

Triple Path Enhanced Neural Architecture Search for Multimodal Fake News Detection

Bo Xu^{1†}, Qiuji Xie^{2†}, Jiahui Zhou³, Linlin Zong^{3*}

¹School of Computer Science and Technology, Dalian University of Technology ²School of Computer Science, Fudan University

³School of Software, Dalian University of Technology

{llzong, xubo}@dlut.edu.cn, qjxie22@m.fudan.edu.cn, zjhjixiang@mail.dlut.edu.cn

Abstract—Multimodal fake news detection has become one of the most crucial issues on social media platforms. Although existing methods have achieved advanced performance, two main challenges persist: (1) Under-performed multimodal news information fusion due to model architecture solidification, and (2) weak generalization ability on partial-modality contained fake news. To meet these challenges, we propose a novel and flexible triple path enhanced neural architecture search model MUSE. MUSE includes two dynamic paths for detecting partial-modality contained fake news and a static path for exploiting potential multimodal correlations. Experimental results show that MUSE achieves stable performance improvement over the baselines.

Index Terms—fake news detection, neural architecture search

I. INTRODUCTION

The emergence of social media has revolutionized information sharing, leading to an alarming surge in deliberately fabricated stories known as fake news [21]. This trend has been further exacerbated by the advent of large language models [9, 13, 27] which can generate a substantial volume of deceptive articles. Fake news often involves various modalities (text, images, and videos etc.), requiring the integration of multimodal information for more accurate detection. However, each modality contains distinct information needing specific analyses, like visual cues in images and linguistic indicators in text. Hence, effectively detecting multimodal fake news on social media remains a critical pursuit for researchers.

Existing fake news detection methods focus on learning effective semantic representations. For example, the attention mechanism [2, 11], modality consistency [18, 24], contrastive pretraining [25] were used to handle the complicated interactions of multimodal data. In addition, Wang et al. [16] introduced event classification; Qi et al. [10] proposed a novel framework using cross-modal correlations. While substantial advancements have been achieved in multimodal false news detection, practical challenges persist, necessitating further optimization of detection models. Firstly, diverse categories of fake news exhibit semantic conflicts and content distortions, as shown in Fig. 1. The integration of complementary information from different modalities remains an open question, as current



Fig. 1: Example for various types of fake news.

models often rely on fixed architectures. Secondly, in real-world scenarios, multimodal news may involve only a subset of available modalities, making it difficult for models trained on partial-modality to adapt to dynamically generated content.

To address these challenges, we propose the MULTimodal neural architecture Search Enhanced fake news detection model (MUSE). MUSE leverages Neural Architecture Search (NAS) techniques [1, 7, 23] to enhance the model's generalization capacity. Specifically, MUSE incorporates two dynamic paths and one static path, offering flexibility in dynamically modifying its network architecture to identify the most effective operators for extracting multimodal fusion features. Our contributions can be summarized as follows:

- MUSE leverages NAS techniques to dynamically adjust network architecture for improved feature fusion and generalization capacity.
- MUSE can adapt to diverse types of fake news. It consistently demonstrates superior performance compared to baselines, particularly for data lacking in modalities.

II. METHOD

A. General Framework

We focus on fake news detection involving text and image modalities. In a dataset $D = \{X_i, Y_i\}_{i=1}^N$ containing N news samples and corresponding labels, the sample X_i typically includes a textual segment T and an image V . To

[†] Equal contribution. ^{*} Corresponding author.

This work was supported in part by the Fundamental Research Funds for the Central Universities (DUT24MS003), and the Liaoning Provincial Natural Science Foundation Joint Fund Program(2023-MSBA-003).

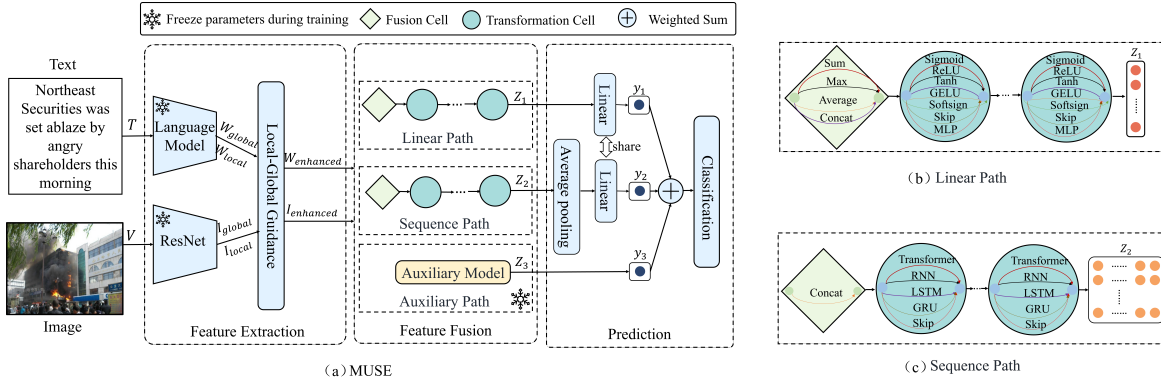


Fig. 2: Overall architecture of the MUSE model.

mirror real-world scenarios, we acknowledge the presence of incomplete data within D . Samples lack text are indicated as $X_i = \{\bar{T}, V\}$, while others lack images are expressed as $X_i = \{T, \bar{V}\}$. Given this context, the research problem entails identifying the mapping $f = \{T/\bar{T}, V/\bar{V}\} \xrightarrow{\theta} Y \in \{0, 1\}$, where θ denotes the model parameters, and $Y = 1$ and $Y = 0$ denote labels for fake and authentic news, respectively.

Given the diverse fake news types and the possibility of missing data, we introduce the MUSE model (Fig. 2(a)). MUSE contains three stages for fake news detection: feature extraction, feature fusion, and news prediction. Details are described in the following sections.

B. Multimodal Feature Extraction

We use the pre-trained language model to extract the text feature matrix, $W_{local} = \{w_1; \dots; w_{K_T}\} \in \mathbb{R}^{K_T \times D_T}$, where K_T is the number of words in T , and D_T is the dimension of feature vectors. Similarly, we use the pre-trained ResNet-50 [4] to extract the regional features of a given image V , namely $I_{local} = \text{ResNet}(V) = \{i_1; \dots; i_{K_V}\} \in \mathbb{R}^{K_V \times D_V}$, where K_V is the number of regions in the image, and D_V is the dimension of the feature vectors.

Since the raw features may contain inadequate information that condenses the context and high-level semantics, we adopt the Global-Local Guidance [12] to further refine the multimodal features. Taking the text features as an example, we use (1) and (2) to obtain the enhanced text feature $W_{enhanced}$, where $I_{global} \in \mathbb{R}^{D_V}$ is the fine-grained feature by applying the average pooling to I_{local} , \odot represents the Hadamard product, FC represents the fully connected layer, $\text{Norm}(\cdot)$ represents L2-normalization.

$$d = W_{local} \odot FC(I_{global}) \quad (1)$$

$$W_{enhanced} = (1 + \text{Norm}(d)) \odot W_{local} \quad (2)$$

Given that some news samples only contain partial modalities, we use (3) to calculate the $W_{enhanced}$ matrix of these samples. We obtain the enhanced image features $I_{enhanced}$ in the same way.

$$W_{enhanced} = \begin{cases} 0 & \text{when input} = \{\bar{T}, V\}, \\ W_{local} & \text{when input} = \{T, \bar{V}\} \end{cases} \quad (3)$$

C. Multimodal Feature Fusion

In this stage, we design two dynamic paths and one static path. The dynamic paths, using DARTS [7], to address the rigidity of model structure. Meanwhile, the static path preserves essential modality correlations throughout training.

1) *Dynamic Path in MUSE*: We optimize architecture with DARTS by adapting it from a multi-input-output graph to a single-input and single-output structure. This change ensures accuracy in detecting fake news, while also reducing the time required for architecture search.

In our modified approach, each dynamic path comprises a linked structure with n cells. Each cell, serving as the building block, searches for the most effective neural network for input data. These cells are categorized into fusion cells and transformation cells based on their search targets. Fusion cells take in enhanced sample features and search for the optimal architecture to fuse these features, while transformation cells search for deep neural architectures to uncover hidden semantic information within the fused features.

Mathematically, each intermediate cell $h^{(j)}$ represents a potential feature representation and connects to preceding cell $h^{(i)}$, $i < j$ through directed operator edges $e^{(i,j)}$. Therefore, each intermediate cell can be represented using (4).

$$h^{(j)} = \sum_{i < j} e^{(i,j)} h^{(i)} \quad (4)$$

To create a continuous search space, each edge $e^{(i,j)}$ is associated with a weight $\alpha_o^{(i,j)}$ for the operator o . The mixed operator applied to the feature representation h is represented by (5).

$$\delta^{(i,j)}(h) = \sum_{o \in O} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in O} \exp(\alpha_{o'}^{(i,j)})} o(h) \quad (5)$$

where O represents the set of candidate operators. Thus, the task simplifies to optimizing a set of continuous variables $\alpha_o^{(i,j)}$. After completing the architecture search, we select the most effective operator $o^{(i,j)}$ with the highest weight on each directed edge $e^{(i,j)}$ and discard other operators to obtain a discrete model, as demonstrated in (6).

$$o^{(i,j)} = \arg\max_{o \in O} \alpha_o^{(i,j)} \quad (6)$$

(1) *Linear path*. Considering different operators can effectively processing certain types of news, we design a linear

TABLE I: Comparison Results on the WEIBO and PHEME Datasets.

Model	WEIBO							PHEME						
	Acc.	Fake News			Real News			Acc.	Fake News			Real News		
		Precision	Recall	F1	Prec.	Recall	F1		Precision	Recall	F1	Precision	Recall	F1
GRU	0.702	0.671	0.794	0.727	0.747	0.609	0.671	0.832	0.782	0.712	0.745	0.855	0.896	0.865
CAMI	0.740	0.736	0.756	0.744	0.747	0.723	0.735	0.779	0.732	0.606	0.663	0.799	0.875	0.835
TextGCN	0.787	0.975	0.573	0.727	0.712	0.985	0.827	0.828	0.775	0.735	0.737	0.827	0.828	0.828
EANN	0.782	0.827	0.697	0.756	0.752	0.863	0.804	0.681	0.685	0.664	0.694	0.701	0.750	0.747
MVAE	0.824	0.854	0.769	0.809	0.802	0.875	0.837	0.852	0.806	0.719	0.760	0.871	0.917	0.893
att_RNN	0.772	0.854	0.656	0.742	0.720	0.889	0.795	0.850	0.791	0.749	0.770	0.876	0.899	0.888
SpotFake*	0.869	0.877	0.859	0.868	0.861	0.879	0.870	0.823	0.743	0.745	0.744	0.864	0.863	0.863
SpotFake+	0.870	0.887	0.849	0.868	0.855	0.892	0.873	0.800	0.730	0.668	0.697	0.832	0.869	0.850
HMCAN	0.885	0.920	0.845	0.881	0.856	0.926	0.890	0.871	0.860	0.781	0.818	0.876	0.924	0.900
MCAN	0.899	0.913	0.889	0.901	0.884	0.909	0.897	-	-	-	-	-	-	-
CAFE	0.840	0.855	0.830	0.842	0.825	0.851	0.837	-	-	-	-	-	-	-
FND-CLIP	0.907	0.914	0.901	0.908	0.914	0.901	0.907	-	-	-	-	-	-	-
MUSE	<i>0.905</i>	0.898	0.915	<i>0.906</i>	0.913	0.896	<i>0.904</i>	0.871	0.860	0.781	0.818	0.876	0.924	0.900
MUSE-discrete	0.914	0.905	0.926	0.916	0.924	0.903	0.913	0.878	0.836	0.838	0.837	0.902	0.902	0.902

operator search path to search for an effective linear neural network operator for one-dimensional feature representations. As shown in Fig. 2(b), the linear path consists of one fusion cell and $n - 1$ transformation cells. The search space of the fusion cell includes four operators: *Sum*, *Max*, *Average* and *Concat*. Taking the *Sum* operator as an example, the fused feature X_{fused} is calculated as follows:

$$X_{fused} = \text{Sum}(FC(W_{enhanced}), FC(I_{enhanced})) \quad (7)$$

The search space of the transformation cell is composed of five commonly used activation functions, namely *Sigmoid*, *ReLU*, *Tanh*, *GELU*, *Softsign*, together with the *Skip* and the *Multi-Layer Perception(MLP)* operator. Taking the *MLP* operator as an example, the outputs can be calculated as follows. After $n - 1$ transformation cells, we obtain the output vector Z_1 of the linear path.

$$X'_{fused} = \begin{cases} \text{Skip}(X_{fused}) \\ \dots \\ \text{MLP}(X_{fused}) \end{cases} \quad (8)$$

$$\text{MLP}(X_{fused}) = FC(\text{Sigmoid}(FC(X_{fused}))) \quad (9)$$

(2) Sequence path. To model multi-dimensional sequence features with richer inter-sequence information, we design a sequence path to search for complex architectures embedded with hidden news relationships. The sequence path consists of one fusion cell and $m - 1$ transformation cells. As shown in Fig. 2(c), the search space of the fusion cell consists of the multi-dimensional concatenation operator *Concat*, and the fused features X_{fused} after fusion cell is computed as follows.

$$X_{fused} = \text{Concat}(FC(W_{enhanced}), FC(I_{enhanced})) \quad (10)$$

For the transformation cell, we use the *Transformer*, *RNN*, *LSTM*, *GRU* and the *Skip* operator as the search space. After $m - 1$ transformation cells, we name the output vector of the sequence path Z_2 .

2) *Static Path in MUSE*: We incorporate a static path aimed at capturing latent details for detection, producing outputs Z_3 that represent the auxiliary information. This auxiliary path remains constant throughout the training process. In practical applications, this path can be predefined and optimized using non-parameter strategies tailored to different task objectives. For example, in datasets with quality concerns, such as the PHEME dataset[28], we address the issue of sample reuse by proposing a sample reference strategy in the static path. Here,

samples are clustered without supervision, and the average prediction score of similar samples within the same batch serves as the outputs Z_3 . We provide two ways to build the static path. For datasets characterized by complete multimodal relationships like the WEIBO dataset [5], we design a similarity Siamese Network as the auxiliary path. In datasets with quality concerns, such as the PHEME dataset [28], where samples are purposefully gathered around five social events, we address the issue of sample reuse by generating adequately similar samples. Here, samples are clustered without supervision, and the average prediction score of similar samples within the same batch serves as the auxiliary information.

TABLE II: Results on the WEIBO_Partial Dataset.

Model	Acc.	Fake News			Real News		
		Precision	Recall	F1	Precision	Recall	F1
MCAN	0.678	0.635	0.607	0.621	0.696	0.721	0.708
Baseline	0.690	0.655	0.640	0.648	0.717	0.730	0.723
MUSE	0.750	0.701	0.764	0.731	0.796	0.734	0.767

D. Fake News Prediction

We combine the output vectors Z_1 , Z_2 , and Z_3 from three different paths using a linear layer as the final classifier to fuse their classification vectors y_1 , y_2 , and y_3 , as shown in (11).

$$y_i = \beta * \text{Scale}(y_1) + \gamma * \text{Scale}(y_2) + \delta * \text{Scale}(y_3) \quad (11)$$

where β , γ and δ are learnable variables. The function $\text{Scale}(\cdot)$ scales the classification vector values of the three paths to a uniform interval. Finally, we use the binary cross entropy loss function as the objective function shown in (12).

$$\text{Loss}_i = -(Y_i * \log(y_i) + (1 - Y_i) \log(1 - y_i)) \quad (12)$$

where Y_i is the true label of the sample X_i , and y_i is the prediction probability of the sample X_i .

III. EXPERIMENTS AND RESULTS

A. Experimental Setup

We experiment on two datasets, WEIBO [5] and PHEME [28]. The Chinese WEIBO dataset consists of fake news from the microblog rumor dispelling system and the real news confirmed by Xinhua News Agency. There are 1.2% news without text and 16.4% news without image. The English PHEME dataset only contains news with missing images. Moreover, to simulate the harsh practical conditions, we extracted 1000 complete news samples from the WEIBO dataset

TABLE III: Results of Ablation Experiment on the WEIBO and PHEME Datasets for Operators.

Dataset	Number of Operators	Accuracy	Fake News			Real News		
			Precision	Recall	F1	Precision	Recall	F1
WEIBO	All Operators	0.904	0.875	0.943	0.908	0.938	0.865	0.900
	6 Operators	0.900	0.918	0.880	0.898	0.884	0.921	0.902
	5 Operators	0.904	0.911	0.897	0.904	0.898	0.912	0.905
	4 Operators	0.903	0.910	0.896	0.903	0.897	0.911	0.904
	3 Operators	0.911	0.913	0.909	0.911	0.909	0.913	0.911
	2 Operators	0.911	0.890	0.939	0.914	0.935	0.884	0.909
	1 Operators	0.914	0.905	0.926	0.916	0.924	0.903	0.913
PHEME	All Operators	0.866	0.865	0.758	0.808	0.866	0.930	0.897
	6 Operators	0.868	0.877	0.857	0.867	0.860	0.880	0.870
	5 Operators	0.871	0.870	0.769	0.816	0.871	0.931	0.900
	4 Operators	0.868	0.843	0.795	0.818	0.882	0.912	0.897
	3 Operators	0.872	0.842	0.811	0.826	0.900	0.909	0.899
	2 Operators	0.878	0.836	0.838	0.837	0.902	0.902	0.902
	1 Operators	0.847	0.817	0.762	0.789	0.864	0.898	0.881

TABLE IV: Ablation study of paths on WEIBO.

Ablation	Acc.	Fake News			Real News		
		Precision	Recall	F1	Precision	Recall	F1
w/o Linear	0.847	0.914	0.766	0.834	0.798	0.928	0.858
w/o Sequence	0.894	0.886	0.906	0.896	0.903	0.883	0.893
w/o Auxiliary	0.844	0.916	0.759	0.830	0.794	0.930	0.856
MUSE	0.905	0.898	0.915	0.906	0.913	0.896	0.904

TABLE V: Ablation study of paths on PHEME.

Ablation	Acc.	Fake News			Real News		
		Precision	Recall	F1	Precision	Recall	F1
w/o Linear	0.870	0.828	0.820	0.824	0.893	0.898	0.896
w/o Sequence	0.855	0.807	0.804	0.806	0.883	0.886	0.885
w/o Auxiliary	0.867	0.836	0.801	0.818	0.885	0.907	0.895
MUSE	0.871	0.860	0.781	0.818	0.876	0.924	0.900

to create a WEIBO_Partial dataset. We randomly remove one modality (text or image) from each news sample to make the whole dataset fully contain partial modality. Eventually, the WEIBO_Partial dataset consists of 500 News with missing texts and 500 news with missing images. We use 800 samples for training and 200 samples for testing.

We compare MUSE with single-modal and multimodal models. Single-modal models consist of GRU [8], CAMI [20], and TextGCN [19]. Multimodal models include EANN [16], MVAE [6], att_RNN [5], SpotFake* [14], SpotFake+ [15], HMCAN [11], MCAN [17], CAFE [3] and FND-CLIP [25]. Moreover, we discretize the search results to obtain the discretized prediction model MUSE-discrete following DARTS.

B. Performance Comparison

Table I shows that MUSE and its discrete variant MUSE-discrete excel in Accuracy and F1 score on WEIBO, surpassing all the baselines. On PHEME, MUSE achieved comparable performance with HMCAN, possibly due to data quality issues such as label imbalance and sample reuse in PHEME. MUSE's advanced performance stems from its adaptive architecture and efficient noise reduction in incomplete data.

Table II shows the results on the WEIBO_Partial dataset. Here we select 1) MCAN [17], a model with slightly weaker detection performance than MUSE on the WEIBO dataset, and 2) a simple baseline model that concatenates the text and image modalities using two linear layers as the compared models. The results show that with the rate of partial modality

contained news increases (16.4% to 50%), the performance difference between these models becomes more obvious. As MCAN focuses on exploiting inter/intra-modal relations, it incorrectly amplifies the noises in the low-quality samples, making the detection performance even 1.2% lower than a simple baseline model. On the contrary, MUSE exhibits strong adaptability to data quality and outperforms all baseline models in terms of Accuracy and F1 score.

C. Ablation Experiments for Operators

To figure out whether the number of retained operators during a discrete process affects the detection performance. According to the size of the weights, we remove the operator with the lowest weight, and then retrain the model using the same settings. The results obtained are shown in Table III.

We can find that reducing the number of retained operators usually increases the model's detection performance. This is because the model needs suitable rather than complex operators. Adding more operators introduces noise into the features when the existing operators are already sufficient for the data. In experiments, we choose to retain one and two operators for the WEIBO and PHEME datasets, respectively.

D. Ablation Experiments for Paths

We conduct the ablation experiments on each path: (a) w/o Linear: MUSE with the linear path removed, (b) w/o Sequence: MUSE with the sequence path removed and (c) w/o Auxiliary: MUSE with the auxiliary path removed. Results in Table IV and Table V indicate that removing any path from MUSE will degrade the detection performance, showing the effectiveness of each path.

IV. CONCLUSION

In this work, we propose a triple path enhanced neural architecture search model MUSE for fake news detection. MUSE incorporates NAS into fake news detection to provide a solution to dynamically adjusting the model architecture according to data characteristics. The experiments demonstrate the effectiveness of MUSE in detecting fake news. In the future, we plan to investigate the application of MUSE in resource-constrained scenarios [22, 26].

REFERENCES

- [1] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45, 2024.
- [2] T. Chen, X. Li, H. Yin, and J. Zhang. Call attention to rumors: Deep attention based recurrent neural networks for early rumor detection. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 40–52. Springer, 2018.
- [3] Y. Chen, D. Li, P. Zhang, J. Sui, Q. Lv, L. Tun, and L. Shang. Cross-modal ambiguity learning for multi-modal fake news detection. In *Proceedings of the ACM Web Conference 2022*, pages 2897–2905, 2022.
- [4] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [5] Z. Jin, J. Cao, H. Guo, Y. Zhang, and J. Luo. Multimodal fusion with recurrent neural networks for rumor detection on microblogs. In *ACM Multimedia*, pages 795–816, 2017.
- [6] D. Khattar, J. S. Goud, M. Gupta, and V. Varma. Mvae: Multimodal variational autoencoder for fake news detection. In *The web conference*, pages 2915–2921, 2019.
- [7] H. Liu, K. Simonyan, and Y. Yang. DARTS: differentiable architecture search. In *ICLR*, 2019.
- [8] J. Ma, W. Gao, P. Mitra, S. Kwon, B. J. Jansen, K. Wong, and M. Cha. Detecting rumors from microblogs with recurrent neural networks. In *IJCAI*, pages 3818–3824, 2016.
- [9] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023. doi: 10.48550/ARXIV.2303.08774.
- [10] P. Qi, J. Cao, X. Li, H. Liu, Q. Sheng, X. Mi, Q. He, Y. Lv, C. Guo, and Y. Yu. Improving fake news detection by using an entity-enhanced framework to fuse diverse multimodal clues. In *ACM Multimedia*, pages 1212–1220, 2021.
- [11] S. Qian, J. Wang, J. Hu, Q. Fang, and C. Xu. Hierarchical multi-modal contextual attention network for fake news detection. In *SIGIR*, pages 153–162, 2021.
- [12] L. Qu, M. Liu, J. Wu, Z. Gao, and L. Nie. Dynamic modality interaction modeling for image-text retrieval. In *SIGIR*, pages 1104–1113, 2021.
- [13] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen, and X. Wang. A comprehensive survey of neural architecture search: Challenges and solutions. *ACM Computing Surveys (CSUR)*, 54(4):1–34, 2021.
- [14] S. Singhal, R. R. Shah, T. Chakraborty, P. Kumaraguru, and S. Satoh. Spotfake: A multi-modal framework for fake news detection. In *2019 IEEE fifth international conference on multimedia big data*, pages 39–47, 2019.
- [15] S. Singhal, A. Kabra, M. Sharma, R. R. Shah, T. Chakraborty, and P. Kumaraguru. Spotfake+: A multi-modal framework for fake news detection via transfer learning. In *AAAI*, pages 13915–13916, 2020.
- [16] Y. Wang, F. Ma, Z. Jin, Y. Yuan, G. Xun, K. Jha, L. Su, and J. Gao. Eann: Event adversarial neural networks for multi-modal fake news detection. In *KDD*, pages 849–857, 2018.
- [17] Y. Wu, P. Zhan, Y. Zhang, L. Wang, and Z. Xu. Multimodal fusion with co-attention networks for fake news detection. In *Findings of ACL*, pages 2560–2569, 2021.
- [18] J. Xue, Y. Wang, Y. Tian, Y. Li, L. Shi, and L. Wei. Detecting fake news by exploring the consistency of multimodal data. *Information Processing & Management*, 58(5):102610, 2021.
- [19] L. Yao, C. Mao, and Y. Luo. Graph convolutional networks for text classification. In *AAAI*, pages 7370–7377, 2019.
- [20] F. Yu, Q. Liu, S. Wu, L. Wang, T. Tan, et al. A convolutional approach for misinformation identification. In *IJCAI*, pages 3901–3907, 2017.
- [21] Z. Zeng, M. Wu, G. Li, X. Li, Z. Huang, and Y. Sha. An explainable multi-view semantic fusion model for multimodal fake news detection. In *ICME*, pages 1235–1240, 2023.
- [22] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.
- [23] M. Zheng, X. Su, S. You, F. Wang, C. Qian, C. Xu, and S. Albanie. Can gpt-4 perform neural architecture search? *arXiv preprint arXiv:2304.10970*, 2023.
- [24] X. Zhou, J. Wu, and R. Zafarani. Safe: Similarity-aware multi-modal fake news detection. In *PAKDD*, pages 354–367, 2020.
- [25] Y. Zhou, Y. Yang, Q. Ying, Z. Qian, and X. Zhang. Multimodal fake news detection via clip-guided learning. In *ICME*, pages 2825–2830, 2023.
- [26] L. Zong, Q. Xie, J. Zhou, P. Wu, X. Zhang, and B. Xu. Fedcmr: Federated cross-modal retrieval. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1672–1676, 2021.
- [27] L. Zong, J. Zhou, W. Lin, X. Liu, X. Zhang, and B. Xu. Unveiling opinion evolution via prompting and diffusion for short video fake news detection. In L. Ku, A. Martins, and V. Srikumar, editors, *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 10817–10826. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.FINDINGS-ACL.642.
- [28] A. Zubiaga, M. Liakata, and R. Procter. Exploiting context for rumour detection in social media. In *International conference on social informatics*, pages 109–123. Springer, 2017.