Claim-Centric And Sentiment Guided Graph Attention Network for Rumour Detection

Sajad Ramezani[1], **Mauajama Firdaus**[1], **Lili Mou**[1+]
[1]Dept. Computing Science, Alberta Machine Intelligence Institute (Amii), University of Alberta
[+]Canada CIFAR AI Chair, Amii
srameza1@ualberta.ca, mauzama.03@gmail.com, doublepower.mou@gmail.com

Abstract

Automatic rumour detection has gained attention due to the influence of social media on individuals and its pervasiveness. In this work, we construct a representation that takes into account the claim in the source tweet, considering both the propagation graph and the accompanying text alongside tweet sentiment. This is achieved through the implementation of a hierarchical attention mechanism, which not only captures the embedding of documents from individual word vectors but also combines these document representations as nodes within the propagation graph. Furthermore, to address potential overfitting concerns, we employ generative models to augment the existing datasets. This involves rephrasing the claims initially made in the source tweet, thereby creating a more diverse and robust dataset. In addition, we augment the dataset with sentiment labels to improve the performance of the rumour detection task. This holistic and refined approach yields a significant enhancement in the performance of our model across three distinct datasets designed for rumour detection. Quantitative and qualitative analysis proves the effectiveness of our methodology, surpassing the achievements of prior methodologies.

Keywords: rumour detection, propagation graph, graph neural network

1. Introduction

Social media exerts a substantial influence on individuals' day-to-day existence, with platforms like Twitter, Instagram, and others claiming an ever-expanding share of people's time. This paradigm shift in information consumption underscores the critical role that social media now plays in shaping our perspectives and choices.

The task of rumor detection (Liu et al., 2015; Rao et al., 2021; Ma et al., 2016) holds significance within the realms of Natural Language Processing (NLP), information retrieval and other related fields of research. In today's digital landscape, where the swift propagation of false information and fake news (Lu and Li, 2020; Kang et al., 2021) is prevalent on social media platforms, the role of rumor detection becomes especially critical. This phenomenon, as observed in instances of misinformation concerning public health, political elections can exert considerable adverse impacts (Friggeri et al., 2014).

Rumour detection involves employing NLP techniques to examine different aspects of textual content, including language patterns, sentiment, source credibility, and audience engagement (Castillo et al., 2011). However, the sheer volume and diversity of rumours circulating on these platforms make manual efforts insufficient (Guo et al., 2022). Consequently, models have been developed to detect and classify rumours, with a particular focus on platforms like Twitter (Ma et al., 2022). Various datasets and benchmarks have been created to assess the effectiveness of different algorithms in the task of rumour detection (Ma et al., 2016).

Early models relied on feature engineering techniques to capture distinctive characteristics. For instance, Ma et al. (2015) extracted content and user features, including average affective scores (Ajao et al., 2019; Zhang et al., 2021; Giachanou et al., 2019), and employed time series analysis to capture temporal variations and the frequency of emojis. These models lack the complexity necessary to capture the nuanced nature of rumours, including their diverse components such as text, images, links, user information, and temporal aspects.

Moreover, how a tweet spreads through different social platforms can provide insightful information regarding its content and purpose. Since false rumours are designed to "go viral," they have different propagation patterns to normal tweets (Ma et al., 2017). Therefore, an effective method to fuse these modalities is needed to make use of all the available information to train more accurate models, which has been the recent trend in the rumour detection frameworks (Sun et al., 2022),(Bian et al., 2020). Graph neural network was proposed to capture propagation structure to aggregate information in retweets and threads (Bian et al., 2020).

Sentiment analysis offers a compelling avenue for enhancing rumor detection methodologies. By incorporating sentiment as a pivotal component, our research endeavors to provide a nuanced and comprehensive approach to rumor detection, enabling a more accurate and contextually aware classification of information circulating in social

LREC-COLING 2024, pages 3235–3241
20-25 May, 2024.

media and online platforms.

The contributions of our work are fourfold. First, we propose a novel graph neural network architecture with two-level attention to capture the nuanced features in the source tweet's claim and the tweet's graph representation. Second, we incorporate the sentiment of each tweet to enhance the performance of the rumor detection task. Third, to overcome overfitting, we augment the datasets via paraphrasing using generative models. Finally, analysis shows that our approach achieves better performance compared to existing methods.

2. Related Work

Early work on automatic rumor detection employed feature engineering, such as analyzing post sentiments and source information (Liu et al., 2015). With the success of deep learning approaches in natural language processing, recurrent neural networks (RNNs) have been utilized to classify rumors based on source tweet content (Chen et al., 2018; Ma et al., 2016). Although highly effective in capturing textual representations, these approaches often overlook the information conveyed in rumor propagation.

Ma et al. (2017) was among the first to incorporate both graph and text modalities by designing tree kernels to compare the distance between source tweet contents and their propagation paths. Various augmented propagation graphs have since been developed, such as graphs with user data nodes (Kang et al., 2021).

In a subsequent study, Ma et al. (2018) enhanced the rumor detection model by employing Recursive Variational Neural Networks (RvNNs) in both bottom-up and top-down directions. Additionally, addressing the issue of information bottleneck in RvNNs, Bian et al. (2020) proposed concatenating the source tweet node representation with other nodes to alleviate this challenge. Lu and Li (2020) introduced the co-attention mechanism to improve the attention given to user features relative to the source tweet.

Our methodology uniquely emphasizes joint representation learning for both the claims made in the source tweets and the associated retweets and comments. Furthermore, it incorporates information from the propagation graph within the end-to-end process for rumor classification.

3. Approach

Our method uses a graph neural network to aggregate representations from all the relevant tweets and captures both propagation structure and textual content to aggregate the information from neighbouring nodes.

3.1. Standard Graph Neural Network

Graph Neural Network (GNN) is an architecture of neural networks based on message passing and aggregation. Given a Graph $G = (V, E)$ and its node features X, GNN updates node representation iteratively based on neighbouring nodes' representations. To update a node representation (h_v), first, each neighbour sends a message $m_{v,u} = f(h_u)$ where $f(.)$ is the message passing function (e.g., linear transformation). Then, the node representation h_v is updated based on messages and its previous representation as follows:

$$a_v = AGG(m_{v,u} \, \forall u \in N(v)) \quad (1)$$

$$h_v = COMBINE(a_v, h_v) \quad (2)$$

The AGG function aggregates messages $m_{v,u}$ from neighbors u of node v into a unified representation a_v, commonly through sum, mean, or max operations. The COMBINE function merges a_v with the current node representation h_v, typically employing neural layers, to update h_v. To classify the graph, a readout function is used (e.g., mean, max pooling) to obtain the graph embedding, which is then fed as input to a classifier.

Node features X are depicted as the initial node representation to the graph. Previous works on for rumour detection used a pre-trained BERT (Sun et al., 2022) transformer to capture document embedding from each document text and then these document embeddings were used as node features as shown in Figure 1. Unlike the previous works, we intend to capture document embedding jointly with the node representation and update them iteratively in training using source tweet text.

3.2. Document Embedding (Word Level Attention)

As illustrated in Figure 1, we use an attention module to feed the source tweet embedding S (captured by a pre-trained BERT transformer) into a shallow neural network (MLP) to get an appropriate query vector Q. By employing the attention mechanism on the obtained query vector and the word embeddings of the tweet's document (indicated by D_v) as values, we obtain the document embedding of the tweet (e_i) such that $D_v = [w_1, ..w_k]$, where w_i is the word embedding for $\forall v \in V$ and k is the maximum length of the document.

$$Q = Relu(W_1^T S + b_1); \quad (3)$$

$$K = Relu(W_2^T D_v + b_2); \quad (4)$$

$$e_v = Softmax(QK^T) D_v. \quad (5)$$

This above formulation captures document embedding e_i as a parameterized weighted average of its words. For batch processing, since each

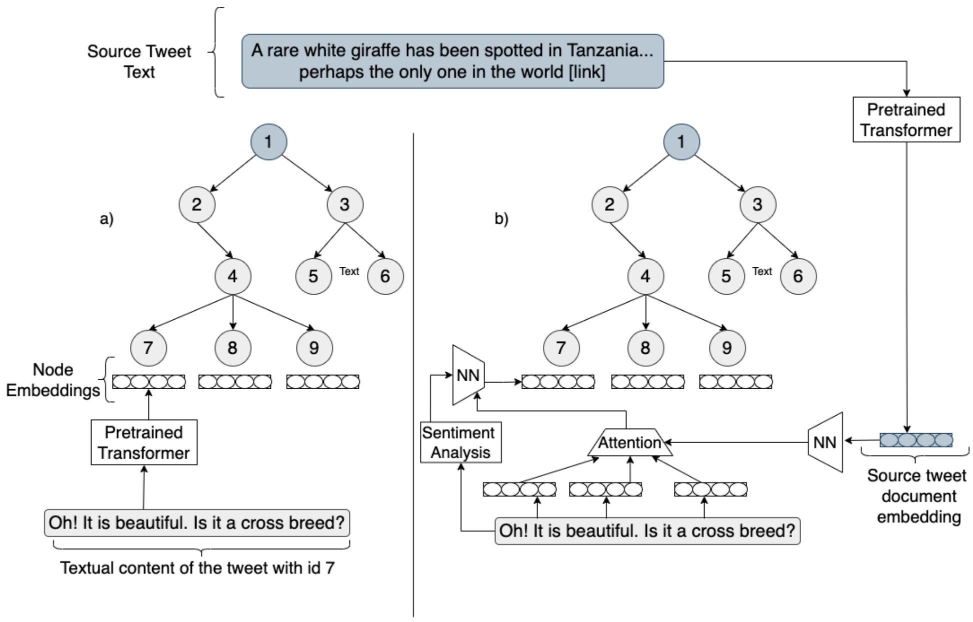

Figure 1: Comparison of (a) fixed pre-trained document embedding as node features (b) parameterized document embedding module based on source tweet attention weights and node sentiment

tweet may have a variable length, we apply zero-padding up to a maximum length of k.

As shown, rumors tend to evoke either positive or negative sentiments and are rarely neutral, as these posts are more likely to go viral (Vosoughi et al., 2018). Therefore, incorporating the sentiment of reactions (i.e., retweets and comments) to the source tweet can enhance the representation for rumor detection. We extracted the sentiment of each reaction using VADER and then used a shallow neural network to combine the document embedding constructed from word-level attention and sentiment.

$$r_v = relu\left(W_3^T(e_v||s_v) + b_3\right) \tag{6}$$

In the above equation, $||$ represents the concatenation operator, e_v is the result of the word-level attention described previously, and s_v is the sentiment vector of node v in one-hot format. In each iteration, the representation of each node is first aggregated with the obtained document embedding: $h_v = AGG(h_v, r_v)$, and then updated according to Equation 1 with incoming messages from its neighbors.

3.3. Graph Attention Network (Document Level Attention)

Velickovic et al. (2017) proposed the Graph Attention Network (GAT), which computes attention coefficients for each node in the graph.

We redesign the attention mechanism to adjust the weights based on the source tweet rather than the node itself. Therefore, instead of using the destination node features to calculate the attention scores h_v we utilize the source tweet embedding S. This modification has proven to be influential, resulting in increased model accuracy, as discussed in the results section. Previous studies have demonstrated that concatenating the source tweet embedding improves the performance of Graph Neural Networks (GNNs) for rumor detection (Bian et al., 2020).

By employing the two proposed attention mechanisms, we enable the model to attend to features at all levels (word and document) based on source tweets, while selectively capturing indicative features through the propagation graph.

3.4. Model Training

As depicted in Figure 1, we utilize the aforementioned two-level attention mechanism along with node sentiment to compute node embeddings. Utilizing a variation of the Graph Attention Network (GAT), we propagate messages throughout the network. Subsequently, employing a max-pooling readout function, we capture the graph representation. Finally, for rumor detection, this graph representation is fed into a shallow neural network to classify into rumor and non-rumor classes.

$$h_g^{(i)} = Max\left(\left\{h_v^{(i)} \mid \forall v \in V\right\}\right) \tag{7}$$

Here, h_g represents the graph representation, and h_g^i denotes the i^{th} dimension of the graph representation vector, computed by taking the maximum value across the i^{th} dimension of each node embedding. The parameters of the graph neural network are trained using the negative log-likelihood loss function.

Statistics	Twitter 15	Twitter 16	PHEME
# Users	276,663	173,487	197,852
# Source Tweets	1490	818	6425
# Non Rumours	374	205	4023
# False Rumours	370	205	2402
# True Rumours	372	205	NA
# Unverified Rumours	374	203	NA

Table 1: Dataset Statistics

4. Dataset and Experiments

4.1. Datasets

To evaluate the effectiveness of the proposed model for rumor detection, we tested it on three publicly available datasets extracted from Twitter. We followed the same setup as previous works (Sun et al., 2022) for each dataset to ensure comparability. Twitter15 and Twitter16, introduced by Ma et al. (2017), contain four classes for rumors (true, false, unverified, and non-rumor). The PHEME dataset has two classes (non-rumor and false rumor) (Zubiaga et al., 2017). The complete statistics of these datasets are provided in Table 1.

4.2. Data Augmentation

Rumor detection tasks, especially those involving benchmark datasets, consistently face the challenge of overfitting. A notable observation is that due to the limited size of rumor detection datasets, Graph Neural Network (GNN) models are particularly susceptible to overfitting (Sun et al., 2022). To address this concern, we have adopted a data augmentation strategy centered around generative models.

Our augmentation strategy involves paraphrasing each rumor source tweet text, as well as its propagation graph. Specifically, we employed the PEGASUS model (Zhang et al., 2020)—a pre-trained transformer renowned for its proficiency in abstractive summarization. For example, an original text reading, ”Major police operation unfolding in downtown Sydney <link>,” was paraphrased as ”There is a major police operation in downtown Sydney.” From each source tweet, we generated two alternative paraphrased versions, effectively tripling the size of our data. During this generation process, we opted for a num_beams parameter value of 5 to ensure diversity in the output paraphrases.

Subsequently, we trained our model using the augmented dataset to enhance its robustness and generalization capabilities. Additionally, we augmented the dataset with sentiment labels, i.e., positive, negative, and neutral, to improve the performance of the rumor detection task.

4.3. Training Details

The sentence and word embeddings are extracted using SentenceBERT (Reimers and Gurevych, 2019). The model is trained using one V100-SXM2 Nvidia GPU for 100 epochs. To construct tweet text embeddings from word embeddings, we employed two shallow neural networks with a dimensionality of 16. Finally, we concatenated the results of sentiment analysis with the document embeddings and passed them through another shallow neural network, resulting in node embeddings of size 768. The proposed graph neural network propagates node embeddings using two document-level attentions with three heads, each with 64 dimensions. For classification, maximum pooling is used as a readout function, along with two fully connected neural networks of size 64 and a dropout layer with a rate of 0.5.

5. Results and Analysis

5.1. Main Results

As depicted in Table 2, our proposed model enhances rumor detection accuracy and achieves state-of-the-art performance on all datasets, particularly on the *PHEME* dataset, which contains five times more records than the other two. Our model achieves a 4.8% increase in accuracy on this dataset. For the other two datasets, it achieves a higher F1 score in the True and False rumor classes, the original classes for the rumor detection task. This improvement can be attributed to two components we introduced: firstly, the parameterization of document embedding enables fine-grained control to represent the node embedding more effectively.

Secondly, incorporating sentiment information directly into the embedding also encourages the model to prioritize it. Furthermore, as the query vector for all attention mechanisms is derived from the source tweet, the model can better capture the most indicative words. This approach is more intuitive than previously proposed methods, such as concatenating the hidden representation of the source tweet with other nodes as suggested in Bian et al. (2020).

Moreover, in addressing overfitting, a significant challenge in the rumor detection task, utilizing generative models to paraphrase source tweet claims proves more effective than previous methods, such as the adversarial feature transformation (AFT) component proposed in GACL (Sun et al., 2022). Additionally, it is more efficient as it is performed only once at the beginning of training, unlike AFT, which is computed at every iteration.

5.2. Ablation Study

In this section, we conducted an ablation experiment to demonstrate the effectiveness of each

Method	PHEME			Twitter 15					Twitter 16				
	Acc	*R F1*	*N F1*	*Acc*	*U F1*	*N F1*	*T F1*	*F F1*	*Acc*	*U F1*	*N F1*	*T F1*	*F F1*
cPTK (Ma et al., 2017)	-	-	-	0.750	0.733	0.804	0.765	0.698	0.732	0.686	0.740	0.836	0.709
RvNN (Ma et al., 2018)	0.763	0.631	0.825	0.723	0.654	0.682	0.821	0.758	0.737	0.708	0.662	0.835	0.743
Bi-GCN (Bian et al., 2020)	0.824	0.741	0.865	0.886	0.864	0.891	0.930	0.860	0.880	0.865	0.847	0.937	0.869
GACL (Sun et al., 2022)	0.850	0.771	0.885	0.901	**0.876**	**0.958**	0.903	0.851	0.920	**0.907**	**0.934**	**0.959**	0.869
Our Method	**0.891**	**0.843**	**0.914**	**0.908**	0.873	0.896	**0.946**	**0.873**	**0.925**	0.883	0.918	**0.959**	**0.883**

Table 2: Main Results. Here R F1:Rumour F1, N F1:Non-Rumour F1, U F1: Unverified Rumour F1, T F1: True Rumour, F F1: False Rumour F1

Source Tweet	Predicted Class	True Class
cdc whistleblower exposes ebola vaccinations containing rfid chips | national report URL via @wpusta	rumour	rumour
why did macklemore delete this tweet!? #macklemorejoinedisis URL	non-rumour	non-rumour
seth rogen has been cast as steve wozniak in the steve jobs biopic starring christian bale. URL URL	unverified rumour	false rumour
morocco says it has arrested a belgian national linked to men who carried out paris attack, in which 130 people died	unverified rumour	true rumour
this little ios 8 bug deletes all your icloud data if you try and reset your settings: URL URL	unverified rumour	unverified rumour
poland says cannot accept migrants under eu quotas after paris attacks. more updates: URL URL	non-rumour	non-rumour

Table 3: Error Analysis

component of the proposed design. We tested the following variations of our model:

1. Without Word level and Document level Attention (WWDA): It utilizes a Graph Attention Network (GAT) alongside a pre-trained transformer model to capture node embeddings from the text content of the tweets. However, this model does not incorporate any of the proposed designs i.e., word and document level attention.

2. Document Level Attention (DLA): It utilized the proposed two-level attention mechanism to capture a more fine-grained node representation regarding the source tweet claim. However, it does not incorporate sentiment and was not trained on the augmented dataset using paraphrasing.

3. Document Level Attention and Sentiment (DLAS): It captures the node embeddings using the two-level attention mechanism, while also considering the sentiment of the tweet. However, it was not trained on the augmented dataset.

4. Our Method: It includes all the different components of the proposed model, including two-level attention, sentiment, and augmented dataset.

The accuracy of the mentioned variations on the PHEME dataset is displayed in Table 4. As shown, our two-level attention mechanism can achieve better node representation and result in more accurate classification. Utilizing sentiment information increases accuracy by 2%, validating our assumption that rumor detection and sentiment analysis are related tasks, and one can benefit from tweets' sentiments to better detect rumors. Lastly, augmenting the dataset by paraphrasing the source tweet enables our model to achieve higher results, with an increase of 1.8% from the version without augmentation. Since real-world datasets are always limited and biased, synthesizing and augmenting data with the help of generative models can help GNN models construct better node representations.

Model Variation	Accuracy
WWDA	0.842
DLA	0.858
DLAS	0.875
Our Method	0.891

Table 4: Ablation Study on the PHEME dataset

5.3. Error Analysis

In our evaluation, we observed certain challenges the model faced in discerning between various rumor stances, as illustrated in Table 3. The main source of discrepancies arose when differentiating between "unverified", "false", and "true" rumor classifications. For instance, the tweet regarding Seth Rogen's purported casting as Steve Wozniak was predicted as an "unverified rumor", yet its true label was a "false rumor".

However, a silver lining in our findings is the model's consistent ability to correctly identify a statement as a rumor or non-rumor. This strength in categorization was particularly evident in our experiments on the PHEME dataset, which is designed to distinguish between these two primary classes.

6. Conclusion and Future Work

In this paper, we propose a novel method for constructing document embeddings based on the source tweet's claim to identify rumors. Additionally, we aggregate document embeddings using a two-level attention mechanism along with sentiment information to enhance the overall task performance. This approach advances the state of the art on three rumor datasets.

In the future, we plan to apply the same method to other node classification tasks with document-level features for nodes and also evaluate transfer learning for these tasks.[1]

[1]we made our code available at

Acknowledgements

We thank all reviewers and chairs for their valuable comments. The research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) under Grant No. RGPIN2020-04465, an Alberta Innovates Project, the Amii Fellow Program, the Canada CIFAR AI Chair Program, a UAHJIC project, a donation from DeepMind, and the Digital Research Alliance of Canada (alliancecan.ca).

Bibliographical References

Oluwaseun Ajao, Deepayan Bhowmik, and Shahrzad Zargari. 2019. Sentiment aware fake news detection on online social networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2019*, pages 2507–2511. IEEE.

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 *The Thirty-Fourth AAAI Conference on Artificial Intelligence 2020*, pages 549–556.

Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th International Conference on World Wide Web, 2011*, pages 675–684.

Tong Chen, Xue Li, Hongzhi Yin, and Jun Zhang. 2018. Call attention to rumors: Deep attention based recurrent neural networks for early rumor detection. In *Trends and Applications in Knowledge Discovery and Data Mining - PAKDD*, pages 40–52.

Adrien Friggeri, Lada A. Adamic, Dean Eckles, and Justin Cheng. 2014. Rumor cascades. In *Proceedings of the Eighth International Conference on Weblogs and Social Media, 2014*.

Anastasia Giachanou, Paolo Rosso, and Fabio Crestani. 2019. Leveraging emotional signals for credibility detection. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2019*, pages 877–880.

Zhijiang Guo, Michael Sejr Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.

https://github.com/sajjadGG/RumorDetection

Zhezhou Kang, Yanan Cao, Yanmin Shang, Tao Liang, Hengzhu Tang, and Lingling Tong. 2021. Fake news detection with heterogenous deep graph convolutional network. In *Advances in Knowledge Discovery and Data Mining: 25th Pacific-Asia Conference, PAKDD 2021*, page 408–420.

Xiaomo Liu, Armineh Nourbakhsh, Quanzhi Li, Rui Fang, and Sameena Shah. 2015. Real-time rumor debunking on twitter. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, 2015*, pages 1867–1870.

Yi-Ju Lu and Cheng-Te Li. 2020. GCAN: graph-aware co-attention networks for explainable fake news detection on social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020*, pages 505–514.

Guanghui Ma, Chunming Hu, Ling Ge, Junfan Chen, Hong Zhang, and Richong Zhang. 2022. Towards robust false information detection on social networks with contrastive learning. In *Proceedings of the 31st ACM International Conference on Information amp; Knowledge Management*, page 1441–1450.

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 *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, page 3818–3824.

Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, page 1751–1754.

Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–717.

Jing Ma, Wei Gao, and Kam-Fai Wong. 2018. Rumor detection on Twitter with tree-structured recursive neural networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1980–1989.

Dongning Rao, Xin Miao, Zhihua Jiang, and Ran Li. 2021. STANKER: stacking network based on level-grained attention-masked BERT for rumor detection on social media. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021*, pages 3347–3363.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 3980–3990.

Tiening Sun, Zhong Qian, Sujun Dong, Peifeng Li, and Qiaoming Zhu. 2022. Rumor detection on social media with graph adversarial contrastive learning. In *WWW '22: The ACM Web Conference 2022*, pages 2789–2797.

Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2017. Graph attention networks. *CoRR*, abs/1710.10903.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. pages 1146–1151.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. PEGASUS: pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, 2020*, pages 11328–11339.

Xueyao Zhang, Juan Cao, Xirong Li, Qiang Sheng, Lei Zhong, and Kai Shu. 2021. Mining dual emotion for fake news detection. In *WWW '21: The Web Conference 2021*, pages 3465–3476.

Arkaitz Zubiaga, Maria Liakata, and Rob Procter. 2017. Exploiting context for rumour detection in social media. In *Social Informatics - 9th International Conference, SocInfo 2017*, pages 109–123.