MPPFND: A Dataset and Analysis of Detecting Fake News with Multi-Platform Propagation

Congyuan Zhao[1,2], Lingwei Wei[1*], Ziming Qin[1,2], Wei Zhou[1], Yunya Song[3], Songlin Hu[1,2*]
[1]Institute of Information Engineering, Chinese Academy of Sciences
[2]School of Cyber Security, University of Chinese Academy of Sciences
[3]Division of Emerging Interdisciplinary Areas, Hong Kong University of Science and Technology
{zhaocongyuan, weilingwei, zhouwei, husonglin}@iie.ac.cn, qinziming21@mails.ucas.ac.cn, yunyasong@ust.hk

Abstract

Fake news spreads widely on social media, leading to numerous negative effects. Most existing detection algorithms focus on analyzing news content and social context to detect fake news. However, these approaches typically detect fake news based on specific platforms, ignoring differences in propagation characteristics across platforms. In this paper, we introduce the MPPFND dataset, which captures propagation structures across multiple platforms. We also describe the commenting and propagation characteristics of different platforms to show that their social contexts have distinct features. We propose a multi-platform fake news detection model (APSL) that uses graph neural networks to extract social context features from various platforms. Experiments show that accounting for cross-platform propagation differences improves fake news detection performance.

Introduction

Fake news is defined as fake news that is deliberately created and disseminated, typically with the intent to mislead(Pennycook et al., 2018). The widespread propagation of fake news on social platforms has had significant negative impacts on society. For instance, during the COVID-19 pandemic, a plethora of fake news caused widespread panic among the public, and incorrect treatment methods adversely affected many lives (Aïmeur et al., 2023; Rocha et al., 2021). Many fact-checking sites[1] are dedicated to having fact-checkers collect accurate information and compare it with news appearing on the internet to identify fake news. However, manual verification is too inefficient to keep up with the vast amount of fake news. Therefore, there is an urgent need for an effective model to automatically detect fake news and curb the spread of fake news.

Currently, many studies mainly use content features and propagation structure to detect fake news: Content feature based models (Kaliyar et al., 2021; Wang et al., 2018; B. Ma et al., 2017; Jin, Cao, Guo, Zhang, Wang, & Luo, 2017) aim to detect fake news through language differences between news contents. Recent studies have explored more fine-grained analyses of news content, such as (Nan et al., 2021; Li et al., 2024) that categorize news into different event domains for fake news detection. Some works incorporate user interactions such as retweets and comments during the propagation. These research model the spread of news as propagation sequence (Liu & Wu, 2018; Khoo et al., 2020) or graph (Silva et al., 2021; Yin et al., 2024,?; Bian et al., 2020; Cui & Jia, 2024) to learn sequential or topological structures for better detection.

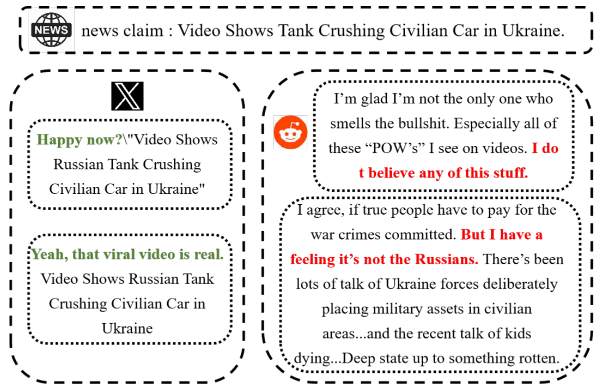

Figure 1: An example of real-world news spreading across multiple social media platforms, where we can observe a significant disparity in how the credibility of the news is assessed by the two platforms.

However, most works usually detect fake news based on the specific platform and obtain limited performance. Some works Voorveld et al. (2018) have shown the difference between various social media platforms. In the real world, a piece of fake news often spreads across multiple platforms through mechanisms such as retweets, that often provide richer and complementary features for detection 1. Therefore, a more detailed analysis and modeling of the propagation structure is necessary.

In this paper, to fill the gap in this area, we aim to conduct an in-depth analysis of the characteristics of multi-platform propagation and explore that these differentiated features aid in detecting fake news. We collect a new fake news dataset that gathers the propagation structure of a news claim across various social media platforms. We list existing popular fake news datasets below and compare them with our dataset repository in Table 1. Based on our constructed dataset, we further analyze the propagation characteristics across platforms, revealing significant variations in both commonalities and differences for fake news detection.

By analyzing our dataset, we find that user engagement ef-

[1]Such as politifact.com, snopes.com etc.
*Corresponding author.

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

Dataset	# Feature			# Social Platform Source	# Language
	Content	Single-platform Propagation	Multi-platform Propagation		
FakeNewsNet (Shu et al., 2020)	√	√	×	X	EN
Weibo16 (Jin, Cao, Guo, Zhang, & Luo, 2017)	√	√	×	Weibo	CH
Weibo21 (Nan et al., 2021)	√	√	×	Weibo	CH
Twitter16 (J. Ma et al., 2016)	√	√	×	X	EN
MCFEND (Li et al., 2024)	√	√	×	Weibo	CH
MPPFND (Ours)	√	√	√	X\Youtube\Reddit	EN

Table 1: Comparison of our dataset with other popular fake news detection datasets.

fectively differentiates between the authenticity of news on YouTube and X, whereas this feature does not show a significant relationship with news authenticity on Reddit. Additionally, the method of using comment sentiment to distinguish news authenticity is applicable to X(Twitter) and Reddit, but shows only moderate effectiveness on YouTube. Thus, it is evident that differential modeling across various platforms is necessary for detecting fake news.

To demonstrate that the platform-specific characteristics we identified contribute to the detection of fake news, we propose a multi-platform detection model that leverages multiple graph neural networks to extract features from propagation structures across different platforms. Our experiments show that the differences in propagation across platforms contribute to improving the effectiveness of fake news detection models.

The contribution of this paper lies in: 1) We construct and release a new MPPFND dataset for fake news detection with multi-platform propagation. It contains 3,500+ claims, along with 440,000+ user engagements across 3 mainstream social platforms.[2] 2) We contribute empirical insights into propagation patterns across different social platforms for fake news detection. Our findings reveal platform-specific and platform-shared propagation patterns on news content, comment style, and user engagement for fake news detection. 3) We evaluate existing content- and propagation-based fake news detection methods for multi-platform fake news detection and further develop an adaptive propagation structure learning network (APSL) to capture platform-adaptive structural features from propagation across platforms for better detection.

Related Work

Fake News Detection

Previous works on fake news detection can be broadly categorized into content-based and propagation-based methods.

Content-Based Methods. Content-based detection identifies recurring patterns and distinctive features in fake news by analyzing news articles (Huang et al., 2023). This includes examining linguistic traits (Choudhary & Arora, 2021), unique writing styles (Yu et al., 2017; Kaliyar et al., 2021), and emotional nuances (Ajao et al., 2019; Giachanou et al., 2019; Zhang et al., 2021). Recent studies have also focused on domain-specific features (Silva et al., 2021; Nan et al., 2021). Additionally, some approaches extend the analysis to multi-modal elements, such as images (Z. Chen et al., 2023; Z. Ma et al., 2024).

Propagation-Based Methods. The spread of fake news on social media has garnered significant attention. Some studies model the propagation process as a sequence (Liu & Wu, 2018; Lin et al., 2023), often using attention mechanisms to capture long-range dependencies (Khoo et al., 2020; Lu & Li, 2020; Ran & Jia, 2023). Furthermore, fake news propagation is sometimes modeled as a tree, making it suitable for graph classification techniques (Tian et al., 2022; Wei et al., 2021; D. Hu et al., 2021; Sheng et al., 2022; Zong et al., 2024; Lu & Li, 2020; Mehta et al., 2022). Some studies explore the role of users in propagation (Dou et al., 2021; Wei et al., 2022), while others focus on the structure of the propagation tree (Cui & Jia, 2024).

Fake News Dataset

Content-centric Datasets initially include news metadata, textual content, and supervised labels, forming the basis for early fake news detection tasks (Nørregaard et al., 2019; Zellers et al., 2019). To reduce bias from semantic and lexical content, factual evidence and multi-modal data are integrated, enhancing the detection framework (Jiang et al., 2020; X. Hu et al., 2023; Yao et al., 2023).

Propagation-based Datasets capture key propagation data, such as retweets, comments, and likes, aiding fake news detection by analyzing propagation patterns (Shu et al., 2020; J. Ma et al., 2016; Jin, Cao, Guo, Zhang, & Luo, 2017). Nan et al. (2021) constructs the large Chinese dataset Weibo-21, adding domain tags to the news metadata and incorporating propagation data. To address the limitations of single-source datasets, Li et al. (2024) develops the multi-source benchmark MCFEND for Chinese fake news detection, using data from social platforms, messaging apps, and traditional news.

However, previous works primarily focus on propagation patterns within a single platform. In contrast, we introduce MPPFND, a pioneering multi-platform dataset that offers richer social context, laying the foundation for deeper investigation into fake news propagation characteristics.

[2]We release our dataset and code at https://github.com/Zhaocongyuan/MPPFND-Dataset.

Fake News Detection with Multi-Platform Propagation

Multi-platform fake news detection aims to detect fake news using propagation structure information from multiple social media platforms. To learn the intricate structure and inherent characteristics between platforms, we first construct a fake news dataset encompassing the propagation structures across multiple platforms (short for **MPPFND**). Subsequently, we conduct a thorough empirical analysis based on this dataset to learn platform-specific and platform-shared propagation patterns for fake news detection.

Dataset Construction

In this subsection, we describe the process of data collection for our dataset.

Data Collection We utilize fact-checking websites[3][4] to obtain news contents and obtain the ground-truth label for fake news based on annotations provided by expert fact checkers. Then, to achieve multi-platform social engagements, we create search queries based on the headlines to retrieve posts on three popular social platforms (i.e., Youtube, X, Reddit). X is a short-text social media platform, which are featured by rapid information diffusion. Reddit are long-text social media platform that can contain rich information. YouTube, as a video-sharing platform, hosts discussions that often include richer video evidence. Additionally, we further fetch the user engagements including comments, reposts and likes towards these posts on the above social platforms. MPPFND encompasses 4965 data samples spanning two semantic topics and three platforms.

Label Mapping Strategy The fact-checking platforms provide a variety of labels. Following Nan et al. (2021), we convert this into a binary classification problem by using the following label mapping strategy. For PolitiFact, which has eight original labels in total, we designate *mostly-true*, *half-true*, and *true* as *True*, while the remaining labels are marked as *False*. For Snopes, which has twelve original labels, we mark *mostly-true* and *true* as *True*, and the rest as *False*.

Analysis of Propagation Properties across Multiple Social Platforms

To explore the differences across different social media platforms, we conduct in-depth analyses based on our MPPFND dataset.

Analysis of Propagation Statistics To investigate propagation structures across platforms, we analyze the proportion of news claims on multiple social media platforms, as shown in Table 2. **The spread of fake and true information shows distinct platform-specific and common traits.** For example, fake news engagement is higher on YouTube and X but lower on Reddit. Fake news also spreads across more platforms than true news, indicating broader reach. **Skewed distributions**, such as fewer nodes on X, highlight potential propagation bias in low-resource platforms for multi-platform fake news detection. Figure 2c shows the distribution of news across single and multiple platforms. **Incomplete cross-platform propagation data** underscores the need for adaptive detection models. **True news is more common on single platforms, while fake news spreads more across multiple platforms.** This suggests fake news has stronger cross-platform diffusion capability, likely due to its sensational or controversial nature, prompting more frequent sharing.

Analysis of Comment Style To further explore comment styles across different platforms, we first analyze the length of comments to determine whether users have fully expressed their opinions. As shown in Figure 2b, we observe that: 1) On X and YouTube, comments on fake claims are shorter

Domain	#Total Nodes	Tree-Width	#Avg. Nodes
Youtube	450,711	774	273.99
Youtube (Fake)	404,348	774	304.48
Youtube (True)	46,363	600	146.26
X	11,059	96	6.00
X (Fake)	9,458	96	6.50
X (True)	1,601	55	4.13
Reddit	35,258	300	323.47
Reddit (Fake)	29,297	300	321.95
Reddit (True)	5,961	200	331.17

Table 2: Summary of data for various domains including total nodes, tree-width, and average nodes per graph.

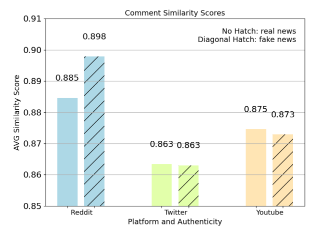

(a) Comment similarity with the claim of each platform in our dataset.

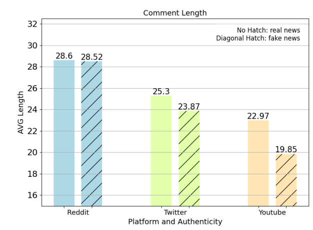

(b) Analysis of comment length on each social platform in our dataset.

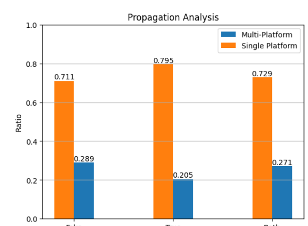

(c) The proportion of news claim across multiple or on a specific platform.

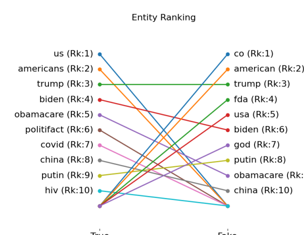

(d) The Top 10 entities most frequently appearing in news and its comments.

Figure 2: (a) Comment similarity, (b) Comment length, (c) News duplication, (d) Entity analysis.

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

[4]https://www.snopes.com/

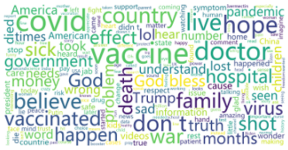

(a) Comments on YouTube

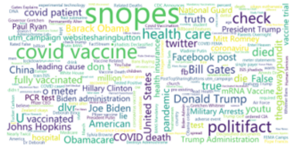

(b) Comments on YouTube

(c) Comments on Reddit

(d) Comments on Reddit

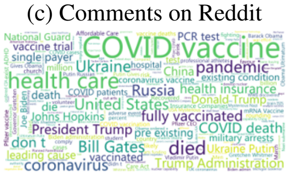

(e) Comments on X

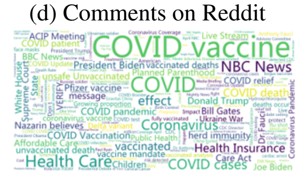

(f) Comments on X

Figure 3: Word cloud of comments in True and Fake claims. Left: Fake claim comments on YouTube, Reddit, and X. Right: True claim comments on YouTube, Reddit, and X.

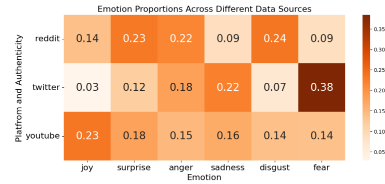

(a) Comment emotion in fake claim.

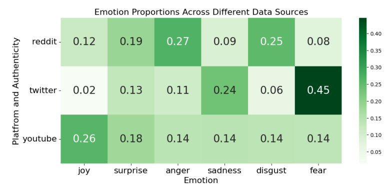

(b) Comment emotion in true claim.

Figure 4: Emotion feature of each social platform comment in news claim. Due to the prevalence of neutral emotions across all platforms, we omit neutral emotions to highlight differences.

than true claims, whereas this phenomenon is not observed on Reddit. This reveals that **discussions on Reddit are more extensive and comprehensive than on other platforms**. 2) **Not all comments across platforms contain sufficient discriminative features for detecting fake news.** Therefore, it is necessary to use different models to extract the features of comments from each platform.

Figure 3 illustrates the differences in keywords used across various platforms. Specifically, X emphasizes short texts and immediacy, hosting real-time discussions and sharing current events more than Reddit and YouTube. In contrast, Reddit and YouTube tend to feature more fragmented content, focusing on hot topics and current affairs. We also observe that both Reddit and X feature discussions on COVID-19 vaccines and political figures (e.g., Biden and Trump). However, Reddit covers more topics related to the Ukraine conflict, Russia, and related actions. Both YouTube and X discuss COVID-19 vaccines, but YouTube places more emphasis on topics related to doctors, countries, and beliefs. Through entity analysis in Figure 2d, we found that **comments on true claims, in contrast to those on fake claims, often incorporate more detailed evidence and positive reinforcement**, including references to prominent figures (e.g., Trump, Putin, Biden) or affirming terms (e.g., support, truth, believe). This highlights a clear distinction between the nature of comments on true versus fake claims. Notably, this trend is consistent across all three platforms, pointing to shared patterns in how information spreads.

Analysis of Comment Emotion and Echo Chambers To explore the emotional tendencies of comments and quantitatively describe the platform's Echo Chambers Alatawi et al. (2021), we extract the emotional features of comments shown in Figure 4. We find that **emotional features on YouTube are difficult to use in distinguishing the credibility of news.** On Reddit and X, there is a significant difference in the proportion of comments expressing different emotions between fake news and real news. However, on YouTube, this difference is almost negligible.

The relevance of comments to news claims is shown in Figure 2a. We find that **Reddit exhibits a stronger Echo Chamber effect.** Comments on Reddit are highly relevant to the claims themselves, followed by YouTube, whereas comments on X show weaker relevance. On Reddit, comments on fake claims are more relevant than comments on true claims; on YouTube, the reverse is observed. This indicates that the Echo Chamber effect on Reddit is more pronounced compared to the other two platforms.

Adaptive Propagation Structure Learning Network

In this section, we propose an efficient Adaptive Propagation Structure Learning Network (APSL) for multi-platform fake news detection. As shown in Figure 5, we use a text feature encoder to extract textual features from news articles and comments. To capture platform-specific features, we design a Platform Adapter for different nodes. Given the varying impact of propagation characteristics across platforms, we employ a multi-platform propagation encoder to extract propagation structure features. To focus on relevant comments, we use a feature fusion module to integrate platform-specific textual and news claim features. Finally, a multilayer perceptron (MLP) consolidates all features, and a Sigmoid function outputs the news authenticity probability, with the model trained using cross-entropy loss.

Multi-platform fake news detection aims to detect fake news using propagation structure information from multiple social media platforms. Formally, given a labeled claim

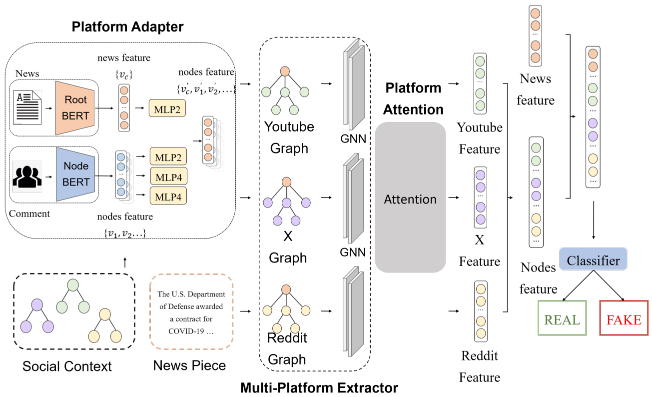

Figure 5: Overview of the proposed framework APSL.

dataset as $D = \{(c, s, y)\}$, where c represents the news claim and y represents the comments on the news claim, and s contains the comment of the news claim. $s = \{s_1, s_2, \ldots, s_m\}$, where m indicates the number of comments. The model aims to predict the True or fake label of the news claim.

Text Encoder

We utilize two text encoders to extract semantic features from the claim and its propagation on social platforms. Given the news content of the claim c and the j-th comment s_j^k on platform k, we use the text encoder to extract text embedding of claim and comments as $\mathbf{c}$ and $\mathbf{s}_j^k$.

Considering the significant differences in comments across multiple platforms, we introduce a learnable vector $\mathbf{p}^k$ to adaptively capture specific semantic features on platform k.

The platform-adaptive representations can be computed as follows:

$$\tilde{\mathbf{s}}_j^k = \text{Softmax}(\mathbf{W}^k(\mathbf{p}^k \cdot \mathbf{s}_j^k) + \mathbf{b}^k) \tag{1}$$

where $\mathbf{W}^k$ and $\mathbf{b}^k$ are trainable parameters for platform k.

Multi-Platform Propagation Encoder

Graph Construction For a news claim, there are multiple propagation graphs across different platforms.For platform k, the propagation graph are denoted as $G^k = \{V^k, E^k\}$. $V^k = \{c, s_1^k, s_2^k, ..., s_M^k\}$ denotes the set of nodes including the claim c and its propagation $s_1^k, ..., s_{M_k}^k$ on the platform k. M_k is the number of comments on the platform. Their initial node embeddings can be computed with $\mathbf{c}$ and $\tilde{\mathbf{s}}_i^k$. $E^k = \{e_1, e_2, \ldots, e_{n_e}\}$ denotes the set of edges, where the edges between nodes are built based on the relationships between reposts and comments.

Learning Platform-specific Propagation Structures In addition, for different platforms, we train a separate GNN for each platform to obtain the corresponding propagation feature representations, which allow us to learn latent structural features from the propagation specific to each platform. We use global add pooling to obtain $\mathbf{h_g}$.

Multi-Platform Propagation Fusion Module Considering that some social context features may be irrelevant to the news claim, we use an attention mechanism to select the portions that are pertinent to the news content. Specifically, after obtaining features from different platforms, we use the textual features of the news claim $\mathbf{c}$ to guide the fusion of multi-platform propagation features $\{\mathbf{h}^1, \mathbf{h}^2, \ldots, \mathbf{h}^{|P|}\}$. The mechanism then extracts the propagation features that are most relevant to the news content. Finally, we concatenate them to form the final propagation feature $\mathbf{h_s}$.

$$\mathbf{h}^k = \text{Softmax}(\frac{\mathbf{c}(\mathbf{h}_g^k)^T}{\sqrt{d_{h_g^k}}})\mathbf{h}_g^k \tag{2}$$

Fake News Detector

We concatenate the claim representation $\mathbf{c}$ and the propagation representations to fuse them into $\mathbf{h}_s$.

Based on the fusion representations, we use an MLP with Sigmoid output to predict the true or fake labels of the news:

$$\hat{y} = \text{Sigmoid}(\mathbf{W}^d(\text{Concat}(\mathbf{c}, \mathbf{h}_s) + \mathbf{b}^d) \tag{3}$$

where $\mathbf{W}^d$ and $\mathbf{b}^d$ are learnable parameters.

Platform-aware Contrastive Learning

Due to the variability and noise in the comments, we aim to capture the general features of each platform. Inspired by the success of self-supervised contrastive learning (He et al., 2020; T. Chen et al., 2020), we use contrastive learning to automatically learn the general propagation features. The objective function is computed as:

$$\mathcal{L}_{PCL} = -\sum_{k \in \mathbf{P}} \frac{1}{|B|} \sum_{i=1}^{|B|} \sum_{j=1}^{|B|} \mathbf{l}_{y^i = y^j} \log \frac{exp(sim(\mathbf{h}_i^k, \mathbf{h}_j^k)/\tau)}{\sum_{m=1}^{|B|} exp(sim(\mathbf{h}_m^k, \mathbf{h}_m^k)/\tau)} \tag{4}$$

where $|B|$ is the batchsize, $\mathbf{l}$ is an indicator. k indicates a specific platform and τ controls the temperature

This model is trained using a binary cross entropy loss function for fake news detection denoted as $\mathcal{L}_{pred}$. We combine all of the loss functions together to jointly train our model, and can be represented as follows:

$$\mathcal{L}_{final} = \mathcal{L}_{pred} + \gamma \mathcal{L}_{PCL} \tag{5}$$

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

Experiments

Experimental Setups

Baselines We compare with content-based and propagation-based fake news detection models. For content-based models, we used BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), EANN Wang et al. (2018) and gpt-4o-mini. The sequence model we use including GRU (Cho et al., 2014) and Transformer (Vaswani et al., 2017). The fake news detection model we use including BiGCN (Bian et al., 2020), UPFD (Dou et al., 2021) and UPSR (Wei et al., 2022).

Model	Accuracy	Precision	Recall	F1
EANN	0.6367	0.6378	0.6359	0.6351
GPT-4o	0.6946	0.6946	0.6944	0.6944
Backbone=BERT				
Bert	0.6707	0.6841	0.6687	0.6629
Transformer	0.6766	0.6829	0.6752	0.6727
GRU	0.6467	0.6498	0.6456	0.6437
UPFD w/ GCN	0.6766	0.6819	0.6754	0.6733
UPFD w/ GAT	0.6766	0.6835	0.6752	0.6724
UPFD w/ SAGE	0.6846	0.6887	0.6835	0.6821
BiGCN	0.6766	0.6783	0.6759	0.6752
UPSR	0.6786	0.7013	0.6762	0.6674
APSL(Ours)	**0.7046**	**0.7045**	**0.7045**	**0.7045**
Backbone=RoBERTa				
RoBERTa	0.6507	0.6595	0.6489	0.6441
Transformer	0.6607	0.6692	0.6590	0.6549
GRU	0.6547	0.6773	0.6521	0.6410
UPFD w/ GCN	0.6427	0.6473	0.6413	0.6385
UPFD w/ GAT	0.6587	0.6745	0.6565	0.6488
UPFD w/ SAGE	0.6727	0.6730	0.6729	0.6726
BiGCN	0.6453	0.6436	0.6448	0.6436
UPSR	0.6627	0.6663	0.6615	0.6598
APSL(Ours)	**0.6846**	**0.6854**	**0.6850**	**0.6845**

Table 3: Multi-platform fake news detection performance in our dataset.

Implementations Considering the imbalance between true and fake labels in the dataset, we sample news with negative labels to ensure data balance. We randomly divide the data into training, validation, and testing sets in a 7:1:2 ratio. The dimension of word embedding vector is fixed at 768. The temperature τ is 0.1, and hyperparameter γ is 0.3.We train the model with Adam optimizer. We conduct three experiments and report the average of the results.

Main Results

The results for fake news detection on the MPPFND dataset with multi-platform propagation are shown in Table 3. From the table, we have the following observations: 1) Propagation-based methods generally outperform content-based ones, highlighting the importance of propagation in fake news detection. 2) Some propagation-based methods, like the BiGCN model, incorporate multi-platform propagation but perform worse, possibly due to insufficient propagation in the MPPFND dataset. UPSR improves the modeling of incomplete propagation, outperforming BiGCN. Our APSL, by leveraging platform differences, models complex platform propagation more effectively, leading to better performance than existing methods. 3) Our model generally outperforms others, suggesting that differentiated modeling of comments across platforms is both useful and necessary for fake news detection.

Ablation Study

Model Component Analysis We analyze the effects of different components in our model, introducing three variants: one without platform-aware contrastive learning objective , one without node enhancement, and one without platform attention. The experimental results are shown in Table 4. From the table, the variants obtain declined performance than the full model, showing the effectiveness of each component. The variant without Node Enhancement struggles to capture platform-specific features, while the one without Platform Attention suffer from significant noise in user comments, leading to suboptimal detection results.

Method	Accuracy	Precision	Recall	F1
APSL	**0.7046**	**0.7045**	**0.7045**	**0.7045**
w/o Platform Contrastive Learning	0.6986	0.7091	0.6970	0.6936
w/o Node Enhancement	0.6846	0.6874	0.6837	0.6827
w/o Platform Attention	0.6727	0.6730	0.6729	0.6726

Table 4: Results of Ablation study results on the MPPFND dataset.

Dataset	Platform	Accuracy	Precision	Recall	F1
APSL (Our Model)	all	**0.7046**	**0.7045**	**0.7045**	**0.7045**
w/o Propagation		0.6707	0.6841	0.6687	0.6629
Single-Platform	X	0.6427	0.6444	0.6433	0.6423
	Reddit	0.6547	0.6546	0.6545	0.6545
	Youtube	0.6766	0.6783	0.6759	0.6752
Double-Platform	w/o X	0.6846	0.6882	0.6836	0.6823
	w/o Reddit	0.6846	0.6874	0.6837	0.6827
	w/o Youtube	0.6687	0.6806	0.6668	0.6615

Table 5: Results of using propagation across different social platforms.

Effect of Modeling Propagation across Different Platforms As shown in Table 5, we further evaluate the impact of different platform propagation methods. As removing propagation data on specific platforms, the detection performance obtains different degrees of degradation, particularly when large amounts of propagation data are missing (e.g., on X and Reddit), leading to worse detection than models trained only on news content.

Conclusion

This paper investigates fake news detection through propagation across multiple platforms. To this end, we introduce a new fake news dataset and conduct an empirical study to analyze propagation patterns across different platforms. Through semantic analysis and propagation feature analysis of the dataset we collected, we summarize the commonalities and differences in news propagation across different social media platforms, providing insights and a foundation for designing a fake news detection model across multiple social media platforms. To further demonstrate the contribution of multi-platform propagation data to detection, we compare the detection performance using data propagated on a single social media platform versus data propagated across multiple platforms. The results validate that multi-platform propagation data helps improve the effectiveness of fake news detection.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. U24A20335), the China Postdoctoral Science Foundation (No. 2024M753481), and the Postdoctoral Fellowship Program of China Postdoctoral Science Foundation (No. GZC20232969).

References

Aïmeur, E., Amri, S., & Brassard, G. (2023). Fake news, disinformation and misinformation in social media: a review. *Social Network Analysis and Mining*, *13*(1), 30.

Ajao, O., Bhowmik, D., & Zargari, S. (2019). Sentiment aware fake news detection on online social networks. In *Icassp 2019-2019 ieee international conference on acoustics, speech and signal processing (icassp)* (pp. 2507–2511).

Alatawi, F., Cheng, L., Tahir, A., Karami, M., Jiang, B., Black, T., & Liu, H. (2021). A survey on echo chambers on social media: Description, detection and mitigation. *arXiv preprint arXiv:2112.05084*.

Bian, T., Xiao, X., Xu, T., Zhao, P., Huang, W., Rong, Y., & Huang, J. (2020). Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 34, pp. 549–556).

Chalnick, A., & Billman, D. (1988). Unsupervised learning of correlational structure. In *Proceedings of the tenth annual conference of the cognitive science society* (pp. 510–516). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607).

Chen, Z., Hu, L., Li, W., Shao, Y., & Nie, L. (2023). Causal intervention and counterfactual reasoning for multi-modal fake news detection. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 627–638).

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Choudhary, A., & Arora, A. (2021). Linguistic feature based learning model for fake news detection and classification. *Expert Systems with Applications*, *169*, 114171.

Cui, C., & Jia, C. (2024). Propagation tree is not deep: Adaptive graph contrastive learning approach for rumor detection. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 38, pp. 73–81).

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dou, Y., Shu, K., Xia, C., Yu, P. S., & Sun, L. (2021). User preference-aware fake news detection. In *Proceedings of the 44th international acm sigir conference on research and development in information retrieval* (pp. 2051–2055).

Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought.* New York: McGraw-Hill.

Giachanou, A., Rosso, P., & Crestani, F. (2019). Leveraging emotional signals for credibility detection. In *Proceedings of the 42nd international acm sigir conference on research and development in information retrieval* (pp. 877–880).

He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition* (pp. 9729–9738).

Hill, J. A. C. (1983). A computational model of language acquisition in the two-year old. *Cognition and Brain Theory*, *6*, 287–317.

Hu, D., Wei, L., Zhou, W., et al. (2021). A rumor detection approach based on multi-relational propagation tree. *Journal of Computer Research and Development*, *58*(7), 1395–1411.

Hu, X., Guo, Z., Chen, J., Wen, L., & Yu, P. S. (2023). Mr2: A benchmark for multimodal retrieval-augmented rumor detection in social media. In *Proceedings of the 46th international acm sigir conference on research and development in information retrieval* (pp. 2901–2912).

Huang, K.-H., McKeown, K., Nakov, P., Choi, Y., & Ji, H. (2023, July). Faking fake news for real fake news detection: Propaganda-loaded training data generation. In A. Rogers, J. Boyd-Graber, & N. Okazaki (Eds.), *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 14571–14589). Toronto, Canada: Association for Computational Linguistics. Retrieved from `https://aclanthology.org/2023.acl-long.815` doi: 10.18653/v1/2023.acl-long.815

Jiang, Y., Bordia, S., Zhong, Z., Dognin, C., Singh, M., & Bansal, M. (2020). Hover: A dataset for many-hop fact extraction and claim verification. *arXiv preprint arXiv:2011.03088*.

Jin, Z., Cao, J., Guo, H., Zhang, Y., & Luo, J. (2017). Multimodal fusion with recurrent neural networks for rumor detection on microblogs. In *Proceedings of the 25th acm international conference on multimedia* (pp. 795–816).

Jin, Z., Cao, J., Guo, H., Zhang, Y., Wang, Y., & Luo, J. (2017). Detection and analysis of 2016 us presidential election related rumors on twitter. In *Social, cultural, and behavioral modeling: 10th international conference, sbp-brims 2017, washington, dc, usa, july 5-8, 2017, proceedings 10* (pp. 14–24).

Kaliyar, R. K., Goswami, A., & Narang, P. (2021). Fake-bert: Fake news detection in social media with a bert-based deep learning approach. *Multimedia tools and applications*, *80*(8), 11765–11788.

Khoo, L. M. S., Chieu, H. L., Qian, Z., & Jiang, J. (2020). Interpretable rumor detection in microblogs by attending to user interactions. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 34, pp. 8783–8790).

Li, Y., He, H., Bai, J., & Wen, D. (2024). Mcfend: a multi-source benchmark dataset for chinese fake news detection. In *Proceedings of the acm on web conference 2024* (pp. 4018–4027).

Lin, H., Yi, P., Ma, J., Jiang, H., Luo, Z., Shi, S., & Liu, R. (2023). Zero-shot rumor detection with propagation structure via prompt learning. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 37, pp. 5213–5221).

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., … Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Liu, Y., & Wu, Y.-F. (2018). Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 32).

Lu, Y.-J., & Li, C.-T. (2020). Gcan: Graph-aware co-attention networks for explainable fake news detection on social media. *arXiv preprint arXiv:2004.11648*.

Ma, B., Lin, D., & Cao, D. (2017). Content representation for microblog rumor detection. In *Advances in computational intelligence systems: Contributions presented at the 16th uk workshop on computational intelligence, september 7–9, 2016, lancaster, uk* (pp. 245–251).

Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B. J., Wong, K.-F., & Cha, M. (2016). Detecting rumors from microblogs with recurrent neural networks.

Ma, Z., Luo, M., Guo, H., Zeng, Z., Hao, Y., & Zhao, X. (2024). Event-radar: Event-driven multi-view learning for multimodal fake news detection. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 5809–5821).

Matlock, T. (2001). *How real is fictive motion?* Doctoral dissertation, Psychology Department, University of California, Santa Cruz.

Mehta, N., Pacheco, M. L., & Goldwasser, D. (2022). Tackling fake news detection by continually improving social context representations using graph neural networks. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 1363–1380).

Nan, Q., Cao, J., Zhu, Y., Wang, Y., & Li, J. (2021). Mdfend: Multi-domain fake news detection. In *Proceedings of the 30th acm international conference on information & knowledge management* (pp. 3343–3347).

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Nørregaard, J., Horne, B. D., & Adalı, S. (2019). Nela-gt-2018: A large multi-labelled news dataset for the study of misinformation in news articles. In *Proceedings of the international aaai conference on web and social media* (Vol. 13, pp. 630–638).

Ohlsson, S., & Langley, P. (1985). *Identifying solution paths in cognitive diagnosis* (Tech. Rep. No. CMU-RI-TR-85-2). Pittsburgh, PA: Carnegie Mellon University, The Robotics Institute.

Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of experimental psychology: general*, *147*(12), 1865.

Ran, H., & Jia, C. (2023). Unsupervised cross-domain rumor detection with contrastive learning and cross-attention. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 37, pp. 13510–13518).

Rocha, Y. M., de Moura, G. A., Desidério, G. A., de Oliveira, C. H., Lourenço, F. D., & de Figueiredo Nicolete, L. D. (2021). The impact of fake news on social media and its influence on health during the covid-19 pandemic: A systematic review. *Journal of Public Health*, 1–10.

Sheng, Q., Cao, J., Zhang, X., Li, R., Wang, D., & Zhu, Y. (2022). Zoom out and observe: News environment perception for fake news detection. *arXiv preprint arXiv:2203.10885*.

Shrager, J., & Langley, P. (Eds.). (1990). *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.

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

Silva, A., Luo, L., Karunasekera, S., & Leckie, C. (2021). Embracing domain differences in fake news: Cross-domain fake news detection using multi-modal data. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 35, pp. 557–565).

Tian, L., Zhang, X. J., & Lau, J. H. (2022). Duck: Rumour detection on social media by modelling user and comment propagation networks. In *Proceedings of the 2022 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 4939–4949).

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

Voorveld, H. A., Van Noort, G., Muntinga, D. G., & Bronner, F. (2018). Engagement with social media and social media advertising: The differentiating role of platform type. *Journal of advertising*, *47*(1), 38–54.

Wang, Y., Ma, F., Jin, Z., Yuan, Y., Xun, G., Jha, K., … 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 (pp. 849–857).

Wei, L., Hu, D., Zhou, W., & Hu, S. (2022). Uncertainty-aware propagation structure reconstruction for fake news detection. In *Proceedings of the 29th international conference on computational linguistics* (pp. 2759–2768).

Wei, L., Hu, D., Zhou, W., et al. (2021). Towards propagation uncertainty: Edge-enhanced bayesian graph convolutional networks for rumor detection. In *ACL* (pp. 3845–3854).

Yao, B. M., Shah, A., Sun, L., Cho, J.-H., & Huang, L. (2023). End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models. In *Proceedings of the 46th international acm sigir conference on research and development in information retrieval* (pp. 2733–2743).

Yin, S., Zhu, P., Wu, L., Gao, C., & Wang, Z. (2024). Gamc: an unsupervised method for fake news detection using graph autoencoder with masking. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 38, pp. 347–355).

Yu, F., Liu, Q., Wu, S., Wang, L., Tan, T., et al. (2017). A convolutional approach for misinformation identification. In *Ijcai* (pp. 3901–3907).

Zellers, R., Holtzman, A., Rashkin, H., Bisk, Y., Farhadi, A., Roesner, F., & Choi, Y. (2019). Defending against neural fake news. *Advances in neural information processing systems*, *32*.

Zhang, X., Cao, J., Li, X., Sheng, Q., Zhong, L., & Shu, K. (2021). Mining dual emotion for fake news detection. In *Proceedings of the web conference 2021* (pp. 3465–3476).

Zong, L., Zhou, J., Lin, W., Liu, X., Zhang, X., & Xu, B. (2024). Unveiling opinion evolution via prompting and diffusion for short video fake news detection. In *Findings of the association for computational linguistics acl 2024* (pp. 10817–10826).