Uncertainty-aware Propagation Structure Reconstruction for Fake News Detection

Lingwei Wei[1,2] and **Dou Hu**[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, zhouwei, husonglin}@iie.ac.cn

Abstract

The widespread of fake news has detrimental societal effects. Recent works model information propagation as graph structure and aggregate structural features from user interactions for fake news detection. However, they usually neglect a broader propagation uncertainty issue, caused by some missing and unreliable interactions during actual spreading, and suffer from learning accurate and diverse structural properties. In this paper, we propose a novel dual graph-based model, Uncertainty-aware Propagation Structure Reconstruction (UPSR) for improving fake news detection. Specifically, after the original propagation modeling, we introduce propagation structure reconstruction to fully explore latent interactions in the actual propagation. We design a novel Gaussian Propagation Estimation to refine the original deterministic node representation by multiple Gaussian distributions and arise latent interactions with KL divergence between distributions in a multi-facet manner. Extensive experiments on two real-world datasets demonstrate the effectiveness and superiority of our model.

1 Introduction

Nowadays, *fake news*[1] has posed detrimental effects on individuals and society. For example, telecommunication towers were burned due to a conspiracy theory linking COVID-19 with 5G technology (Ahmed et al., 2020). To help mitigate the negative effects caused by fake news, it's critical to develop automatic methods to detect fake news.

Existing works generally leverage the user interactions (*e.g.,* retweet) and shared content in a social media conversation thread to detect fake news. The key principle behind such work is that users on social media share opinions, conjectures and evidence for checking fake news. Some studies (Ruchansky et al., 2017; Ma et al., 2016) flatten the conversation in a chronological order to catch linguistic and temporal features from the propagation sequence, which does not make better use of the network properties. Some works (Ma et al., 2018; Kumar and Carley, 2019; Khoo et al., 2020; Ma and Gao, 2020) build the conversation thread with a tree structure to capture the structural patterns from the interactions of information propagation. Driven by the success of graph neural networks (Kipf and Welling, 2017), recent methods (Bian et al., 2020; Hu et al., 2021; Lin et al., 2021) regard the conversation thread as a graph structure and aggregate informative neighbors to learn a good representation for detection.

However, most methods usually assume that the propagation structure is deterministic and complete at some point. In the real world, it is often the case that each sample describes a partial propagation structure that includes some missing and unreliable interactions due to various reasons such as personal privacy protection and profit-driven social bots (Shao et al., 2018). This fact contributes to the propagation uncertainty issue and makes it challenging to discover effective structural patterns for fake news detection. Wei et al. (2021) learned relational bias to alleviate the negative effect of unreliable interactions. But they only focus on explicit interactions between a tweet and its direct retweets. Thus, they still ignore some latent interactions that are not connected but may share similar stances that are useful to debunk false information. These vital but missing latent interactions in the social media conversation thread are also key to driving the propagation uncertainty issue. Thus, how to model the propagation uncertainty issue and learn effective structure-property is a practical research topic to enhance fake news detection.

An intuitive way is to reconstruct the original propagation structure to capture all possible interactions between posting nodes. We argue that, in

* Corresponding author.

[1]We adopts a broad definition, i.e., *fake news is false news where news broadly includes claims, statements, posts, among other types of information* (Zhou and Zafarani, 2020).

Proceedings of the 29th International Conference on Computational Linguistics, pages 2759–2768
October 12–17, 2022.

the propagation, many retweets that subconsciously promote each other (such as similar stances or emotions). Hu et al. (2021); Lin et al. (2021) have shown the positive gains of implicit interactions between sibling retweets from the same tweet. Beyond their assumptions, we make the attempt to investigate more potential interactions of all postings in the propagation structure, not limited to sibling retweets. Besides, previous works (Wei et al., 2021; Hu et al., 2021; Lin et al., 2021) usually measure interactions by learning deterministic embedding of each tweet, which may be insufficient to depict potential interactions accurately and comprehensively for uncertain propagation. Therefore, it is desirable to study potential interactions from multiple underlying facets, which can reflect their fuzzy stances, emotions, and other factors.

In this paper, we investigate a broader propagation uncertainty issue caused by missing and unreliable interactions. Towards this issue, we develop a novel dual graph-based model, named Uncertainty-aware Propagation Structure Reconstruction (UPSR), to adaptively learn accurate and diverse structural properties. Specifically, inspired by Chen et al. (2020), we first utilize deep graph convolutional networks to fully model long-range interactions in the original propagation. Then, instead of directly using deterministic node representations for reconstruction, we design a novel Gaussian Propagation Estimation to sample node representations from multiple Gaussian distributions where the covariance enables the model to reduce noisy interactions. We measure the Kullback-Leibler (KL) divergence between distributions in a multi-facet manner to update the propagation structure. Based on the reconstructed graph, we apply root-aware graph convolutional networks to aggregate features based on the learned latent interactions. UPSR's dual graph structure can not only learn accurate structural information in the original propagation but also capture diverse structural patterns in the reconstructed propagation. Finally, we exploit the dual-graph representation to identify fake news.

We conduct extensive experiments on two real-world public datasets. The experimental results show that UPSR significantly outperforms the state-of-the-art models, indicating the effectiveness for fake news detection. The core contributions of this paper are summarized as follows:

- To handle a broader propagation uncertainty issue caused by missing and unreliable relations, we propose a novel Uncertainty-aware Propagation Structure Reconstruction (UPSR) to learn accurate and diverse structural properties for fake news detection.
- We design a Gaussian Propagation Estimation (GPE) to reconstruct latent propagation structure by measuring KL divergence between different Gaussian distributions of retweets.
- We evaluate the model on two real-world benchmark datasets. Experimental results demonstrate the effectiveness and superiority of the proposed model.

2 Related Work

In the literature, some works (Jiang et al., 2019; Shu et al., 2019b; Mishra, 2020; Nguyen et al., 2020) leverage user characteristics to assist detection. As user information is not allowed recorded in many cases, we mainly focus on detecting fake news based on *text* and *propagation*.

Text-based fake news detection approaches (Mihalcea and Strapparava, 2009) emphasize investigating the truthfulness of news content by extracting its textual features. Early works relied on feature engineering to capture textual characteristics, *e.g.,* topic features (Castillo et al., 2011), writing styles and consistency (Popat, 2017; Potthast et al., 2018). After the emergence of deep learning, many works (Ma et al., 2016; Ruchansky et al., 2017; Karimi and Tang, 2019) apply various neural networks to automatically learn rich semantic or syntactic features from the source news and its retweets to detect fake news.

Propagation-based fake news detection approaches take advantage of the information related to the dissemination of a news article. Many empirical studies (Vosoughi et al., 2018; Jang et al., 2018) have shown that compared to real news, fake news has deeper propagation structures, and reaches a wider audience. Shu et al. (2019a) jointly learned the sequential effect of comments and correlation between source news and the corresponding comments. To capture structural propagation patterns, Ma et al. (2016) constructed a tree-structured neural network to model the propagation structure. Khoo et al. (2020) adopted Transformer (Vaswani et al., 2017) to learn long-distance interactions. Recently, Bian et al. (2020) regarded the propagation as a graph and applied two graph convolutional networks (GCNs) (Kipf and Welling, 2017) to learn structural patterns from two distinct directed graphs. Hu et al. (2021); Lin et al. (2021) further explored multi-relational interactions in the propagation graph. Wei et al. (2021, 2022) focused on the propagation uncertainty and learned robust structural features.

Differences with Existing Models. 1) The aforementioned graph-based models (Bian et al., 2020; Hu et al., 2021) are shallow structure, limiting to explore latent interactions in a deeper propagation. Inspired by Chen et al. (2020), we stack more graph layers to explore long-range interactions in propagation. 2) Most approaches learn latent structural features on statics propagation trees/graphs. They may be disturbed by missing and unreliable behaviors easily, leading to a broader propagation uncertainty issue. This paper designs modules to reconstruct original propagation and explore more latent interactions from multiple facets.

3 Problem Statement

Formally, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a propagation structure, where $\mathcal{V} = \{r, c_1, ..., c_N\}$ is a set of nodes representing the source news r and its retweets $c_1, ..., c_N$. $\mathcal{E}$ refers to a set of explicit interactive behaviors, *e.g.,* retweet. Define the embedding of the source news r as $\mathbf{r} \in \mathbb{R}^{d_0}$, and that of a retweet $\mathbf{c}_i \in \mathbb{R}^{d_0}$, where d_0 is the dimensionality of textual features. Each propagation 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 ($y_i = 1$). Fake news detection task can be seen as to learn a classifier f from the labeled set, *i.e.,* $f : \mathcal{G} \rightarrow y$.

4 The Proposed Model

In this section, we propose a novel dual graph-based model, **UPSR**, to fully model long-range dependencies in the original propagation and explore rich latent dependencies in the corresponding reconstructed propagation.

4.1 Overview

The overview architecture of UPSR is presented in Figure 1. Firstly, given the input text and propagation structure, we apply deep graph convolutions to learn long-range interactions in the original propagation. To better alleviate the propagation uncertainty issue, we design a Gaussian Propagation Estimation to reconstruct the propagation to

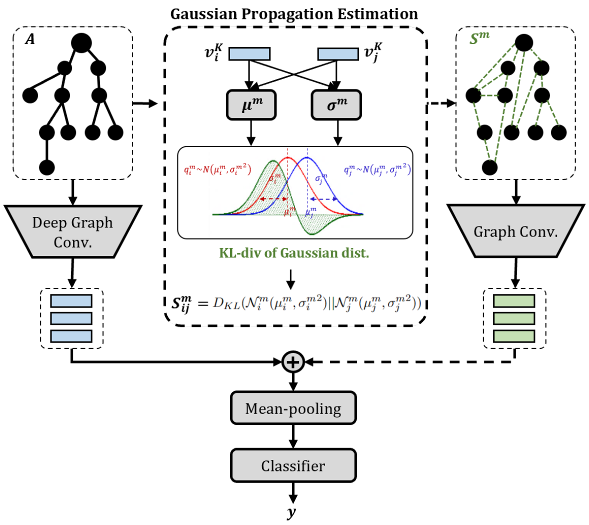

Figure 1: The overall architecture of UPSR.

discover more potential interactions. Then, based on the reconstructed propagation, we further aggregate node features with the guidance of latent connections. Finally, both node representations encoded in the original and latent propagation are concatenated for fake news classification.

4.2 Original Propagation Modeling

Vosoughi et al. (2018) have verified that fake news diffused significantly farther, deeper, and more broadly than the truth. Thus, modeling long-range interactions in the propagation are critical to differentiate fake news and true news. Inspired by (Chen et al., 2020), we develop a deep graph convolutional network to capture long-range interactions in the original propagation.

4.2.1 Graph Construction

First, we construct an undirected graph for each propagation structure to aggregate bi-directional interactions comprehensively. Formally, a propagation structure can be represented as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ denotes a set of tweet nodes including source news r and its retweets $c_1, \ldots, c_n$. $\mathcal{E}$ is a set of propagation behaviors. The edge weights are set to 1 if there is an edge between two nodes, *i.e.,* $\mathbf{A}_{ij} = 1$.

4.2.2 Learning Long-Range Interactions in the Original Propagation Graph

Chen et al. (2020) improved traditional graph convolutional networks by introducing the initial residual connection and an identity mapping to enable stack multiple graph layers, which has shown promising performance on recent downstream applications (Hu et al., 2022). For information propagation, Vosoughi et al. (2018); Jang et al. (2018) have shown that compared to real news, fake news has deeper propagation structures, and reaches a wider audience. Therefore, we apply deep graph convolutional networks (Chen et al., 2020) on an undirected graph to fully capture this kind of long-range dependencies between two nodes in the original propagation.

Given the undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the graph convolution at the k-th layer is defined as Eq. (1). A residual connection to the first layer $\mathbf{V}^{(0)}$ is added to the representation $\tilde{\mathbf{P}}\mathbf{V}^{(k)}$ and an identity mapping $\mathbf{I}$ is added to the weight matrix $\mathbf{W}_t^{(k)}$. $\mathbf{V}^{(0)}$ is initialized with the input embedding, *i.e.,* $\mathbf{V}^{(0)} = [\mathbf{r}, \mathbf{c}_1, ..., \mathbf{c}_N]$.

$$\mathbf{V}^{(k+1)} = \sigma\left(((1-\alpha_k)\tilde{\mathbf{P}}\mathbf{V}^{(k)} + \alpha_k\mathbf{V}^{(0)})((1-\beta_k)\mathbf{I}_n + \beta_k\mathbf{W}_t^{(k)})\right), \tag{1}$$

where $\tilde{\mathbf{P}} = (\mathbf{D}+\mathbf{I})^{-1/2}(\mathbf{A}+\mathbf{I})(\mathbf{D}+\mathbf{I})^{-1/2}$ is the renormalized graph Laplacian matrix (Kipf and Welling, 2017). $\mathbf{A}$ is the original adjacency matrix of $\mathcal{G}$. $\mathbf{D}$ is the diagonal degree matrix, and $\mathbf{I}$ is the identity matrix. α_k, β_k are two hyperparameters. In experiments, $\alpha_k = 0.1$ to make node representations consist of at least a fraction of the input features even if we stack many layers. Let $\beta_k = \log(\frac{\eta}{k}+1)$ to ensure the decay of the weight matrix adaptively increases when stacking more layers. η is also a hyperparameter. $\mathbf{W}_t^{(k)}$ is the k-th weight matrix. σ denotes the activation function.

Based on the above modifications, we can stack many graph layers to capture long distant connections in the original propagation and provide more accurate node representations for the subsequent reconstructed propagation modeling. We denote the number of graph layers as K and final node representations as $\mathbf{V}^{(K)} = \{\mathbf{v}_r^{(K)}, \mathbf{v}_1^{(K)}, ..., \mathbf{v}_N^{(K)}\}$.

4.3 Reconstructed Propagation Modeling

To explore diverse structural patterns, we reconstruct the original propagation for finding more latent interactions and then encode the reconstructed propagation graph for improving detection.

4.3.1 Gaussian Propagation Estimation

We design a Gaussian Propagation Estimation (GPE) to reconstruct the original propagation from multiple facets. Instead of directly measuring the original deterministic embedding of each tweet, the GPE module generate samples stochastic node representations from multiple Gaussian distributions. It can depict potential interactions accurately and comprehensively for uncertain propagation.

Formally, given the deterministic embedding $\mathbf{v}_i^{(K)}$ of each node v_i, the uncertainty-aware node representations is defined as distributional estimation parameterised with estimated mean $\boldsymbol{\mu}_i^m$ and estimated variance $\boldsymbol{\sigma}_i^m$,

$$\begin{aligned} \{\boldsymbol{\mu}_i^1, \boldsymbol{\mu}_i^2, ..., \boldsymbol{\mu}_i^M\} &= g_\theta(\mathbf{v}_i^{(K)}) \\ \{\boldsymbol{\sigma}_i^1, \boldsymbol{\sigma}_i^2, ..., \boldsymbol{\sigma}_i^M\} &= \phi(g'_\theta(\mathbf{v}_i^{(K)})), \end{aligned} \tag{2}$$

where M is a parameter representing the number of facets to estimate uncertain effects of nodes. g_θ and g'_θ are two trainable neural networks such as a multilayer perception (MLP). ϕ is a non-linear activation function. $\{\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2, ..., \boldsymbol{\sigma}^M\}$ indicate the uncertainty of tweets which impacts others in a multi-facet manner. Then, the node representations $\mathbf{Q}^m = \{\mathbf{q}_r^m, \mathbf{q}_1^m, ..., \mathbf{q}_N^m\}$ at the m-th view latent propagation can be sampled from $\mathcal{N}_i^m(\boldsymbol{\mu}_i^m, \boldsymbol{\sigma}_i^{m2})$,

$$\mathbf{q}_i^m = \boldsymbol{\mu}_i^m + \epsilon\boldsymbol{\sigma}_i^m, \epsilon \in \mathcal{N}(\mathbf{0}, \mathbf{I}). \tag{3}$$

Then, GPE measures the latent interactions between nodes with KL divergence between distributions from multiple underlying facets. The edge weight between node v_i and node v_j on the m-th view reconstructed graph is computed as,

$$\mathbf{S}_{ij}^m = D_{KL}(\mathcal{N}_i^m(\boldsymbol{\mu}_i^m, \boldsymbol{\sigma}_i^{m2})||\mathcal{N}_j^m(\mu_j^m, \boldsymbol{\sigma}_j^{m2})). \tag{4}$$

According to the above computations, we can obtain multi-view refined node representations $\{\mathbf{Q}^1, \mathbf{Q}^2, ..., \mathbf{Q}^M\}$ and the corresponding adjacent matrices $\{\mathbf{S}^1, \mathbf{S}^2, ..., \mathbf{S}^M\}$. They enable the model to learn uncertain effects of nodes in multiple reconstructed directed graphs.

4.3.2 Re-Learning Potential Interactions in the Reconstructed Propagation Graph

Based on these reconstructed graphs, we further apply two-layer graph convolutions to capture different potential interactions between two tweets. The message-passing is defined as,

$$\mathbf{U}^m = \sigma\left(\hat{\mathbf{S}}^m(\sigma\left(\hat{\mathbf{S}}^m\mathbf{Q}^m\mathbf{W}_g^{(0)}\right))\mathbf{W}_g^{(1)}\right), \tag{5}$$

where $\hat{\mathbf{S}}$ represents the normalization of adjacency matrix $\mathbf{S}$. $\mathbf{W}_g^{(0)}$ and $\mathbf{W}_g^{(1)}$ are learnable parameter matrices in the first and second graph layer.

Inspired by Bian et al. (2020), we concatenate hidden feature vectors of each node with that of the root node after each graph convolution operation to emphasize the vital role of source news in the propagation. Then, the final representation of nodes in the reconstructed graph is computed as,

$$\mathbf{Z} = \mathbf{W}_z[\mathbf{U}^1; \mathbf{U}^2; ...; \mathbf{U}^M] + \mathbf{b}_z, \quad (6)$$

where $\mathbf{W}_z$ and $\mathbf{b}_z$ are trainable parameters.

Through the above dual graph structure, we can not only learn long-range interactions in the original propagation but also capture potential interactions between uncertain tweets.

We aggregate node representations in the graph to form the graph representations. Given node representations $\mathbf{V}$ in the original propagation and node representations $\mathbf{Z}$ in the reconstructed graph, the graph representation is computed as,

$$\mathbf{O} = meanpooling([\mathbf{V}; \mathbf{Z}]), \quad (7)$$

where $meanpooling(\cdot)$ refers to the mean-pooling aggregating function.

4.4 Fake News Detection and Training

Based on the concatenation of two distinct graph representations, label probabilities of all classes can be defined by a full connection layer and a softmax function, *i.e.,*

$$\hat{\mathbf{y}} = softmax\left(\mathbf{W}_o\mathbf{O} + \mathbf{b}_o\right), \quad (8)$$

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

We optimize the fake news classification loss function calculated by the cross-entropy criterion, i.e.,

$$\mathcal{L} = -\mathbf{y}\log(\hat{\mathbf{y}}) - (1-\mathbf{y})\log(1-\hat{\mathbf{y}}), \quad (9)$$

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

5 Experiments

In this section, we experimentally evaluate the performance of our proposed model for fake news detection.

5.1 Datasets

The dataset statistics are shown in Table 1. **PolitiFact** and **GossipCop** datasets are released by FakeNewsNet (Shu et al., 2020). Samples are collected from *PolitiFact*[2] and *GossipCop*[3], which are two websites for fact-checking political and celebrity news, respectively. We follow the same procedure as Shu et al. (2019a) to split each dataset, *i.e.,* randomly choose 75% of the news as the training data while keeping the rest as the test data.

Dataset	PolitiFact	GossipCop
# News	314	5,464
# True News	157	2,732
# Fake News	157	2,732
# Retweets	40,740	308,798
# Avg. Nodes per Graph	131	58
# Avg. Breadth per Graph	73.62	44.35
# Avg. Depth per Graph	3.75	2.51

Table 1: The statistics of two benchmark datasets.

5.2 Experimental Setups

Since the fake news detection is a classification task, we choose accuracy (Acc), prevision (P), recall (R), and macro-average F1 scores (F1) to measure the performance of each model.

All experiments are conducted on a single GeForce RTX 3080Ti. For the input features of text contents, we follow (Dou et al., 2021) and consider 300-dimensional word2vec vectors (Mikolov et al., 2013), which are pretrained on a large corpus with 680k words by spaCy (Honnibal and Montani, 2017), *i.e.,* $d_0 = 300$. The dimension of hidden vectors is set to 64. We train all models via backpropagation and a wildly used stochastic gradient descent named Adam (Kingma and Ba, 2015). The learning rate is set to 0.001 and 0.0005 for *PolitiFact* and *GossipCop*, respectively. The training process is iterated upon 200 epochs and early stopping (Yuan et al., 2007) is applied when the validation loss stops decreasing by 10 epochs. The final result is the average performance over 5 repeats.

5.3 Comparison Methods

Text-based fake news detection methods include: **mGRU** (Ma et al., 2016) uses an RNN to capture temporal-linguistic patterns recognized from sequences of retweets. **CSI** (Ruchansky et al., 2017) learns the sequential retweet features by employing an LSTM. *Propagation-based* fake news detection methods include: **GCNFN** (Monti et al., 2019) models the propagation structure as a graph

[2] https://www.politifact.com/
[3] https://www.gossipcop.com/

Method	PolitiFact				GossipCop			
	Acc	P	R	F1	Acc	P	R	F1
mGRU (Ma et al., 2016)	0.754	0.800	0.666	0.744	0.859	0.845	0.881	0.859
CSI (Ruchansky et al., 2017)	0.734	0.672	0.550	0.688	0.866	0.892	0.840	0.866
GAT (Velickovic et al., 2018)	0.861	0.848	0.883	0.853	0.958	0.957	0.959	0.957
GCNFN (Monti et al., 2019)	0.856	0.862	0.851	0.849	0.886	0.892	0.881	0.883
PLAN (Khoo et al., 2020)	0.868	0.861	0.879	0.858	0.962	0.960	0.945	0.953
BiGCN (Bian et al., 2020)	0.861	0.865	0.877	0.853	0.959	0.959	0.959	0.958
RumorGCN (Hu et al., 2021)	0.891	0.901	0.875	0.888	0.968	0.965	0.971	0.968
EBGCN (Wei et al., 2021)	0.896	0.898	0.909	0.891	0.964	0.966	0.962	0.963
UPSR	**0.914**	**0.911**	**0.917**	**0.910**	**0.977**	**0.980**	**0.974**	**0.976**

Table 2: Model performance for fake news detection on *PolitiFact* and *GossipCop*. The best result is in bold-face.

and uses GCN to encode the propagation graph. We implemented the model by removing profile information for fair comparison. **GAT** (Velickovic et al., 2018) applies graph attention networks to encode the propagation structure. **PLAN** (Khoo et al., 2020) uses the multi-head attention mechanism to model long-distance interaction in the propagation structure. **BiGCN** (Bian et al., 2020) employs two GCNs to model the propagation graph and dispersion graph. **RumorGCN** (Hu et al., 2021) learns multi-relational dependencies from the propagation by using Relational GCNs. **EBGCN** (Wei et al., 2021), a graph-based model, focuses on the uncertainty issue in the propagation structure from a probability perspective.

5.4 Fake News Detection Results

The overall performance for fake news detection is reported in Table 2. From them, we we have the following key observations:

1) Text-based methods achieve inferior performance than propagation-based methods. It indicates that propagation patterns are more beneficial to detect fake news since fake news publishers always deliberately distort the text content of news. 2) PLAN captures long-range interactions in the propagation sequence with attention modules and obtains moderate results, even outperforming some shallow graph-based models. However, they still could not effectively distill latent interactions hidden in the propagation sequence and thus obtain limited performance. 3) EBGCN and RumorGCN achieve sub-optimal performance on PolitiFact and GossipCop, respectively. It makes sense as RumorGCN considers potential interactions from sibling nodes; while EBGCN explores robust interactions in an adjusted propagation tree, which can provide more effective structural information for detection. Nevertheless, their shallow networks make it hard to model long-distance interactions in the propagation, and thus they cannot be adaptive for news that has a deeper propagation structure. 4) Our UPSR yields consistently better performance than all the baselines on both datasets. The benefit mainly comes in two-fold. First, deep graph convolutions enable the model to focus on long-range interactions in the original propagation modeling. Second, UPSR further encodes the reconstructed propagation based on uncertainty-aware node representations, which can effectively capture more potential interactions between retweets and learn diverse structural patterns for detection.

Methods	PolitiFact		GossipCop	
	Acc	F1	Acc	F1
UPSR	**0.914**	**0.910**	**0.977**	**0.976**
- w/o Root	0.891	0.886	0.974	0.973
- w/o GPE	0.904	0.894	0.972	0.961
- w/o OPM	0.828	0.817	0.975	0.974
- w/o RPM	0.873	0.867	0.962	0.961
UPSR$_{\text{GCN}}$	0.891	0.886	0.972	0.971
UPSR$_{\text{GAT}}$	0.899	0.894	0.973	0.973
UPSR$_{\text{BiGCN}}$	0.886	0.880	0.974	0.973

Table 3: Results of ablation study and component analysis. The best result is in bold-face.

6 Discussion

In this section, we conduct more experiemtns to further understand the performance of UPSR.

6.1 Ablation Study

We conduct an ablation study to evaluate key components in UPSR. 1) **w/o Root** indicates that encoding the reconstructed propagation graph does not explicitly consider the influence of source news. 2) **w/o GPE** removes Gaussian Propagation Estimation module and measures cosine-similarity between two node embedding. 3) **w/o OPM** refers to removing the original propagation modeling and directly reconstructing the propagation according to input textual features. 4) **w/o RPM** is removing the whole reconstructed propagation modeling.

The results of the ablation study are shown in the first block of Table 3. The full model yields the best performance in terms of accuracy and F1 score. 1) Without the consideration of source news influence in the reconstructed propagation modeling, the performance of **w/o Root** slightly reduces on both datasets, showing the vital role of source news in the propagation. 2) **w/o GPE** is obviously inferior to the full model, verifying that estimating propagation structure with multiple facets can successfully adapt to the uncertain effect of retweets and enable to derive accurate potential interactions. 3) When removing the complete reconstructed propagation modeling, **w/o RPM** obtains the inferior performance in terms of two evaluation metrics, which proves the effectiveness of the propagation reconstruction. 4) After removing the original propagation modeling, the performance of **w/o OPM** also drops significantly. This is intuitive since learning from explicit interactions between retweets in the original propagation could lead to relatively comprehensive representations, which enables GPE to explore more effective interactions.

6.2 Comparison with Different Original Propagation Modeling Modules

We further replace the deep graph convolutional network in the original propagation modeling with the following alternatives. 1) **UPSR**$_{\text{GCN}}$ adopts vanilla two-layer GCNs (Kipf and Welling, 2017) to model the original propagation. 2) **UPSR**$_{\text{GAT}}$ replaces with vanilla two-layer GATs (Velickovic et al., 2018). 3) **UPSR**$_{\text{BiGCN}}$ follows (Bian et al., 2020) to apply bi-directional GCNs .

The results are reported in the second block of Table 3. The degradation performance of these variants indicates the superiority of our model, which can capture long-range interactions in the propagation by stacking multiple graph convolutions. Besides, **UPSR** and its variants **UPSR**$_{\text{GCN}}$, **UPSR**$_{\text{GAT}}$, **UPSR**$_{\text{BiGCN}}$ consistently outperform the corresponding single graph models on both datasets. The reason is that the dual graph framework can not only learn interactions in the original propagation but also capture potential interactions between uncertain tweets.

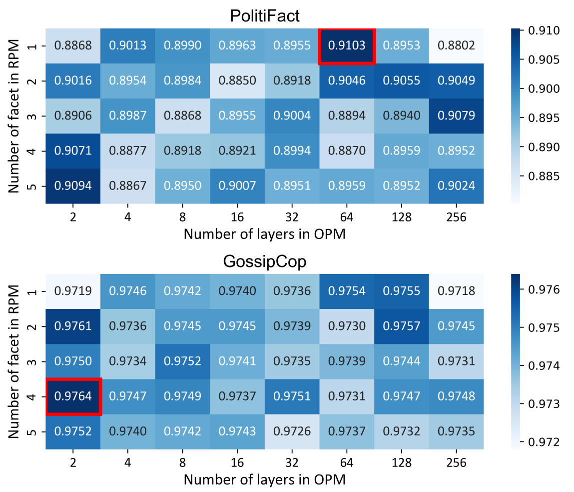

Figure 2: F1 scores against different hyperparameters.

6.3 Parameter Analysis

Figure 2 explores the performance of UPSR against two vital parameters, *i.e.,* different numbers of layers in the original propagation modeling (OPM), and different numbers of facets in the reconstructed propagation modeling (RPM).

Effect of Graph Layers in Original Propagation Modeling. To investigate whether our model can benefit from the multi-layer propagation in the original propagation modeling, we vary the number of graph convolutional layers in the range of {2, 4, 8, 16, 32, 64, 128, 256}. The best setting is 64 and 2 on PolitiFact and GossipCop, respectively. Propagation structures are deeper on PolitiFact and thus more graph layers are needed to capture long-range interactions between nodes. The continual increase of the layer number even harms the performance. This might be caused by the overfitting issue.

Effect of Number of Facets in Reconstructed Propagation Modeling. To investigate whether our model can benefit from the multi-facet estimation for uncertainty, we vary the number of facets in the range of {1, 2, 3, 4, 5}. The optimal setting is 1 and 4 on PolitiFact and GossipCop datasets, respectively. These results indicate that estimating nodes from multiple facets is more profitable for detecting celebrity-related fake news, which can boost to capture latent interactions between two nodes sufficiently. Besides, dependencies between retweets under celebrity news may be more com-

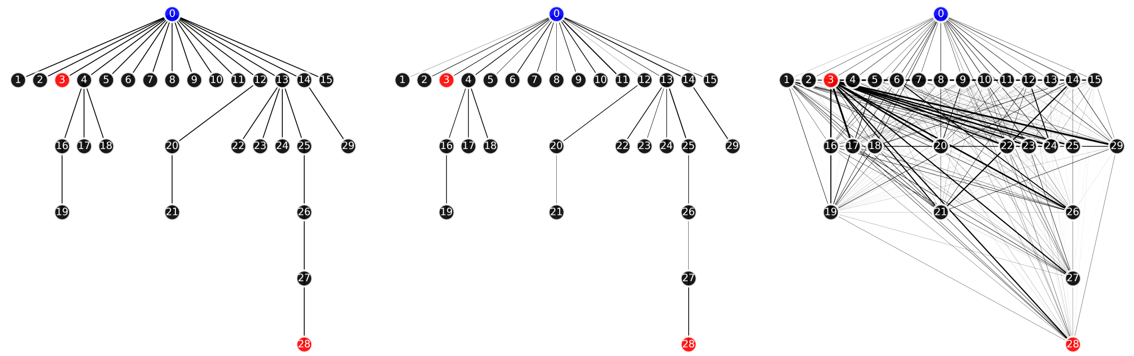

(a) Original Propagation (b) Adjusted by EBGCN (c) Reconstructed by UPSR

Figure 3: A case study of fake news on *PolitiFact*, which is missed by BiGCN and EBGCN but detected by UPSR. Node 0 refers to the source news and other nodes are its retweets. The breadth of the propagation is 15 and the depth of the propagation is 5. The edge width represents the weight of interactions.

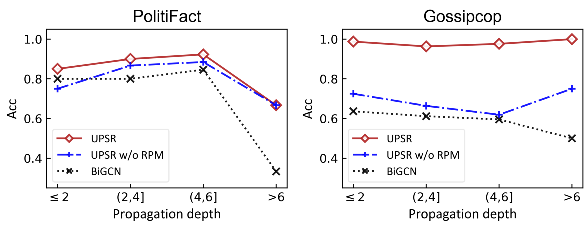

Figure 4: Performance on propagation structures with different depths. Y-axis refers to the accuracy score.

plex and more facets need to be considered.

6.4 Propagation Depth Analysis

Figure 4 shows performance on propagation structures with different depths. From the figure, the performance of BiGCN for detecting deeper propagation clearly decreases on both datasets. This reveals that fake news detection is more challenging with the deeper propagation which usually reflects vital potential interactions between users. Compared with BiGCN, UPSR and its variant obtain better performance in recognizing deeper propagation. This indicates that the original propagation modeling can effectively capture longer-range interactions in the original propagation for fake news detection. Moreover, UPSR achieves a considerable improvement over almost any range of propagation depth. We speculate, through estimating uncertain effects of retweets to reconstructing the original propagation, UPSR can further capture more potential interactions between two nodes and learn better representations for detection. Thus, UPSR is not sensitive to propagation depth and can be adaptive for both shallow and deep propagation.

6.5 Case Study

Figure 3 visualizes a propagation structure of a piece of *fake* news from PolitiFact. The news is misclassified by BiGCN and EBGCN but is detected by our model successfully.

Previous shallow graph networks (*e.g.,* BiGCN, EBGCN) would ignore the distant connections such as the interaction between node 3 and 28 and can only capture local structural propagation information. Through reconstructing the original propagation, UPSR alleviates this issue to some extent and aggregates more effective information in the graph via reconstructed edges between two distant nodes. Besides, compared with Figure 3(b) and 3(c), EBGCN dealt with noisy edges by adaptively adjusting weights of explicit edges. However, they solely focus on explicit edges and limit the message-passing in the graph. Different from their model, UPSR not only is robust to these noisy edges but also captures more valuable potential interactions between nodes to improve detection.

7 Conclusion

This paper has studied a broader propagation uncertainty issue in fake news detection. We propose a novel Uncertainty-aware Propagation Structure Reconstruction (UPSR) to jointly model long-range and potential interactions in the uncertain propagation. Gaussian Propagation Estimation (GPE) is developed to reconstruct latent propagation by adapting the inherent uncertain effect of retweets in the propagation. Experiments conducted on two real-world benchmarks have shown that UPSR outperforms recent detection methods.

In the future, we will focus on improving the detection performance of our model in scenarios where training propagation data is limited.

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

Wasim Ahmed, Josep Vidal-Alaball, Joseph Downing, Francesc López Seguí, et al. 2020. Covid-19 and the 5g conspiracy theory: social network analysis of twitter data. *Journal of medical internet research*, 22(5):e19458.

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.

Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *WWW*, pages 675–684. ACM.

Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. 2020. Simple and deep graph convolutional networks. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 1725–1735. PMLR.

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.

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. 2022. 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.

S. Mo Jang, Tieming Geng, Jo-Yun Queenie Li, Ruofan Xia, Chin-Tser Huang, Hwalbin Kim, and Jijun Tang. 2018. A computational approach for examining the roots and spreading patterns of fake news: Evolution tree analysis. *Comput. Hum. Behav.*, 84:103–113.

Shengyi Jiang, Xiaoting Chen, Liming Zhang, Sutong Chen, and Haonan Liu. 2019. User-characteristic enhanced model for fake news detection in social media. In *NLPCC (1)*, volume 11838 of *Lecture Notes in Computer Science*, pages 634–646. Springer.

Hamid Karimi and Jiliang Tang. 2019. Learning hierarchical discourse-level structure for fake news detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3432–3442. Association for Computational Linguistics.

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 (Poster)*. OpenReview.net.

Sumeet Kumar and Kathleen M. Carley. 2019. Tree lstms with convolution units to predict stance and rumor veracity in social media conversations. In *ACL (1)*, pages 5047–5058.

Hongzhan Lin, Jing Ma, Mingfei Cheng, Zhiwei Yang, Liangliang Chen, and Guang Chen. 2021. Rumor detection on twitter with claim-guided hierarchical graph attention networks. In *EMNLP (1)*, pages 10035–10047. Association for Computational Linguistics.

Jing Ma and Wei Gao. 2020. Debunking rumors on twitter with tree transformer. In *COLING*, pages 5455–5466. International Committee on 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. 2018. Rumor detection on twitter with tree-structured recursive neural networks. In *ACL (1)*, pages 1980–1989. Association for Computational Linguistics.

Rada Mihalcea and Carlo Strapparava. 2009. The lie detector: Explorations in the automatic recognition of deceptive language. In *ACL/IJCNLP (2)*, pages 309–312. The Association for Computer Linguistics.

Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *ICLR (Workshop Poster)*.

Rahul Mishra. 2020. Fake news detection using higher-order user to user mutual-attention progression in propagation paths. In *CVPR Workshops*, pages 2775–2783. Computer Vision Foundation / IEEE.

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)*.

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.

Kashyap Popat. 2017. Assessing the credibility of claims on the web. In *WWW (Companion Volume)*, 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 (1)*, pages 231–240. Association for Computational Linguistics.

Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. CSI: A hybrid deep model for fake news detection. In *CIKM*, pages 797–806. ACM.

Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. The spread of low-credibility content by social bots. *Nature communications*, 9(1):1–9.

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. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data*, 8(3):171–188.

Kai Shu, Xinyi Zhou, Suhang Wang, Reza Zafarani, and Huan Liu. 2019b. The role of user profiles for fake news detection. In *ASONAM*, pages 436–439. ACM.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.

Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *ICLR (Poster)*. OpenReview.net.

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 (1)*, pages 3845–3854. Association for Computational Linguistics.

Yao Yuan, Lorenzo Rosasco, and Andrea Caponnetto. 2007. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289 – 315.

Xinyi Zhou and Reza Zafarani. 2020. A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Comput. Surv.*, 53(5):109:1–109:40.