Fake News Detection Using Temporal Snapshots in Graph Neural Networks

Sampanna Sharma, Yi Li

Abstract—**The widespread propagation of misinformation in social media has been a significant barrier to information credibility and trust among the general public. It is crucial to address this problem in order to stop misinformation from twisting public conversations. Despite many deep learning methods having been employed for the detection of fake news, one task remains challenging: capturing complex and ever-evolving interactions present in social media. In this paper, we use dynamic Graph Neural Networks for fake news detection on Twitter. We study news articles' content and the evolving patterns of user engagement to provide a more accurate and context-aware classification of fake news, and its early detection, and stop misinformation from spreading. Our results illustrate that the combination of temporal dynamics with network structure creates a more complete approach to improving information reliability on social media and combating misinformation.**

Index Terms—**graph neural networks, dynamic graphs, temporal modeling, fake news detection, early detection, social media interactions**

I. Introduction

THE fake news on social media is one of the major issues in the digital age. Fake news that spreads rapidly and is intentionally generated to deceive or influence public opinions can lead to serious consequences, including the undermining of democratic processes and worsening of public health emergencies [1], [2]. Since the spread of counterfeit information on social networking sites is much faster as compared to the actual information, its detection and reduction become more challenging [3]. While social media is an ever-increasing means of news material, this increases difficulties in distinguishing between real and misleading information, thus requiring more comprehensive methods of detection.

Traditional methods for false news detection have thus far mainly depended on static features: sentiment analysis, lexical patterns, and source reliability [4]. While these methods perform well in a controlled environment, they invariably fail in the real world because of the scale and speed at which data come in. Moreover, these techniques do not consider the complex structure of the social network and time-related contexts in which information is shared over social platforms [5]. Consequently, they find it challenging to adjust to the ever-changing and developing social media landscapes.

More recently, advances in machine learning, especially those on Graph Neural Networks (GNNs), have created new routes to investigate the complexities in fake news detection [6]. The techniques of GNNs are of particular relevance in dealing with this problem because they can model the relational and structural information that is encoded in social networks [7], [8]. GNNs can capture complicated patterns that characterize the spread by representing users, posts, and their interactions as nodes and edges in a graph. However, most of the GNN-based approaches have regarded these graphs as static, failing to consider the temporal dynamics so vital to understanding how fake news propagates [9]. Misinformation spreads dynamically, with user engagement and interaction patterns changing rapidly over time. Ignoring this temporal aspect can result in less effective models in real-time detection scenarios [10].

Sampanna Sharma and Yi Li are with the Department of Computer Science, The University of Texas at Tyler, Tyler, Texas, 75799 USA (e-mail: yli@uttyler.edu; ssharma8@partriots.uttyler.edu).

To address these limitations, we employ a Heterogeneous Dynamic Graph Neural Network method that specifically considers the temporal progression of tweet interactions to enhance the detection of fake news. In the method we use, users, tweets, and their interactions are depicted as nodes and edges in a dynamic graph that changes over time. This will inherently allow the model to learn from news content as well as changes in user behavior trends, resulting in much more accurate fake news classification with better contextual understanding. Temporal dynamics are added, which is critical because it reflects the complex nature of social media interactions in the real world; the importance and exposure of information change in just a blink of an eye.

The structure of this article is as follows: Section II reviews relevant literature, while Section III provides a detailed description of the proposed approach. The experimental setup is outlined in Section IV, and Section V presents the results. Finally, Section VI offers concluding remarks.

II. Related Works

The field of fake news detection has progressed from content-based analysis to more sophisticated models that incorporate multi-modal data, social network structures, and temporal dynamics. Early approaches focused on linguistic and stylometric features using NLP techniques such as word frequency, syntactic structures, and readability metrics. While these methods showed promise, they often struggled with more complex fake news that mimicked legitimate writing styles [5], [11]. Sentiment analysis was later introduced to detect emotionally charged language, commonly found in sensational fake news, though its effectiveness was limited to content with explicit emotional cues [12]. The advent of deep learning, with models like CNNs and RNNs, further improved fake news detection by capturing local and sequential dependencies in text, though challenges remained with context-dependent language [13].

User behavior and social context became critical in enhancing detection models. Behavioral patterns such as content sharing frequency and user activity provided insights into the likelihood of fake news dissemination [14], while social influence models leveraged network dynamics to account for the role of peer influence and echo chambers in spreading misinformation [15]. GNNs further advanced the field by modeling social relationships, capturing both network structure and user interactions, though early models were limited by static network representations [16].

More recently, multi-modal methods have combined data from different formats-text, images, and videos-into unified models, which permits more powerful fake news detection across media types [17]. Hybrid models, which merge NLP techniques with network analysis or deep learning, turn out to be particularly effective, addressing both the content and the social context of news articles [18]. Temporal Graph Neural Networks brought temporal dynamics to fake news detection, modeling how social interactions evolve over time [10], while event-based models focused on the identification of misinformation bursts during critical events [19].

We extend these works with a dynamic graph where nodes represent news articles, users, and tweets, and edges capture interactions like retweets, replies, and authorship. We model the temporal evolution of these interactions, thus making possible more accurate and context-aware fake news classification in the fast-changing environment of social media.

III. Proposed Approach

In this paper, we solve the problem of fake news detection by constructing a dynamic heterogeneous graph evolving along with time. We denote the dynamic graph as:

$$G(t) = (\mathcal{V}(t), \mathcal{E}(t)) \tag{1}$$

where $G(t)$ is the dynamic graph at time t, $\mathcal{V}(t)$ denotes the set of nodes, and $\mathcal{E}(t)$ denotes the set of edges at time t. Our goal is to predict the class label $\hat{y}(v_i, t)$ for each node v_i in $\mathcal{V}(t)$.

A. Dynamic Heterogeneous Graph Construction

We construct a dynamic heterogeneous graph $G(t)$ where the set of nodes $\mathcal{V}(t)$ is divided into three main types of nodes, namely *news articles*, *tweets*, and *users*. Each type of node $v \in \mathcal{V}(t)$ has a feature vector $\mathbf{X}(v)$ which captures the properties useful for classification. Namely:

- News articles v_n are represented by their textual content embeddings $\mathbf{X}(v_n)$ obtained form natural language processing models.
- Tweets v_t are described by their textual embeddings $\mathbf{X}(v_t)$ representing the semantic meaning of the tweet's content.
- Users v_u are described by attributes including credibility scores, follower count, and engagement metrics $\mathbf{X}(v_u)$.

Edges $\mathcal{E}(t)$ model interactions between nodes and are undirected, with examples including "news articles replied to by tweets" and "tweets posted by users." These edges are essential for tracking the propagation of information along with the influence the latter has on user activities. Fig 1 illustrates the dynamic relationships between news articles, tweets, and users across three consecutive time steps.

The temporal aspect is captured by taking snapshots of $G(t)$ at discrete instants of time t such that each snapshot represents the state of the graph at time t. Such temporal succession will enable us to observe the evolution occurring in the dissemination of fake news. New interactions that occur, like a user retweeting a tweet, will be part of future snapshots, preserving the correct temporal order. This methodology allows for changes in the model both structurally and feature-based over time.

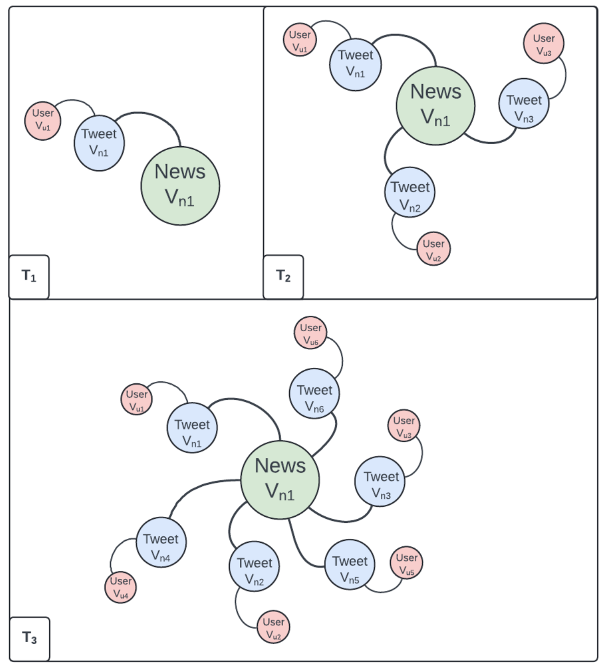

Fig. 1. Dynamic Evolution of News, Tweet, and User Interactions in a Temporal Graph(T3 >T2 >T1)

B. Heterogeneous Graph Convolutional LSTM Model

To model both spatial and temporal dependencies in dynamic heterogeneous graphs, we propose the Heterogeneous Graph Convolutional Long Short-Term Memory (HeteroGCLSTM) network. This model integrates heterogeneous graph convolutional layers with LSTM cells to effectively address the complexity of heterogeneous graph structures and their temporal dynamics. It has the following keys components:

1) **Heterogeneous Graph Convolutional Layers**: These layers perform convolutional operations over heterogeneous graphs, taking into account different node and edge types. For each node type τ, the graph convolution aggregates information from neighboring nodes using:

$$h_\tau^{l+1} = \sigma\left(\sum_{r \in R} W_r \cdot \mathrm{AGG}_r\left(\{h_v^l : v \in N_r(u)\}\right)\right) \tag{2}$$

where h_τ^l represents the node features at layer l, W_r is a learnable weight matrix for relation r, and AGG_r

is an aggregation function. This mechanism preserves semantic differences between node types, such as news articles versus tweets, ensuring accurate propagation of information through the graph structure.

2) **Temporal Processing via LSTM Cells**: The LSTM component in our model is designed to capture temporal dependencies across the dynamic graph snapshots. At every step in time, it takes the inputs regarding evolving structure and feature information about nodes, such that it is able to recognize patterns of temporal variation in changing interactions and features.
By integrating temporal patterns with graph-based features, the LSTM links the historical and current states of the graph. This allows the model to adapt to dynamic user behavior, tweet propagation, and content evolution. It enables the effective capture of long-term dependencies and contextual changes that are necessary for the accurate classification of news articles as real or fake in a fast-evolving social media environment. Fig. 2 illustrates the overall architecture, highlighting the integration of the graph convolution and LSTM components.
3) **Fully Connected Layer**: Lastly the fully connected layer combines temporal context from the LSTM cells and graph-based features to produce the predictions. For a news article n at time t, the final prediction is computed as:

$$\hat{y}(n,t) = \sigma(W_p \cdot h_t^n + b_p) \tag{3}$$

where h_t^n is the final hidden state for the news node n, and σ is the sigmoid activation function.

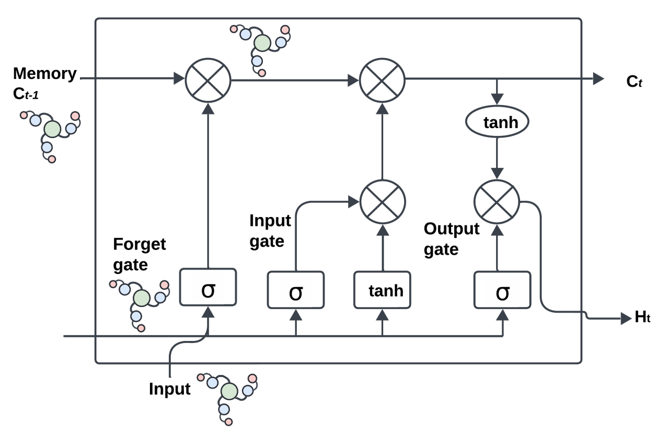

Fig. 2. High-level framework of the Graph Convolutional LSTM Model

IV. EXPERIMENTAL SETUP

To evaluate the performance of the proposed HeteroGCLSTM model, a thorough experimental setup is created for data preparation, model training, and performance evaluation. This section discusses each element of the experimental arrangement in detail.

A. Dataset

We utilize *TruthSeeker2023* dataset [20] for real and fake content being propagated on social media. It consists of three main entities: *news articles*, *tweets*, and *users*. Each instance in the dataset is annotated and includes the complete text of the news articles, related tweets, engagement metrics, and timestamps to track information dissemination. User data includes credibility score, social influence indicators (such as number of followers and followees), metadata, and interaction patterns (e.g., retweets, mentions) that form structural edges of the network.

Prior to model training, we perform several pre-processing steps:

- **Textual Data Encoding**: News articles and tweets are converted into dense vector embeddings using pre-trained BERT sentence transformer [21]. This transformation captures the semantic content of the text and facilitates effective feature extraction.
- **User Attribute Extraction**: User attributes, including credibility scores are normalized and encoded into numerical feature vectors.
- **Graph Construction**: We build the dynamic heterogeneous graph $G(t)$ by mapping entities (news articles, tweets, users) to unique identifiers and create edges based on interactions for each temporal snapshot t.

B. Model Training

We train the HeteroGCLSTM model on the training set and the training process involves the following steps:

- **Loss Function**: We utilize a cross-entropy loss function to evaluate the difference between the predicted and true class labels of news articles. This loss function is appropriate for binary classification tasks, thus helping us find the best parameters of this model. The cross-entropy loss function is defined as:

$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log(p_i) + (1-y_i)\log(1-p_i)\right] \tag{4}$$

  where N is the number of samples, p_i is the predicted probability of the positive class for the i-th sample, and y_i is the true label of the i-th sample.
- **Optimization Algorithm**: The Adam optimizer [22] is used to minimize the loss function. Adam has been selected because it provides adaptive learning rates and is effective for large datasets.
- **Training Procedure**: Our model is trained for various epochs. Early stopping criterion according to the validation loss prevents overfitting. We track the training progress based on accuracy and loss metrics regarding the validation set.

C. Performance Evaluation

To comprehensively evaluate our model's performance, we employ the following metrics:

- **Accuracy**: The proportion of correctly classified articles out of the total articles evaluated.
- **Precision**: The proportion of true positives among all instances predicted as positive.

- **Recall**: The proportion of true positives among all actual positive instances identified.
- **F1-Score**: The harmonic mean of precision and recall, balancing both false positives and negatives.

To further validate the efficacy of the model, we compared HeteroGCLSTM against several state-of-the-art GNN architectures. Moreover, we did an analysis of the temporal performance of the model to assess its adaptiveness regarding changes in the graph structure and features over time.

V. Results

This section demonstrates how the HeteroGCLSTM model performs in fake news detection through its dynamic graph structure using results on the TruthSeeker2023 dataset. A comparison with other GNN architectures and robustness across various stages of temporal evolution is also provided.

A. Performance Metrics

Our model yielded the results summarized in Table I, indicating its performance for the test set based on various metrics such as accuracy, precision, recall, and F1-score.

TABLE I
Performance Metrics of the HeteroGCLSTM Model

Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
87.89	87.91	87.89	87.90

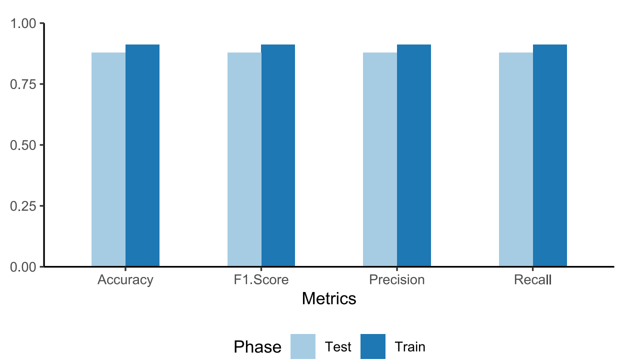

Fig. 3. Comparison of Train and Test Performance

Fig 3 illustrates the Accuracy, Precision, Recall, and F1 Score for the model in both the Train and Test phases. It is clear that when this model achieves an accuracy of 87.89%, it means that almost 88% of news articles will correctly be classified as real or fake. This can well establish its ability to model structural and temporal aspects of the graph, which makes the classification more accurate and context-aware. Therefore, the model has an accuracy of 87.91%, which makes it highly reliable in those cases where the model predicts that a news article is actually fake. The recall also shows the strong ability of this model, with 87.89%, as it can detect most of the fake news articles, minimizing false negatives. The F1-score, with 87.90%, keeps a good balance between them, thus confirming that the model maintains a strong balance between precision and recall.

B. Comparison with Other GNN Architectures

To demonstrate the superiority of our dynamic heterogeneous GNN model, we compared its performance against three other GNN architectures: static heterogeneous, dynamic homogeneous, and static homogeneous GNNs. The static models utilize Graph Attention Networks (GATs) to capture node dependencies, while the dynamic model employs a Graph Convolutional Long Short-Term Memory (GCLSTM) network to model temporal evolution.

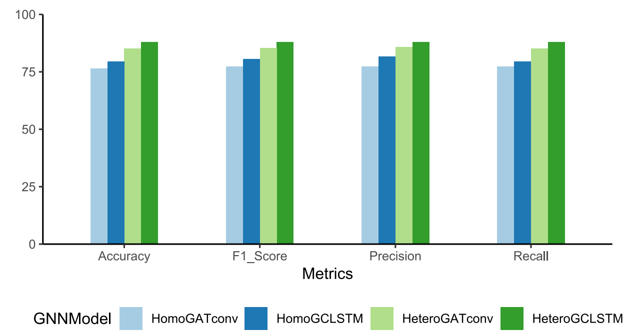

Fig. 4. Performance comparison of different GNN architectures

As shown in 4, the dynamic heterogeneous GNN demonstrates a notable advantage over the other models in terms of the evaluated metrics. The static heterogeneous GNN, while it effectively captures the diversity of node and edge types, appears to struggle with incorporating the temporal evolution of news propagation and user interactions. Conversely, the dynamic homogeneous GNN is capable of modeling temporal patterns; however, it cannot distinguish among various node and edge types, which may contribute to its comparatively lower precision and recall. The static homogeneous GNN, on the other hand, exhibits the least favorable performance, as it does not adequately address either structural heterogeneity or temporal dynamics, both of which are essential for effective fake news detection.

C. Real-time Fake News Detection

A key advantage of our approach is the ability to detect fake news as it evolves. Due to the support for dynamic graph updates, the model can make predictions at each temporal snapshot, allowing for early detection of misinformation before it fully propagates. We evaluated the model with a part of the initial snapshot graph to determine its effectiveness as the graph evolves. We tested for 40%, 60%, 80%, and 100% of the snapshots; the results demonstrate that the model performs well even with varying degrees of graph evolution. Fig 5 shows the performance metrics at different stages of graph evolution

As shown in Figure 5, our model achieved high accuracy even during the early stages of graph evolution. Specifically, the model attained an accuracy of 83% when the graph had evolved to only 40%, and this increased to 88% at the final stage of evolution. This indicates that our model is capable of early detection of fake news without requiring the graph to fully evolve, which is crucial for real-time fake news detection.

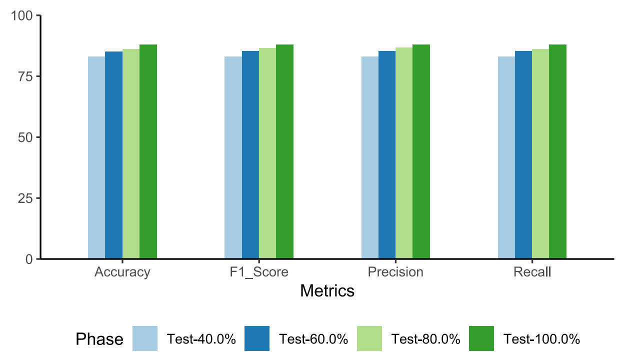

Fig. 5. Performance metrics at different stages

D. Error Analysis

The model faced notable challenges in three key categories: complex verification requirements, source credibility assessment, and statistical validation. False negatives were prevalent in scenarios requiring cross-referencing multiple sources, verifying technical claims, or addressing content complexity, such as policy verification and historical context. Conversely, false positives primarily stemmed from statistical misrepresentation and authority misattribution, where misleading statistics or incorrect statements about institutional policies were misclassified. Table II provides a sample of misclassified news stories, highlighting the nuanced nature of these errors.

TABLE II
MISCLASSIFIED NEWS STORIES

News Story
False Negatives (True stories classified as false)
No city in the state can quarantine itself without state approval.
Small trials to test convalescent plasma therapy for coronavirus patients seem to have had some degree of success
The United States is 'the oldest democracy' in the world.
False Positives (False stories classified as true)
There are 43 states that have now passed voter suppression laws.
Texas has lower taxes and less regulations than anywhere in the U.S.
The majority of minimum wage workers...are mothers.

VI. CONCLUSION

In this paper, we applied a Heterogeneous Graph Convolutional LSTM model for fake news detection, leveraging both the structural and temporal dynamics of news dissemination. By constructing a dynamic graph that evolves over time, we were able to capture complex interactions between news articles, users, and tweets.

The evaluation of the model on the TruthSeeker2023 dataset revealed that our model excels in various performance metrics. It achieved an accuracy of 87.89%, with precision, recall, and F1-score all exceeding 87%, demonstrating its effectiveness in accurately identifying fake news. Notably, the model performed well even as the graph evolved, showing its capability to provide reliable predictions at different stages of graph development. Furthermore, a comparative analysis with other GNN architectures underscored the advantages of our dynamic heterogeneous approach, highlighting its ability to capture both structural diversity and temporal dynamics. In future work, we plan to extend our model by integrating external knowledge sources and exploring optimization strategies to enhance scalability for larger and more complex datasets.

REFERENCES

[1] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *science*, vol. 359, no. 6380, pp. 1146–1151, 2018.
[2] D. M. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild *et al.*, "The science of fake news," *Science*, vol. 359, no. 6380, pp. 1094–1096, 2018.
[3] C. Shao, G. L. Ciampaglia, O. Varol, A. Flammini, and F. Menczer, "The spread of fake news by social bots," *arXiv preprint arXiv:1707.07592*, vol. 96, no. 104, p. 14, 2017.
[4] H. Allcott and M. Gentzkow, "Social media and fake news in the 2016 election," *Journal of economic perspectives*, vol. 31, no. 2, pp. 211–236, 2017.
[5] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *ACM SIGKDD explorations newsletter*, vol. 19, no. 1, pp. 22–36, 2017.
[6] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
[7] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
[8] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *arXiv preprint arXiv:1710.10903*, 2017.
[9] X. Zhou, J. Wu, and R. Zafarani, "Safe: similarity-aware multi-modal fake news detection (2020)," *Preprint. arXiv*, vol. 200304981, no. 2, 2020.
[10] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, "Temporal graph networks for deep learning on dynamic graphs," *arXiv preprint arXiv:2006.10637*, 2020.
[11] Z. Zhou, H. Guan, M. M. Bhat, and J. Hsu, "Fake news detection via nlp is vulnerable to adversarial attacks," *arXiv preprint arXiv:1901.09657*, 2019.
[12] H. Rashkin, E. Choi, J. Y. Jang, S. Volkova, and Y. Choi, "Truth of varying shades: Analyzing language in fake news and political fact-checking," in *Proceedings of the 2017 conference on empirical methods in natural language processing*, 2017, pp. 2931–2937.
[13] Z. Jin, J. Cao, H. Guo, Y. Zhang, and J. Luo, "Multimodal fusion with recurrent neural networks for rumor detection on microblogs," in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 795–816.
[14] K. Shu, S. Wang, and H. Liu, "Beyond news contents: The role of social context for fake news detection," in *Proceedings of the twelfth ACM international conference on web search and data mining*, 2019, pp. 312–320.
[15] C. Budak, D. Agrawal, and A. El Abbadi, "Limiting the spread of misinformation in social networks," in *Proceedings of the 20th international conference on World wide web*, 2011, pp. 665–674.
[16] F. Monti, F. Frasca, D. Eynard, D. Mannion, and M. Bronstein, "Fake news detection on social media using geometric deep learning. arxiv 2019," *arXiv preprint arXiv:1902.06673*, 2019.
[17] Y. Wang, Y. Ma, Z. Zhou, Y. Guan, and H. Liu, "Amf: Attention-based multi-modal fusion for fake news detection," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 2072–2080.
[18] K. Shu, L. Cui, S. Wang, D. Lee, and H. Liu, "defend: Explainable fake news detection," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 395–405.
[19] J. Ma, W. Gao, and K.-F. Wong, "Detecting rumors from microblogs with recurrent neural networks," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016, pp. 3818–3824.
[20] S. Dadkhah, X. Zhang, A. G. Weismann, A. Firouzi, and A. A. Ghorbani, "The largest social media ground-truth dataset for real/fake content: Truthseeker," *IEEE Transactions on Computational Social Systems*, 2023.
[21] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. [Online]. Available: https://arxiv.org/abs/1908.10084
[22] D. P. Kingma, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.