlations across different propagation trees, which are supportive for identifying fake news for two-fold reason. 1) Most fake news is usually published deliberately around a specific hot event / topic and then widely disseminated in reality (Frigeri et al., 2014; Nourbakhsh et al., 2015). Hence, potential semantical connections may exist across them. 2) Some spreaders were social bots that are manipulated by a malicious group (Shu et al., 2020a). These deliberate and organized behaviors during propagation may lead to similar structural patterns across two propagation trees. We believe that capturing these vital **semantic** and **structural** characteristics across propagation trees is beneficial to understanding the target propagation tree, accordingly to make more accurate detection.

Yuan et al. (2019) built a global user-tweet heterogeneous graph according to similar participants but ignored semantic relations. Huang et al. (2020) introduced word nodes and constructed a user-tweet-word graph to capture fine-grained semantic relations between source news. However, these fine-grained (e.g., word-level) correlations between news contents may compensate for semantic information to some extent since some fine-grained words may suffer from the polysemy problem (Neelakantan et al., 2014). Given two examples, the real news is *Donald Trump: ...Senate I believe really wants to get something done because Obamacare is dead...Obamacare is absolutely dead*; the fake news is *Donald Trump was pronounced dead this morning following what some are describing as a violent heart attack...*

* Corresponding author.

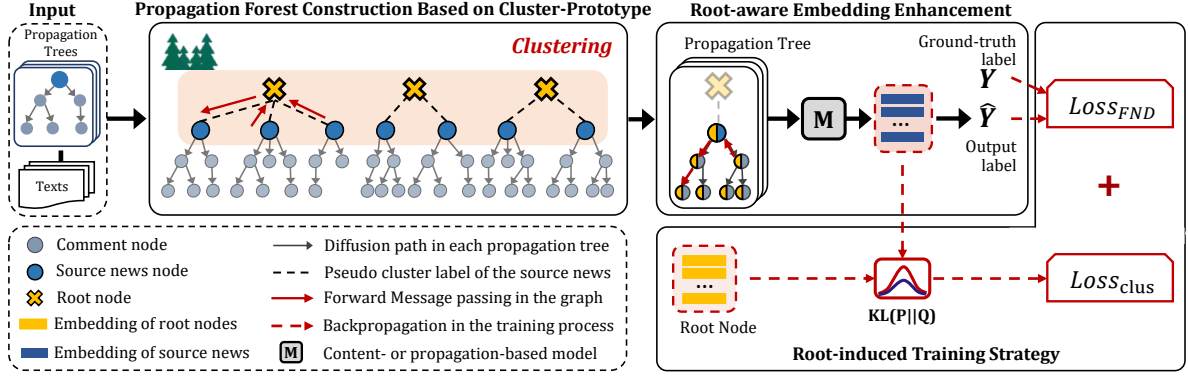


Figure 1: Overall architecture of Unified Propagation Forest-based (UniPF) framework.

Both *Donald Trump* and *dead* were mentioned in both news, but sentence-level semantics are obviously different. The former news shows Donald Trump saying that Obamacare is abolished. But the latter falsely reported Donald Trump was pronounced dead. Thus, there still lacks a unified coarse-grained paradigm that considers effective semantic and structural correlations cross propagation trees simultaneously for detection.

In this paper, we develop a new concept of *Propagation Forest* that combines all propagation trees to explore latent semantic and structural correlations between propagation trees. Under the propagation forest, we propose a general **Unified Propagation Forest-based framework (UniPF)** to enhance target sample’s embedding by exploring latent semantic and structural correlations between propagation trees. Root nodes in the propagation forest are generated according to prototypes (Snell et al., 2017) in a clustering manner. A prototype is a typical embedding for a cluster of related propagation trees. These root nodes can be interpreted as coarse-grained topics or events on social media in reality. Through aggregating neighbor features in the propagation forest, potential patterns implied in similar sub-propagation trees can be captured by performing a high-level transformation of the propagation forest. Besides, a new root-induced training strategy is designed to enhance the quality of the propagation forest by modeling the consistency of generated prototypical root nodes and observed propagation trees. UniPF is of great scalability and can easily be extended in any propagation-based or content-based approaches. We conduct experiments on four benchmarks to assess our framework. The results consistently suggest that UniPF can successfully strengthen fake news detection performance whilst being scalable.

The following are main contributions: 1) We develop a novel concept of *Propagation Forest* that bridges all propagation trees in a semantic-clustering manner. 2) We propose a general UniPF, to enhance fake news detection by fully exploring latent both semantic and structural correlations between similar propagation trees in the generated propagation forest. 3) We devise a root-induced training strategy to guarantee high-quality of prototypical roots in propagation forest. 4) Experiments on four benchmarks consistently demonstrate the scalability and effectiveness of UniPF.

2 Problem Statement

Let $\mathcal{D} = \{(x_i, G_i^{\text{tree}})\}_{i=1}^{|\mathcal{D}|}$ be fake news detection dataset with $|\mathcal{D}|$ samples. Each sample includes a specific source news x_i and its unique propagation tree G_i^{tree} . Text-based fake news detection techniques mainly use x ; while propagation-based techniques use x and G^{tree} . $G_i^{\text{tree}} = (V_i^{\text{tree}}, E_i^{\text{tree}})$ refers to the corresponding propagation tree of i -th source news. $V_i^{\text{tree}} = \{x_i\} \cup \mathcal{C}_i$ is a set of nodes representing the source news x_i and comments \mathcal{C}_i . E_i^{tree} refers to a set of directed edges based on anonymous propagation behaviors, e.g., retweet or comment. Define the embedding of the source news x_i as $\mathbf{x}_i \in \mathbb{R}^{d_0}$, and that of a comment $c_{ij} \in \mathcal{C}_i$ as $\mathbf{c}_{ij} \in \mathbb{R}^{d_0}$, where d_0 is the dimensionality of textual features. Each sample is annotated with a ground-truth label $y_i \in \{0, 1\}$. We formulate the fake news detection problem as a binary classification problem, i.e., each sample can be real ($y_i = 0$) or fake ($y_i = 1$), and learn a classifier f from the labeled set, i.e., $f : \mathcal{D} \rightarrow \mathcal{Y}$.

3 Methodology

This section offers Unified Propagation Forest-based (UniPF) framework to boost fake news detec-




Figure 2: The workflow of constructing a propagation forest based on the input.

tion by examining latent correlations across propagation trees in a semantic-clustering manner. Figure 1 depicts an overview of UniPF framework. It consists of three key components, propagation forest construction, prototype-aware embedding enhancement, and root-induced training strategy.

3.1 Propagation Forest Construction Based on Cluster-Prototype

Given training samples, we perform semantic-aware prototype clustering to generate root nodes, which are *ancestors* for propagation subtrees and are representative for propagation trees with a similar structure as well as semantics. According to pseudo labels of clustering, a propagation forest is developed to combine all propagation trees.

3.1.1 Prototype Generation

Since most fake news is published around a specific event or hot topic and is widely disseminated. it is intuitive to find root nodes via clustering. The goal of this module is to group the entire propagation trees to generate prototypical root nodes of the propagation forest in a clustering manner. Inspired by K-Means (Arthur and Vassilvitskii, 2007), we perform semantic clustering to find several prototypes. The number of clusters K is the only parameter required by the algorithm.

The workflow of constructing a propagation forest based on the input is shown in Figure 2. Specifically, at the preliminary stage, each source news of propagation tree is projected into a unified semantic space. Then, we exploit the clustering tool to process the semantic representations of samples and obtain their pseudo clustering labels to build the propagation forest.

Given the textual embedding $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{|D|}]$ of samples, we find K prototypes and assign propagation trees to a prototype so as to minimize the potential function ϕ . And K prototypes can be

denoted as $\{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_K\}$.

$$\phi = \frac{\sum_{z=1}^Z \|\mathbf{x}_z - \mathbf{m}\|^2}{Z \log(Z + \beta)}, \quad (1)$$

where $\mathbf{m} \in \mathbb{R}^{d_0}$ refers to a prototype. β is a smooth variable that prohibits minor clusters from receiving an excessively high ϕ . Z is the number of data points covered by each prototype. We minimize the function in two distinct principles: 1) summation of squared distance between a target propagation tree and a cluster’s nucleus (i.e., prototype) is minimal; 2) each cluster covers more key-points, namely a higher value of Z . Within clusters, the less variance there is, the more uniform the data points are.

3.1.2 Propagation Forest Construction

Based on the above generated prototypes and propagation trees, we build a propagation forest to combine propagation trees. The propagation forest is formulated as an undirected graph structure, denoted as $G^{\text{Forest}} = (V^{\text{Forest}}, E^{\text{Forest}})$, where V^{Forest} indicates a node set in G^{Forest} and E^{Forest} refers to a undirected edge set in G^{Forest} .

Nodes. There are three node types in propagation forest graph. $V^{\text{Forest}} = V^S \cup V^R \cup V^C$, where 1) source news nodes represent the source news of a propagation tree, denoted as $V^S = \{x_1, \dots, x_N\}$; 2) root nodes represent prototypes, which can be interpreted as latent topics or events on social media, denoted as $V^R = \{m_1, \dots, m_K\}$; 3) comment nodes refer to subsequent retweets given the source news, $V^C = \{\mathcal{C}_i\}_{i=1}^N$, where \mathcal{C}_i is the set of comment nodes of the i -th propagation tree.

Edges. There are two types of edges in the propagation forest graph. 1) For connections between source news nodes and root nodes, we define undirected edges based on pseudo cluster labels, i.e., each propagation tree is connected with the subordinate cluster. The edge weights are defined as the probability that the source news is assigned to the root node. 2) For connections between source

news nodes and comment nodes, and connections between two comment nodes, we define directed edges based on retweet relations. For the i -th propagation tree, if the j -th comment c_{ij} retweeted the source news x_i , there will be an directed edge $x_i \rightarrow c_{ij}$; if the q -th comment c_{iq} retweeted the j -th comment c_{ij} , there will be an directed edge $c_{ij} \rightarrow c_{iq}$. The edge weights are set to 1 if there is a directed edge between two nodes.

3.2 Root-aware Embedding Enhancement

Based on the propagation forest, we perform message passing between root nodes and source news nodes with graph-based transformations to explore latent correlations cross similar propagation trees. Then, we enhance the original message passing in the target propagation tree to boost the understanding of information propagation with the shared semantics and structure from other propagation trees.

3.2.1 Modeling Latent Correlations in the Propagation Forest

We transform the propagation forest using a differentiable message passing method to explore rich correlations between propagation trees. Motivated by graph convolutional networks (Kipf and Welling, 2017), $\mathbf{v}_i^{(1)}$ was calculated for source news node v_i in the first layer by aggregating neighborhood information (i.e., neighbors indicated by the subordinate prototype root nodes) using the transformation,

$$\mathbf{v}_i^{(1)} = \sigma\left(\sum_{j \in \mathcal{N}_i} a_{ij} \mathbf{W}^{(1)} \mathbf{v}_j^{(0)} + a_{ii} \mathbf{W}_0^{(1)} \mathbf{v}_i^{(0)}\right), \quad (2)$$

for $i = 1, \dots, N + K$,

where $\mathbf{v}^{(0)}$ is initialized with \mathbf{m} for root nodes and \mathbf{x} for other nodes. $\mathbf{W}^{(1)}$ and $\mathbf{W}_0^{(1)}$ are trainable parameters. σ is ReLU activation function. \mathcal{N}_i denotes neighbouring indices of node v_i .

Based on the output, another neighborhood-based transformation is applied on source news nodes to integrate shared features from the root nodes. The computations can be defined as,

$$\mathbf{v}_i^{(2)} = \sigma\left(\sum_{j \in \mathcal{N}_i} a_{ij} \mathbf{W}^{(2)} \mathbf{v}_j^{(1)} + a_{ii} \mathbf{W}_0^{(2)} \mathbf{v}_i^{(1)}\right), \quad (3)$$

for $i = 1, \dots, N + K$,

where $\mathbf{W}^{(2)}$ and $\mathbf{W}_0^{(2)}$ are learnable parameters. This stack of transformations effectively accumulates a normalized sum of information from similar propagation trees in the propagation forest.

In this way, latent correlations can be captured by extracting and aggregating effective information from the node’s neighbors. These root nodes can not only transmit semantic and structural features of adjacent samples, but also further integrate features of similar clusters.

3.2.2 Improving Fake News Detection

Next, we exploit the embedding of source news $\mathbf{v}_i^{(2)}$ learned by the above network to enhance the embedding in the target propagation tree. To highlight the impact of the source news during propagation, we make a concatenation and apply a fully connected layer to compute the enhanced embedding of source news \mathbf{x}'_i and comments $\mathbf{C}'_i = [\mathbf{c}'_{i1}, \mathbf{c}'_{i2}, \dots, \mathbf{c}'_{iN}]$.

$$\begin{aligned} \mathbf{x}'_i &= \mathbf{W}_e [\mathbf{x}_i; \mathbf{v}_i^{(2)}] + \mathbf{b}_e, \\ \mathbf{c}'_{ij} &= \mathbf{W}_e [\mathbf{c}_{ij}; \mathbf{v}_i^{(2)}] + \mathbf{b}_e, \end{aligned} \quad (4)$$

where \mathbf{W}_e denotes the transform matrix and \mathbf{b}_e denotes the bias term. The representation of the i -th propagation tree can be further encoded by the existing fake news detection models.

$$\mathbf{u}_i = \text{Model}(\mathbf{x}'_i, \mathbf{C}'_i, G_i^{\text{tree}}), \quad (5)$$

where $\text{Model}(\cdot)$ refers to a base detection model, given textual content and its propagation tree. Then, the label probability $\hat{\mathbf{y}}$ is computed as:

$$\hat{\mathbf{y}}_i = \text{Softmax}(\mathbf{W}_u \mathbf{u}_i + \mathbf{b}_u), \quad (6)$$

where \mathbf{W}_u and \mathbf{b}_u are learnable parameter matrices.

3.3 Root-induced Training Strategy

To guarantee the quality of generated root hubs, we design the root-induced training strategy to constraint the consistency of the representation of the target propagation tree to its prototypical root node in the propagation forest. The strategy can take the “confident” clustering assignments as soft labels, and be assist to guide the optimizing procedure.

3.3.1 Supervised Classification Loss

We minimize the fake news classification loss calculated by the Cross-entropy criterion. That is,

$$\mathcal{L}_{\text{FND}} = -\mathbf{y} \log(\hat{\mathbf{y}}) - (1 - \mathbf{y}) \log(1 - \hat{\mathbf{y}}), \quad (7)$$

where \mathbf{y} indicates the ground-truth label and $\hat{\mathbf{y}}$ represents the prediction label.

3.3.2 Unsupervised Consistency Loss

To improve the embedding of prototypes, inspired by Xie et al. (2016), we examine the similarity between propagation tree representation \mathbf{u} and prototype \mathbf{m} via a Student’s t -distribution.

$$\mathbf{q}_{ik} = \frac{(1 + \|\mathbf{u}_i - \mathbf{m}_k\|^2/\alpha)^{-\frac{\alpha+1}{2}}}{\sum_k^K (1 + \|\mathbf{u}_i - \mathbf{m}_k\|^2/\alpha)^{-\frac{\alpha+1}{2}}}, \quad (8)$$

where α are degrees of freedom of Student’s t -distribution and \mathbf{q}_{ik} can be interpreted as the probability of assigning the i -th sample to prototype m_k , which is a soft assignment. This option can accommodate a variety of scaled clusters while still being computationally efficient.

Then, leveraging propagation tree nodes, we progressively update root nodes by gaining knowledge from high confidence predictions. As a consequence, we outline our goal as a Kullback-Leibler divergence between smooth assignments \mathbf{q} and accessory distribution \mathbf{p} :

$$\mathcal{L}_{\text{Clus}} = \text{KL}(\mathbf{P} \parallel \mathbf{Q}) = \sum_i \sum_k \mathbf{p}_{ik} \log \frac{\mathbf{p}_{ik}}{\mathbf{q}_{ik}}. \quad (9)$$

Considering \mathbf{q} are smooth assignments, using softer probabilistic objectives \mathbf{p} appears more natural and adaptable. The auxiliary distribution \mathbf{p} possesses have three attributes: 1) improve the purity of center clusters; 2) concentrate on highly relevant propagation trees; 3) standardize contribution of each centroid to reduce potential negative risk that larger-scale clusters may obfuscate latent feature space. By increasing \mathbf{q}_i to the second power and then standardizing by cluster size, \mathbf{p}_i is derived:

$$\mathbf{p}_{ik} = \frac{\mathbf{q}_{ik}^2 / \sum_i \mathbf{q}_{ik}}{\sum_{k'} \mathbf{q}_{ik'}^2 / \sum_i \mathbf{q}_{ik'}}. \quad (10)$$

3.3.3 Joint Training Procedure

To recap, we optimize UniPF framework during training through reducing supervised cross-entropy objective of labeled data \mathcal{L}_{FND} and unsupervised consistency objective of unlabeled root nodes $\mathcal{L}_{\text{Clus}}$,

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}_{\text{FND}} + \lambda \mathcal{L}_{\text{Clus}}, \quad (11)$$

where λ is a trade-off hyper-parameter. Θ is all trainable parameters of the model.

4 Experiments

In this section, we experimentally evaluate the performance of our proposed UniPF.

Dataset	Total	# Real News	# Fake News
PolitiFact	314	157	157
GossipCop	5,464	2,732	2,732
Twitter15	712	372	370
Twitter16	410	205	205

Table 1: The statistics of four public datasets.

4.1 Experimental Setup

Datasets We conduct experiments on four public benchmarks. The dataset statistics are shown in Table 1. **PolitiFact** and **GossipCop** are released by FakeNewsNet (Shu et al., 2020a). Samples are collected from *PolitiFact*¹ and *GossipCop*², which are two websites for fact-checking political and celebrity news, respectively. Each sample contains text content of source news and comments, and diffusion relations between anonymous posts. The ground truth label is obtained according to journalists and domain experts. We follow Shu et al. (2019a) and randomly select 75% of news as training data and the remaining as test data. **Twitter15** and **Twitter16**, released by Ma et al. (2017), contain a collection of source news on *Twitter*³ in 2015 (Liu et al., 2015) and 2016 (Ma et al., 2016), respectively. Each sample contains text content of source news and comments, and diffusion relations between comments. Each sample is tagged as non-rumor, false rumor, true rumor, or unverified rumor based on veracity tags in rumor debunking websites. Following Lu and Li (2020), we choose only “true” and “fake” labels as the ground truth. Following Bian et al. (2020); Wei et al. (2021), we conduct 5-fold cross validation for the two datasets.

Comparison Methods *Text*-based fake news detection methods include: **mGRU** (Ma et al., 2016) uses an recurrent neural network to capture sequential features from user comments. **CSI** (Ruchansky et al., 2017) learns the sequential retweet features by employing an LSTM network. *Propagation*-based fake news detection methods include: **GCNFN** (Monti et al., 2019) models the propagation tree as a graph and uses GCN to encode propagation graphs. **BiGCN** (Bian et al., 2020) employs two distinct GCNs to model propagation directed graph and dispersion directed graph, respectively. **RumorGCN** (Hu et al., 2021) learns multi-relational dependencies from the propagation

¹<https://www.politifact.com/>

²<https://www.gossipcop.com/>

³<https://twitter.com/>

Method		PolitiFact				GossipCop			
		Acc	P	R	F1	Acc	P	R	F1
Text-based	mGRU	0.754	0.800	0.666	0.744	0.859	0.845	0.881	0.859
	UniPF-mGRU	0.772[†]	0.846[†]	0.733[†]	0.771[†]	0.905[†]	0.894[†]	0.926[†]	0.905[†]
	CSI	0.734	0.672	0.550	0.688	0.866	0.892	0.840	0.866
	UniPF-CSI	0.760[†]	0.783[†]	0.800[†]	0.754[†]	0.933[†]	0.932[†]	0.933[†]	0.932[†]
Propagation-based	GCNFN	0.856	0.862	0.851	0.849	0.886	0.892	0.881	0.883
	UniPF-GCNFN	0.886[†]	0.902[†]	0.867	0.882[†]	0.958[†]	0.963[†]	0.954[†]	0.958[†]
	BiGCN	0.861	0.865	0.877	0.853	0.959	0.959	0.959	0.958
	UniPF-BiGCN	0.906[†]	0.901[†]	0.906[†]	0.904[†]	0.968[†]	0.968[†]	0.969[†]	0.968[†]
	RumorGCN	0.891	0.901	0.875	0.888	0.968	0.965	0.971	0.968
	UniPF-RumorGCN	0.899[†]	0.911[†]	0.881[†]	0.895[†]	0.968	0.986[†]	0.952	0.969
EBGCN	0.896	0.898	0.909	0.891	0.964	0.966	0.962	0.963	
UniPF-EBGCN	0.911[†]	0.912[†]	0.904	0.909[†]	0.966[†]	0.965	0.973[†]	0.966[†]	

Table 2: Model performance for fake news detection on *PolitiFact* and *GossipCop* datasets. “[†]” marks superior results compared to the corresponding base model, which are significant at level $p < 0.05$ based on t -test.

Method		Twitter15				Twitter16			
		Acc	P	R	F1	Acc	P	R	F1
Text-based	mGRU	0.881	0.888	0.860	0.880	0.893	0.906	0.874	0.893
	UniPF-mGRU	0.923[†]	0.931[†]	0.912[†]	0.923[†]	0.928[†]	0.925[†]	0.928[†]	0.927[†]
	CSI	0.911	0.926	0.900	0.919	0.914	0.912	0.886	0.912
	UniPF-CSI	0.931[†]	0.939[†]	0.921[†]	0.929[†]	0.927[†]	0.954[†]	0.911[†]	0.926[†]
Propagation-based	GCNFN	0.927	0.920	0.936	0.927	0.926	0.927	0.927	0.926
	UniPF-GCNFN	0.938[†]	0.956[†]	0.917	0.938[†]	0.938[†]	0.929[†]	0.951[†]	0.938[†]
	BiGCN	0.942	0.950	0.934	0.942	0.936	0.937	0.936	0.936
	UniPF-BiGCN	0.952[†]	0.959[†]	0.945[†]	0.952[†]	0.955[†]	0.962[†]	0.951[†]	0.955[†]
	RumorGCN	0.952	0.951	0.948	0.952	0.945	0.944	0.944	0.945
	UniPF-RumorGCN	0.959[†]	0.936	0.987[†]	0.959[†]	0.956[†]	0.956[†]	0.948[†]	0.952[†]
EBGCN	0.949	0.947	0.945	0.949	0.943	0.953	0.937	0.944	
UniPF-EBGCN	0.955[†]	0.971[†]	0.941	0.955[†]	0.963[†]	0.951[†]	0.976[†]	0.963[†]	

Table 3: Model performance for fake news detection on *Twitter15* and *Twitter16* datasets. “[†]” marks superior results compared to the corresponding base model, which are significant at level $p < 0.05$ based on t -test.

tree by using relational GCNs. **EBGCN** (Wei et al., 2021), a graph-based model to learn propagation uncertainty in a probability manner, encodes propagation tree with edge-enhanced Bayesian networks.

Evaluation Metrics Since fake news detection is formulated as a classification problem, we evaluate models with four commonly used metrics: Accuracy (Acc), Precision (P), Recall (R), and F1 scores.

Implementation All experiments are conducted on a single Ge-Force RTX 3080Ti. For *PolitiFact* and *GossipCop*, we follow Dou et al. (2021) and consider 300-dimensional word2vec vectors (Mikolov et al., 2013) to initialize word embeddings, which are pretrained on a large corpus with 680k words by spaCy (Honnibal and Montani, 2017). For *Twitter15* and *Twitter16*, following Ma et al. (2018); Bian et al. (2020), we consider TF-

IDF features as a text encoder. We train models via back-propagation and Adam optimizer (Kingma and Ba, 2015). When validation loss stops reducing by 10 epochs, the training will be stopped earlier (Yuan et al., 2007). The maximum value of iterations is 200.

4.2 Results of Fake News Detection

Table 2 and Table 3 show the results of fake news detection on four datasets. Through employing UniPF, the performance of all baselines is improved to different extents, which shows the effectiveness of the framework. Several conclusions can be drawn as follows:

1) Propagation-based methods are often superior to those only using text information. This indicates the importance of learning structural features for fake news detection. As fake news publishers may

Method	PolitiFact				GossipCop			
	Acc	P	R	F1	Acc	P	R	F1
UniPF-BiGCN	0.906	0.901	0.906	0.904	0.968	0.968	0.969	0.968
- w/o SC-G	0.868	0.864	0.867	0.866	0.963	0.959	0.968	0.963
- w/o RT	0.889	0.869	0.913	0.886	0.967	0.968	0.967	0.966
UniPFRS-G-BiGCN	0.871	0.871	0.871	0.869	0.963	0.964	0.964	0.963
- w/o RT	0.868	0.866	0.857	0.865	0.951	0.944	0.956	0.951

Table 4: The ablation study results on *PolitiFact* and *GossipCop* datasets.

Method	Twitter15				Twitter16			
	Acc	P	R	F1	Acc	P	R	F1
UniPF-BiGCN	0.952	0.959	0.945	0.952	0.955	0.962	0.951	0.955
- w/o SC-G	0.951	0.952	0.951	0.950	0.947	0.953	0.942	0.946
- w/o RT	0.949	0.956	0.942	0.949	0.952	0.956	0.948	0.951
UniPFRS-G-BiGCN	0.945	0.950	0.939	0.945	0.951	0.930	0.976	0.951
- w/o RT	0.939	0.945	0.934	0.939	0.949	0.951	0.951	0.950

Table 5: The ablation study results on *Twitter15* and *Twitter16* datasets.

deliberately rub off on the heated topic or disguise themselves by imitating other users, it is challenging to learn informative indicators from noisy texts.

2) In text-based fake news detection methods, UniPF significantly boosts existing models' performance in terms of most metrics. It suggests that fully exploiting deep semantic correlations between similar propagation trees can provide positive auxiliary information to detect fake news. UniPF formulates propagation forest as graph and retrieves semantic-related patterns from textual contents of other propagation trees. These patterns can purified noisy fine-grained correlations obtained by original base model and accordingly learn more easily discernible semantic features for detection.

3) In propagation-based fake news detection models, UniPF framework can also consistently improve the corresponding base models. These base models only capture structural features from the target propagation tree. UniPF effectively explores semantic-structural correlations between propagation trees with the guidance of prototypical root nodes in the propagation forest, thus more discriminative patterns are injected into the representation of the target sample.

4.3 Ablation Study

We perform ablation studies by comparing with: **w/o SC-G** refers to removing the graph-based transformations and learning embedding of root nodes with a simple fully-connected layer. **w/o RT** refers to Root-induced Training strategy is removed in the training process. **UniPFRS-G** performs Randomly-

Sampling to generate root nodes and apply Graph-based transformations to explore deep correlations in the propagation forest.

The results are shown in Table 4 and Table 5. The full UniPF consistently yields the best performance, which shows the effectiveness of the proposed framework for enhancing the detection performance.

1) UniPF w/o SC-G removing graph-based transformations is inferior to UniPF. It shows the embedding of the target sample can be enhanced by explicitly modeling latent correlations in the propagation forest.

2) After removing the root-induced unsupervised clustering loss, results of both UniPF w/o RT and UniPFRS-G w/o RT are reduced. Both variants generate prototypical root nodes in a semantic-clustering or random-sampling way. It implies the efficacy of our training strategy, which ensures the quality of prototypical root nodes. Accordingly, more effective semantic-structural features can be injected to boost the detection performance.

3) UniPFRS-G ignoring semantic-aware clustering achieves poor performance on four datasets, which shows the effectiveness of semantic clustering in our model. The fact reveals more beneficial correlations can be explored between semantically similar propagation trees.

4.4 Parameter Analysis

Figure 3 shows the performance against different numbers of root nodes K in the propagation forest. From results, UniPF is influenced by the number of

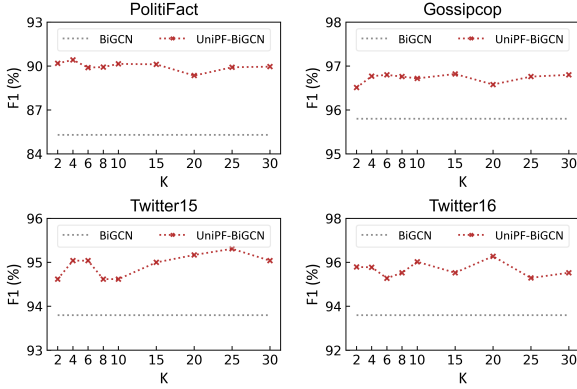


Figure 3: F1 score against different values of K .

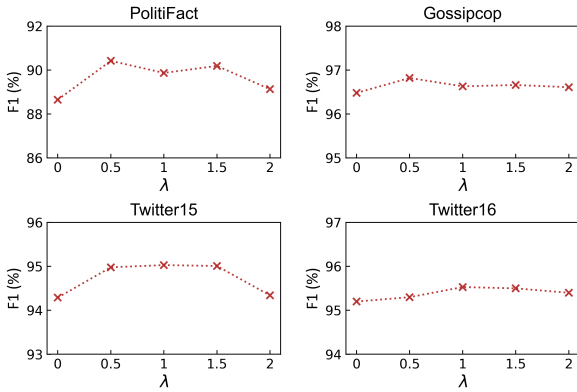


Figure 4: F1 score against different values of λ .

root nodes since these root nodes act as propagation tree’s ancestor in the forest and play a pivot role to incorporate correlations between similar propagation trees. Under too smaller number of root nodes, some redundant or noisy features may be introduced at the same time; while a larger number of root nodes also hurts the performance since assigning more root nodes means propagation trees are independent. The optimal value is 4, 15, 25, and 20 on *PolitiFact*, *GossipCop*, *Twitter15* and *Twitter16* datasets, respectively.

Figure 4 reports the performance against different values of λ , a trade-off hyper-parameter of the unsupervised node-induced clustering loss in the root-induced training strategy. Performance first gets better with the increase of λ . The improvements demonstrate the effectiveness of the root-induced training strategy, which can further ensure high-quality root nodes in the propagation forest. The best setting of 0.5, 0.5, 1, and 1 on *PolitiFact*, *GossipCop*, *Twitter15* and *Twitter16* datasets, respectively. However, performance declines when applying a large value of λ . Because the model with too large λ would pay more attention to se-




Figure 5: Performance with different textual features on *PolitiFact* and *GossipCop* datasets.

Model	Number of Parameters	Average Time
BiGCN	427.5k	2.35s
UniPF-BiGCN	575.2k	2.47s

Table 6: Efficiency analysis on the *PolitiFact* dataset.

semantic clustering based on all propagation trees and cannot effectively exploit structural information of individual propagation trees.

4.5 Feature Analysis

BERT (Devlin et al., 2019) has shown powerful ability to encode textual features and been applied in many NLP tasks (Hu et al., 2022a,b). To investigate the robustness and universality of UniPF, we follow Dou et al. (2021) and evaluate the performance against different feature extractors to explore that given different textual features, whether UniPF can still fully explore latent correlations between propagation trees. Figure 5 shows the fake news detection performance using two different textual features, i.e., Word2Vec (Mikolov et al., 2013) and BERT (Devlin et al., 2019).

From the results, the proposed framework can effectively boost the performance of the baseline model whether BERT or Word2Vec is used to initialize textual features, which reveals the robustness and effectiveness of the proposed UniPF framework. Besides, textual features extracted by BERT are better than Word2Vec. The fact shows contextual features can be captured by the pre-trained model from a large corpus.

4.6 Efficiency Analysis

As shown in Table 6, after employing UniPF, the number of model parameters is 34.5% more than that of BiGCN. The time difference is about 0.12s. Although UniPF would introduce a small number of parameters, the average spent time of each epoch is similar to the base model. We argue it is worthwhile to cost a very small amount of computing resources to provide better results.

5 Related Work

Some works exploit user information to detect fake news such as user credibility (Yuan et al., 2020), user profiles (Shu et al., 2019b) and user social relations (Shu et al., 2019a). However, due to user privacy protection, user’s information is not allowed recorded on social media. This paper focuses on detecting fake news based on *anonymous propagation*. We review related works in two groups: content- and propagation-based detection methods.

5.1 Content-based Detection Methods

Content-based detection approaches investigate the truthfulness of news content by extracting its textual features. Early works (Castillo et al., 2011; Zhao et al., 2015; Popat, 2017; Potthast et al., 2018; Ajao et al., 2019) reply on feature engineering to capture textual attributes. After the emergence of deep learning, many works (Ma et al., 2016; Ruchansky et al., 2017; Shu et al., 2019a; Zhou et al., 2020; Kaliyar et al., 2021) apply various neural networks to automatically learn semantic features from source content and its comments to detect fake news. In reality, fake news is usually published around a specific event/topic, leading to potentially similar semantic characteristics. However, most of them usually extract fine-grained features from words but may suffer from polysemy problem and are not effective to learn shared semantic features. This paper explores potential semantics via generating cluster-prototype and integrating information under the propagation forest, fully exploring shared and identifiable semantic features to boost detection.

Understanding complicated dissemination patterns from propagation trees is also critical since this gives valuable hints into the discovery of fake news. Prior studies (Castillo et al., 2011; Vosoughi et al., 2018; Zhao et al., 2021) focus on several propagation-related criteria such as total number of nodes in a propagation tree, propagation tree depth and breadth. Within a deep learning framework, (Ma et al., 2016; Shu et al., 2020b; Khoo et al., 2020) learn the representation of the target propagation tree with neural networks to capture geometrical spreading patterns. More recently, many mainstream studies (Hu et al., 2019; Dong et al., 2019; Nguyen et al., 2020; Bian et al., 2020; Hu et al., 2021; Silva et al., 2020; Wei et al., 2021, 2022; Silva et al., 2021) regard the propagation tree as a graph structure and apply various graph-based

techniques to learn richer structural features. Although obtaining promising results to some extent, they usually assume propagation trees are independent and ignore vital high-level correlations among propagation. To alleviate this issue, we develop the propagation forest and propose a generic framework to fully explore latent semantic and structural correlations across propagation.

6 Conclusion

In this paper, we investigate deep correlations between propagation trees and propose a generic framework UniPF for improving fake news detection. UniPF builds a propagation forest to naturally combine propagation trees in a semantic-clustering manner. The representation of the target sample is enhanced with shared patterns from similar propagation trees and prototypical root nodes. Besides, a node-induced training strategy is designed for guaranteeing the consistency of the representation of propagation trees and newly prototypical root nodes. Experiments on four datasets consistently prove the scalability and effectiveness of UniPF.

For future work, we will explore more available information (such as images) for perfecting propagation forest to enhance detection.

Acknowledgments

We thank our anonymous reviewers for their helpful comments. This work was supported by the National Natural Science Foundation of China under Grant No.6210071416.

References

- Sadia Afroz, Michael Brennan, and Rachel Greenstadt. 2012. Detecting hoaxes, frauds, and deception in writing style online. In *IEEE Symposium on Security and Privacy*, pages 461–475. IEEE Computer Society.
- Oluwaseun Ajao, Deepayan Bhowmik, and Shahrzad Zargari. 2019. Sentiment aware fake news detection on online social networks. In *ICASSP*, pages 2507–2511. IEEE.
- David Arthur and Sergei Vassilvitskii. 2007. k-means++: the advantages of careful seeding. In *SODA*, pages 1027–1035. SIAM.
- Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *AAAI*, pages 549–556. AAAI Press.

- Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *WWW*, pages 675–684. ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.
- N. Difonzo and P. Bordia. 2007. Rumor, gossip and urban legends. *Diogenes*, 54(1):19–35.
- Ming Dong, Bolong Zheng, Nguyen Quoc Viet Hung, Han Su, and GuoHui Li. 2019. Multiple rumor source detection with graph convolutional networks. In *CIKM*, pages 569–578. ACM.
- Yingtong Dou, Kai Shu, Congying Xia, Philip S. Yu, and Lichao Sun. 2021. User preference-aware fake news detection. In *SIGIR*, pages 2051–2055. ACM.
- Adrien Friggeri, Lada Adamic, Dean Eckles, and Justin Cheng. 2014. Rumor cascades.
- Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*, 7(1).
- Dou Hu, Xiaolong Hou, Lingwei Wei, Lian-Xin Jiang, and Yang Mo. 2022a. MM-DFN: multimodal dynamic fusion network for emotion recognition in conversations. In *ICASSP*, pages 7037–7041. IEEE.
- Dou Hu, Lingwei Wei, Wei Zhou, Xiaoyong Huai, Jizhong Han, and Songlin Hu. 2021. A rumor detection approach based on multi-relational propagation tree. *Journal of Computer Research and Development*, 58(7):1395–1411.
- Dou Hu, Mengyuan Zhou, Xiyang Du, Mengfei Yuan, Jin Zhi, Lian-Xin Jiang, Yang Mo, and Xiaofeng Shi. 2022b. PALI-NLP at semeval-2022 task 4: Discriminative fine-tuning of transformers for patronizing and condescending language detection. In *SemEval@NAACL*, pages 335–343. Association for Computational Linguistics.
- Guoyong Hu, Ye Ding, Shuhan Qi, Xuan Wang, and Qing Liao. 2019. Multi-depth graph convolutional networks for fake news detection. In *NLPCC*, volume 11838 of *Lecture Notes in Computer Science*, pages 698–710. Springer.
- Qi Huang, Junshuai Yu, Jia Wu, and Bin Wang. 2020. Heterogeneous graph attention networks for early detection of rumors on twitter. In *IJCNN*, pages 1–8.
- Jaroslav Jankowski, Piotr Bartkow, Patryk Pazura, and Kamil Bortko. 2020. Evaluation of the costs of delayed campaigns for limiting the spread of negative content, panic and rumours in complex networks. In *ICCS*, volume 12140 of *Lecture Notes in Computer Science*, pages 291–304. Springer.
- Zhiwei Jin, Juan Cao, Han Guo, Yongdong Zhang, Yu Wang, and Jiebo Luo. 2017. Detection and analysis of 2016 US presidential election related rumors on twitter. In *SBP-BRIMS*, volume 10354 of *Lecture Notes in Computer Science*, pages 14–24. Springer.
- Rohit Kumar Kaliyar, Anurag Goswami, and Pratik Narang. 2021. Fakebert: Fake news detection in social media with a bert-based deep learning approach. *Multim. Tools Appl.*, 80(8):11765–11788.
- Ling Min Serena Khoo, Hai Leong Chieu, Zhong Qian, and Jing Jiang. 2020. Interpretable rumor detection in microblogs by attending to user interactions. In *AAAI*, pages 8783–8790. AAAI Press.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*. OpenReview.net.
- Xiaomo Liu, Armineh Nourbakhsh, Quanzhi Li, Rui Fang, and Sameena Shah. 2015. Real-time rumor debunking on twitter. In *CIKM*, pages 1867–1870. ACM.
- Yi-Ju Lu and Cheng-Te Li. 2020. GCAN: graph-aware co-attention networks for explainable fake news detection on social media. In *ACL*, pages 505–514. Association for Computational Linguistics.
- Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J. Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *IJCAI*, pages 3818–3824. IJCAI/AAAI Press.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *ACL*, pages 708–717. Association for Computational Linguistics.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2018. Rumor detection on twitter with tree-structured recursive neural networks. In *ACL*, pages 1980–1989. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *ICLR (Workshop Poster)*.
- Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M. Bronstein. 2019. Fake news detection on social media using geometric deep learning. In *ICLR (Workshop)*.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *EMNLP*, pages 1059–1069.
- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. FANG: leveraging social context for fake news detection using graph representation. In *CIKM*, pages 1165–1174. ACM.

- Armineh Nourbakhsh, Xiaomo Liu, Sameena Shah, Rui Fang, Mohammad Mahdi Ghassemi, and Quanzhi Li. 2015. Newsworthy rumor events: A case study of twitter. In *ICDM Workshops*, pages 27–32. IEEE Computer Society.
- Kashyap Papat. 2017. Assessing the credibility of claims on the web. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 735–739. ACM.
- Martin Potthast, Johannes Kiesel, Kevin Reinartz, Janek Bevendorff, and Benno Stein. 2018. A stylometric inquiry into hyperpartisan and fake news. In *ACL*, pages 231–240.
- Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. CSI: A hybrid deep model for fake news detection. In *CIKM*, pages 797–806. ACM.
- Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019a. defend: Explainable fake news detection. In *KDD*, pages 395–405. ACM.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020a. 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.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, and Huan Liu. 2020b. Hierarchical propagation networks for fake news detection: Investigation and exploitation. In *ICWSM*, pages 626–637.
- Kai Shu, Xinyi Zhou, Suhang Wang, Reza Zafarani, and Huan Liu. 2019b. The role of user profiles for fake news detection. In *International Conference on Advances in Social Networks Analysis and Mining*, pages 436–439. ACM.
- Amila Silva, Yi Han, Ling Luo, Shanika Karunasekera, and Christopher Leckie. 2020. Embedding partial propagation network for fake news early detection. In *CIKM (Workshops)*, volume 2699 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Amila Silva, Yi Han, Ling Luo, Shanika Karunasekera, and Christopher Leckie. 2021. Propagation2vec: Embedding partial propagation networks for explainable fake news early detection. *Inf. Process. Manag.*, 58(5):102618.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. Prototypical networks for few-shot learning. In *NIPS*, pages 4077–4087.
- Chenguang Song, Kai Shu, and Bin Wu. 2021. Temporally evolving graph neural network for fake news detection. *Inf. Process. Manag.*, 58(6):102712.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Lingwei Wei, Dou Hu, Wei Zhou, Xin Wang, and Songlin Hu. 2022. Modeling the uncertainty of information propagation for rumor detection: A neuro-fuzzy approach. *IEEE Transactions on Neural Networks and Learning Systems*.
- Lingwei Wei, Dou Hu, Wei Zhou, Zhaojuan Yue, and Songlin Hu. 2021. Towards propagation uncertainty: Edge-enhanced bayesian graph convolutional networks for rumor detection. In *ACL/IJCNLP*, pages 3845–3854.
- Junyuan Xie, Ross B. Girshick, and Ali Farhadi. 2016. Unsupervised deep embedding for clustering analysis. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 478–487. JMLR.org.
- Chunyuan Yuan, Qianwen Ma, Wei Zhou, Jizhong Han, and Songlin Hu. 2019. Jointly embedding the local and global relations of heterogeneous graph for rumor detection. In *ICDM*, pages 796–805. IEEE.
- Chunyuan Yuan, Qianwen Ma, Wei Zhou, Jizhong Han, and Songlin Hu. 2020. Early detection of fake news by utilizing the credibility of news, publishers, and users based on weakly supervised learning. In *COLING*, pages 5444–5454. International Committee on Computational Linguistics.
- Yao Yuan, Lorenzo Rosasco, and Andrea Caponnetto. 2007. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289 – 315.
- Yuehua Zhao, Jingwei Da, and Jiaqi Yan. 2021. Detecting health misinformation in online health communities: Incorporating behavioral features into machine learning based approaches. *Inf. Process. Manag.*, 58(1):102390.
- Zhe Zhao, Paul Resnick, and Qiaozhu Mei. 2015. Enquiring minds: Early detection of rumors in social media from enquiry posts. In *WWW*, pages 1395–1405. ACM.
- Cheng Zhou, Kai Li, and Yanhong Lu. 2021. Linguistic characteristics and the dissemination of misinformation in social media: The moderating effect of information richness. *Inf. Process. Manag.*, 58(6):102679.
- Xinyi Zhou, Jindi Wu, and Reza Zafarani. 2020. SAFE: similarity-aware multi-modal fake news detection. In *PAKDD*, volume 12085 of *Lecture Notes in Computer Science*, pages 354–367. Springer.