



FuDFEND: Fuzzy-Domain for Multi-domain Fake News Detection

Chaoqi Liang, Yu Zhang^(✉), Xinyuan Li, Jinyu Zhang, and Yongqi Yu

Harbin Institute of Technology, Harbin, China

{120L030706, 7203610518, 7203610523, 7203610515}@stu.hit.edu.cn,
zhangyu@ir.hit.edu.cn

Abstract. On the Internet, fake news exists in various domain (e.g., education, health). Since news in different domains has different features, researchers have begun to use single domain label for fake news detection recently. Existing works show that using single domain label can improve the accuracy of fake news detection model. However, there are two problems in previous works. Firstly, they ignore that a piece of news may have features from different domains. The single domain label focuses only on the features of one domain. This may reduce the performance of the model. Secondly, their model cannot transfer the domain knowledge to the other dataset without domain label. In this paper, we propose a novel model, FuDFEND, which solves the limitations above by introducing the fuzzy inference mechanism. Specifically, FuDFEND utilizes a neural network to fit the fuzzy inference process which constructs a fuzzy domain label for each news item. Then, the feature extraction module uses the fuzzy domain label to extract the multi-domain features of the news and obtain the total feature representation. Finally, the discriminator module uses the total feature representation to discriminate whether the news item is fake news. The results on the Weibo21 show that our model works better than the model using only single domain label. In addition, our model transfers domain knowledge better to Thu dataset which has no domain label.

Keywords: Fake news detection · Social media · Multi-domain

1 Introduction

With the development of the Internet, social media platforms such as Sina Weibo and Twitter have become the main source of information. Fake news is also widely spread on social platforms. The fake news on social media is usually a short text, a picture and a short video that they can be understood in a few seconds. At the same time, the fake news incites people's emotions and stimulates users to forward, so they can be widely spread. Tavernise [1] showed that important social events were affected due to moderated fake news campaigns. The spread of fake news may result in people's panic and social dislocation. The quality of content on social media platforms has suffered greatly due to the spread of fake news, misinformation and unverifiable facts [2]. Therefore, it is meaningful to research into social media fake news detection today.

At present, there are two methods of social media fake news detection. One is social background detection method, and the other is content-based detection method [3]. Social background method aims to study users' social network structure, users' personal information, microblog forwarding and reply relationship and rumor's propagation patterns. The content-based research method aims to detect the text, voice and video carried by microblog news. For example, Mouratidis et al. [4] analyzes linguistic features to distinguish between fake news and real news. Our work focuses on the content-based method.

On the Internet, fake news arises in many different domains. The task of detecting fake news in multiple domains is called Multi-domain Fake News Detection (MFND). There are two main challenges in MFND. First, the performance of techniques generally drops if news records are coming from different domains (e.g., politics, entertainment), and a possible explanation of this could be the rather unique content and style of each domain [5]. For example, during the 2020 US election, political fake news and Covid19 fake news were widespread simultaneously. It is difficult to detect both political fake news and Covid19 fake news at the same time, because of the different or even conflicting features between political fake news and Covid19 news. Thus, a variety of previous works [6–10] focused on rumor detection in a single domain. But herein lies another problem. In a single-domain, there may be too little data to train a good model. Therefore, Nan et al. [11] proposed a Multi-domain Fake News Detection Model (MDFEND) which use single domain label to solve the above two problems. Single domain label is to describe that a piece of news belongs to a certain domain, such as science and technology, education, health and so on. The model receives news content and single domain label as input. Utilizing these data, the model can extract the common features of fake news in all different domains, and can distinguish the specific features of certain domain by domain label.

Table 1. The difference between single domain label and fuzzy domain label.

News content	Single domain label	Fuzzy domain label
The new generation of iphone has added numbers of new technologies such as face recognition, which are popular with consumers. The next day, Apple shares rose sharply	Technology	Technology: 70% Business: 20% Finance: 10%

We build this paper on that work [11] by recognizing that there are two problems. (i) A piece of news may have features of several domains. There is an example showed in Table 1. This news has the features from three domains. Single domain label can't help the model extract this kind of news with multi domain features. (ii) Their model can't be transferred to datasets without domain label. Therefore, we propose a novel model, Fuzzy-domain Fake News Detection Model (FuDFEND). FuDFEND is an improved model based on the MDFEND proposed in [11]. The improvement is made by introducing a fuzzy inference mechanism into the model. The fuzzy inference mechanism can solve the above two problems. The fuzzy inference mechanism constructs a fuzzy domain

label for each news item. Compared with single domain label, fuzzy domain label can better describe the domain features of news, so that help the model better extract the multi-domain features of news. We demonstrated this by the experiment on the weibo21. In addition, FuDFEND has better transfer ability than MDFEND. We will illustrate this through experiment on Thu dataset.

Our summarized contributions with this paper are:

- We propose a novel model, Fuzzy-domain Fake News Detection Model (FuDFEND), which can extract the multi-domain features of news content by the fuzzy domain label that fuzzy inference mechanism generate. The results on the Weibo21 show that FuDFEND works better than the model, MDFEND, using only single domain label.
- In order to describe the multi-domain features of news, we introduced fuzzy inference mechanism to the multi-domain fake news detection task. The fuzzy domain label constructed by fuzzy inference mechanism can more accurately describe the multi-domain features that news has.
- We solve the problem that model can't transfer to the dataset without domain label. We evaluate FuDFEND on Thu dataset which has no domain label. Experimental results show that FuDFEND can transfer domain knowledge better by utilizing the fuzzy inference mechanism.

2 Related Work

2.1 Fake News Detection Methods

Researchers have come up with number of ways to detect fake news.

Content-Based Methods. In [12], the authors proposed an ensemble classification model for fake news detection. Their model obtains relevant features from a fake news dataset and then uses an ensemble model to classify the extracted features, but these works are based only on news texts. In [13], the authors study and analyze whether fake news can be distinguished from mainstream news by text writing style and whether fake news can be detected only by writing style by building a model. Rawat et al. [2] proposed a method to automatically collect fake news detection tasks online. For each piece of news data, they collect evidence and generate their summaries as another input to the model to help detect news text. In [14] author found images have different distribution patterns and statistically distinctive patterns for fake and real news. It reveals that images in the real news are more diverse and much denser than those in the fake news. Therefore, they extract image features from the visual features and overall statistics of images in news events. Alonso-Bartolome et al. [15] used the early fusion method to fuse text and image information for rumor detection. In [16], an innovative RNN with attention mechanism (att-RNN) is proposed for effective multimodal feature fusion. In [17], author propose a Multi-domain Visual Neural Network (MVNN) framework consisting of three main parts to mix all the features together.

Social Background Detection Method. These researches mainly focused on propagation patterns, publisher-news relations and user-news interactions. In [18], the authors examined numerous features related to user, linguistic, network and temporal characteristics of rumors. They studied the spread patterns of rumors over time and the ability to track the precise changes in the predictive power of rumor features. However, the events of this volume (111 rumor and non-rumors) may not be sufficient for summarizing the prediction performance, and the algorithm takes every feature into account did not achieve good results in short observation periods. Shu et al. [19] examined the correlation of publisher bias, news stance, and relevant user engagement. They studied the novel problem of exploiting social context for fake news detection and proposed a tri-relationship embedding framework TriFN, which model relations and interactions between publishers, news, and users simultaneously for fake news classification. But it requires the social context information to be included in the data. Meanwhile, there may be significant differences in social relations in different domains, and targeted detection in different domains is a promising concept. Alrubaian et al. [20] proposed a credibility analysis system consisting of a model based on reputation, a feature ranking algorithm, a credibility evaluation classifier engine, and a user expertise model, to assess the accuracy of information on Twitter so it can stop the promotion of disinformation.

2.2 Multi-domain Rumor Task

In [5], the author introduced two new English fake news datasets (FakeNewsAMT, Celebrity) covering six news domains, and the author analyzes the different performance of his model in different news domains. Nan et al. [11] construct Weibo21, a multi-domain fake news dataset in Chinese. Weibo21 contains more than 9000 pieces of real and fake news from nine domains and was manually marked the domain label for each piece of news. And they designed a Multi-domain Fake News Detection Model (MDFEND) for Weibo21. The model can extract the common features of fake news in all different domains, and can distinguish the specific features of certain domain through single domain label. However, this model relies on manually annotated datasets, which not only contain news text and true and false information, but also need to input the domain category to which the news belongs, then select the corresponding vector in the domain matrix according to the type of domain to calculate the weight of each domain experts (TextCNN), so it requires a lot of effort to obtain the data set that match the model. In the following we will propose our improvement strategy on the base of this model.

3 FuDFEND: Fuzzy-Domain Fake News Detection Model

FuDFEND is an improved model based on the MDFEND proposed in [11]. The improvement is made by introducing a fuzzy inference mechanism into the model. The fuzzy inference mechanism includes two modules: Membership Function and Gate.




Fig. 1. Overall framework of FuDFEND.

Overall framework of FuDFEND is showed in Fig. 1. The working process of the model is as follows. Enter the news text into BERT [21, 22] to obtain a series of word embeddings $\mathbf{W} = [w_{[CLS]}, w_1, w_2, \dots, w_n, w_{[SEP]}]$. $\mathbf{W}$ is input into a mixture of experts to extract the features of different domains. Enter $\mathbf{W}$ into Membership Function to obtain fuzzy domain label. Then, Gate generates weight scores α by inputting fuzzy domain label. The output experts are weighted and summed by the weight scores α , so that obtain the total feature representation $\mathbf{v}$. The classifier module uses the total feature representation $\mathbf{v}$ to discriminate whether the news item is fake news.

3.1 Membership Function

In traditional set theory, we define a set with a definite condition. For example, we define a tall set as those who are taller than or equal to 190 cm. The other belong to short set. If a person is 189 cm tall, very close to 190 cm but still belongs to short set. Obviously, the traditional way of describing sets is very crude. In order to describe the set more accurately, LA Zadeh proposed the concept of fuzzy sets in 1965. A fuzzy set is characterized by a Membership (characteristic) Function which assigns to each object a membership grade between zero and one [23]. As shown in Fig. 2, we can use a Membership Function to describe high set and short set. If a person is 189 cm tall, he belongs to a high set with a membership grade of 0.9 and belongs to a short set with a membership grade of 0.1.




Fig. 2. An example of fuzzy set.

A piece of news may have features of different domains. Here is an example that “The new generation of iphone has added a number of new technologies such as face recognition, which are popular with consumers. The next day, Apple shares rose sharply.” This news has the features of both science domain and financial domain. Therefore, we need to use fuzzy sets to more precisely measure the domains that news belongs to.

We will show how to use neural network to fit the Membership Function for the sets of news domains. The Membership Function consist of an *GRU*, a multi-layer perception (*MLP*) and a softmax function. Membership Function can generate a nine-dimensional membership grade vector $\mathbf{g}$. We call this vector $\mathbf{g}$ as the fuzzy domain label. During training, we train it as training a classifier for news domains. When we use it, we regard it as a Membership Function. The specific process is as follows. For a piece of news, we put the news’ word embeddings $\mathbf{W}$ into an *GRU*. The output $\mathbf{h}$ of *GRU* is fed into *MLP* which output is a nine-dimensional vector. Use softmax function to normalize the output of *MLP*. Then we obtain a vector $\mathbf{g}$ which has nine dimensions and sum of all dimensions is 1. The value of each dimension of the vector represents the membership grade of nine news domains (Science, Military, Education, Disasters, Politics, Health, Finance, Entertainment). We denote the Membership Function as $M(\cdot; \cdot)$, and $\theta, \theta_1, \theta_2$ is the parameters in the Membership Function, *GRU*, *MLP*. $\mathbf{g}$ represents the fuzzy domain label:

$$\mathbf{g} = M(\mathbf{W}; \theta) = \text{softmax}(\text{MLP}(\text{GRU}(\mathbf{W}; \theta_1); \theta_2)) \quad (1)$$

Membership Function’s parameter is pre-trained. During the training of rumor detection task, the parameter of Membership Function is fixed. Membership Function is train as a classifier, its accuracy rate can only reach about 82% in the end. But as a Membership Function, We cannot think that its error rate is 18%, but should be understood as the correction of labels by large models. Because some domain labels labeled by the human are not necessarily suitable for machine. The machine reclassified things that it thought had the same features. Further, the outputs of Membership Function, $\mathbf{g}$, is a probability distribution, which can provide enough noise for the model.

Figure 3 shows an execution of the Membership Function. Membership Function analysis shows that the content of the news has the features of entertainment domain and social domain.




Fig. 3. An execution of the Membership Function. The single domain label of this news is “Society”. The content of this news is that “# Dad said his daughter is not a vegetative person but a sleeping beauty # when I saw this topic, I was moved by the story of Mizuho in *The House Where the Mermaid Sleeps*. Mizuho’s mother always believed that Mizuho could hear what she said. She dressed and dressed Mizuho every day. Mizuho was like falling asleep. Of course, at the end of the book, when choosing between life and death, she chose “living” to donate Mizuho’s heart. This is the first time I realized the human body.”

3.2 Feature Extraction

We use Mixture-of-Expert [24–26] to extract features of news’ content. Each expert network is a TextCNN [27] in our model. Each expert has its own area of expertise and is good at extracting the features of certain domain. An expert network can be denoted by $E_i(\cdot; \cdot)$ ($1 \leq i \leq T$). T is a hyperparameter, representing the number of the experts. r_i denote the output of the expert. φ_i are the trainable parameters of an expert:

$$r_i = E_i(\mathbf{W}; \varphi_i) = \text{TextCNN}_i(\mathbf{W}; \varphi_i) \quad (2)$$

3.3 Domain Gate

Because different experts can extract the features of different domains, we input the fuzzy domain label into the Domain Gate, to obtain the weight score. The weight score consists of T positive real numbers, corresponding to T experts. Then we use the weight score to aggregate the features extracted by experts. We denote the Domain Gate as $\text{Gate}(\cdot; \cdot)$. ψ is the trainable parameters of the Domain Gate. α represent the weight score:

$$\alpha = (\alpha_1, \alpha_2, \dots, \alpha_T) = \text{Gate}(\mathbf{g}; \psi) = \text{softmax}(\text{MLP}(\mathbf{g}; \psi)) \quad (3)$$

3.4 Fake News Prediction and Loss Function

Through the weight score, we aggregate the features extracted by experts to obtain the total features vector v :

$$v = \sum_{i=1}^T \alpha_i r_i \quad (4)$$

Input the total features vector into a binary classifier which is MLP with a sigmoid output layer. And then we get our model's predicted value, $\hat{y}$. The value $\hat{y}$ is between zero and one and it is the probability that the news is fake. The larger the value of $\hat{y}$ is, the more likely the model is to conclude that the news is fake:

$$\hat{y} = \text{sigmoid}(MLP(\mathbf{v}, \xi)) \quad (5)$$

ξ represent the trainable parameters of the MLP.

We use y_i to represent value label and $\hat{y}_i$ to represent predicted value. We employ Binary Cross-Entropy Loss for training:

$$L = - \sum_{i=1}^N (y_i \log \hat{y}_i + (1 - y_i) \log (1 - \hat{y}_i)) \quad (6)$$

4 Experiment

FuDFEND is an improved model based on the MDFEND proposed in [11]. The improvement is made by introducing a fuzzy inference mechanism into the model. In order to show the effect of fuzzy inference mechanism, we compare our model FuDFEND with baseline method which mentioned in [11], especially MDFEND which hasn't fuzzy inference mechanism.

4.1 Dataset

In this section, we introduce the two datasets we use.

Weibo21 was released in this work [11] last year. It is a Chinese multi-domain fake news dataset. It consists of 4,488 pieces of fake news and 4,640 pieces of real news, covering nine different news domains: **Science, Military, Education, Accidents, Politics, Health, Finance, Entertainment**. The number of true and fake news contained in each domain is shown in Table 2.

Table 2. Data Statistics of Weibo21.

Domain	Science	Military	Education	Accidents	Politics
Real	143	121	243	185	306
Fake	93	222	248	591	546
All	236	343	491	776	852
Domain	Health	Finance	Entertainment	Society	All
Real	485	959	1000	1198	4640
Fake	515	362	440	1471	4488
All	1000	1321	1440	2669	9128

Table 3. Five Samples of Weibo21.

content	Single domain label	Fake label
【三星.....年推出柔性折叠屏的产品，你会入手么？	Science	0
#四川.....化工厂泄漏#紧急撤离！！大家注意安全	Accidents	1
#广州.....料公布！居然是日本间谍！】@广东热搜	Politics	1
《哈利.....面都是回忆[悲伤]http://t.cn/RURkC7V	Entertainment	0
#人民.....。求@平安北京@通州警方在线给说法！！	Society	1

Each piece of news is labeled a specific domain. We show a few examples of the Weibo21 dataset in Table 3.

Another dataset we used is the Chinese Rumor Dataset released by thunlp [28, 29]. The news in this dataset has no domain label. The dataset consists of two parts. One is called rumors_V170613 [28], including 31669 rumors. Another named CEO_Dataset [29], including 1538 rumors and 1849 real news. We found that some rumors in CEO_Dataset coincide with the data in Weibo21, so we randomly selected 1560 pieces of fake news from rumors_V170613 and combined the fake news with the real news in the CEO_Dataset to form a new dataset that we called Thu dataset. We performed zero-shot learning on Thu dataset to test the model’s capability of domain knowledge transfer.

4.2 Experiment Setting

BERT [21, 22] is fixed during training. The dimension of BERT’s output embedding vectors is fixed to 768. The max length of the sentence is 170. We employ the Adam [30] optimizer for our training. The learning rate is $5e-4$. In order to enhance the credibility of our experiments, the process is performed for 5 times and the average f1-score is reported.

4.3 Train Membership Function and FuDFEND

As mentioned previously, we trained the Membership Function as a nine-classification classifier for domain label. We randomly selected 70% of the data from Weibo21 as the training set and the remain as the verification set. The mini-batch size is 32. After training 7 epochs using the crossentropy loss function, the f1-score on the validation set reaches 82.31%. We take the Membership Function from the 7th epoch training, because it has the highest f1-score. As training FuDFEND, the Membership Function is fixed. The mini-batch size is 64.

4.4 Experiment on Weibo21

Nan et al. [11] have showed the effectiveness of their multi-domain fake news detection model (MDFEND). Further, the results in Table 4 support that our idea is right. The fake news detection performance of FuDFEND compared with MDFEND demonstrates

that the fuzzy domain label generated by fuzzy inference mechanism can better help the model extract the features for fake news detection task.

Table 4. Multi-domain Fake News Detection Performance on Weibo21 (F1-score).

Model	Science	Military	Education	Accidents	Politics
TextCNN_single	0.7470	0.7780	0.8882	0.8310	0.8694
BiGRU_single	0.4876	0.7169	0.7067	0.7625	0.8477
BERT_single	0.8192	0.7795	0.8136	0.7885	0.8188
TextCNN_all	0.7254	0.8839	0.8362	0.8222	0.8561
BiGRU_all	0.7269	0.8724	0.8138	0.7935	0.8356
BERT_all	0.7777	0.9072	0.8331	0.8512	0.8366
EANN	0.8225	0.9274	0.8624	0.8666	0.8705
MMOE	0.8755	0.9112	0.8706	0.8770	0.8620
MOSE	0.8502	0.8858	0.8815	0.8672	0.8808
EDDFN	0.8186	0.9137	0.8676	0.8786	0.8478
MDFEND	0.8301	0.9389	0.8917	0.9003	0.8865
FuDFEND (our)	0.8133	0.9468	0.8917	0.9059	0.9013
Model	Health	Finance	Entertainment	Society	All
TextCNN_single	0.9053	0.7909	0.8591	0.8727	0.8380
BiGRU_single	0.8378	0.8109	0.8308	0.6067	0.7342
BERT_single	0.8909	0.8464	0.8638	0.8242	0.8272
TextCNN_all	0.8768	0.8638	0.8456	0.8540	0.8686
BiGRU_all	0.8868	0.8291	0.8629	0.8485	0.8595
BERT_all	0.9090	0.8735	0.8769	0.8577	0.8795
EANN	0.9150	0.8710	0.8957	0.8877	0.8975
MMOE	0.9364	0.8567	0.8886	0.8750	0.8947
MOSE	0.9179	0.8672	0.8913	0.8729	0.8939
EDDFN	0.9379	0.8636	0.8832	0.8689	0.8919
MDFEND	0.9400	0.8951	0.9066	0.8980	0.9137
FuDFEND (our)	0.9417	0.8901	0.9161	0.9174	0.9213

4.5 Experiment on Thu Dataset

The news in Thu dataset hasn't domain label. FuDFEND has Membership Function module, so it only needs to enter news content and does not need to enter domain label. To demonstrate the transfer learning capabilities of our models, we use MDFEND as a

baseline and used the hyperparameters provided in [11] to train MDFEND. MDFEND requires domain label, so we use Membership Function to label the single domain label for the dataset. The results are showed in Table 5. The results demonstrate that FuDFEND has good transfer ability with the help of fuzzy inference mechanism.

Table 5. Multi-domain Fake News Detection Performance on Thu dataset (F1-score).

Model	Science	Military	Education	Accidents	Politics
MDFEND	0.8157	0.8407	0.7995	0.6878	0.7938
FuDFEND (our)	0.8760	0.8301	0.8568	0.7129	0.8148
Model	Health	Finance	Entertainment	Society	All
MDFEND	0.8364	0.7535	0.8047	0.8642	0.8731
FuDFEND (our)	0.8753	0.7930	0.8492	0.8653	0.8900

5 Conclusion

Previous work has demonstrated that the use of single domain label can effectively improve the performance of fake news detection models. However, for news containing multi-domain features, models using single domain label do not synthesize multi-domain features well. In this work, we propose FuDFEND and provide a set of experiments to demonstrate that: (i) Fuzzy domain label more accurately portrays the multi-domain features of news content, so it can help the model to better extract the multi-domain features of news. (ii) Fuzzy inference mechanisms can be very helpful for model to transfer domain knowledge to dataset without domain label.

6 Future Work

We have received valuable comments from reviewers. For example, analyze good and bad cases and test more datasets and indicators. These suggestions give us great inspiration. However, due to time constraints and limited space, we cannot solve these problems one by one in this paper. We will study them in our future work.

References

1. Tavernise, S.: As fake news spreads lies, more readers shrug at the truth. The New York Times. Accessed 26 Jan 2017
2. Rawat, M., Kanojia, D.: Automated Evidence Collection for Fake News Detection. [arXiv: 2112.06507](#) (2021)
3. Hangloo, S., Arora, B.: Fake news detection tools and methods — a review. Int. J. Adv. Innovat. Res. **8**(2) (IX), 100–108 (2021)

4. Mouratidis, D., Nikiforos, M.N., Kermanidis, K.L.: Deep learning for fake news detection in a pairwise textual input schema. *Computation* **9**(2), 20 (2021)
5. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., et al.: Automatic detection of fake news. In: Association for Computational Linguistics, pp. 3391–3401 (2017)
6. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on Twitter. In: WWW, pp. 675–684 (2011)
7. Jin, Z., Cao, J., Guo, H., Zhang, Y., Wang, Y., Luo, J.: Detection and analysis of 2016 US presidential election related rumors on Twitter. In: Lee, D., Lin, Y.-R., Osgood, N., Thomson, R. (eds.) SBP-BRiMS 2017. LNCS, vol. 10354, pp. 14–24. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-60240-0_2
8. Kwon, S., et al.: Prominent Features of Rumor Propagation in Online Social Media, pp. 1103–1108. IEEE (2013)
9. Ma, B., Lin, D., Cao, D.: Content representation for microblog rumor detection. In: Angelov, P., Gegov, A., Jayne, C., Shen, Q. (eds.) Advances in Computational Intelligence Systems. AISC, vol. 513, pp. 245–251. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-46562-3_16
10. Ma, J., et al.: Detecting rumors from microblogs with recurrent neural networks. *IJCA*. 3818–3824 (2016)
11. Nan, Q., et al.: MDFEND: Multi-domain Fake News Detection, pp. 3343–3347. ACM (2022)
12. Hakak, S., Alazab, M., Khan, S., Gadekallu, T.R., Maddikunta, P.K.R., Khan, W.Z.: An ensemble machine learning approach through effective feature extraction to classify fake news. *Futur. Gen. Comput. Syst.* **117**, 47–58 (2021)
13. Potthast, M., Kiesel, J., Reinartz, K., et al.: A stylometric inquiry into hyperpartisan and fake news. In: Association for Computational Linguistics, pp. 231–240 (2018)
14. Jin, Z., Cao, J., Zhang, Y., et al.: Novel visual and statistical image features for microblogs news verification. *IEEE Trans. Multim.* **19**(3), 598–608 (2016)
15. Alonso-Bartolome, S., Segura-Bedmar, I.: Multimodal Fake News Detection. [arXiv:2112.04831](https://arxiv.org/abs/2112.04831) (2021)
16. Jin, Z.W., et al.: Multimodal Fusion with Recurrent Neural Networks for Rumor Detection on Microblogs, pp. 795–816. ACM (2017)
17. Qi, P., Cao, J., Yang, T., et al.: Exploiting Multi-domain Visual Information for Fake News Detection, pp. 518–527. IEEE (2019)
18. Kaist, K.S., Kaist, C.M., Snu, J.K.: Rumor detection over varying time windows. *PLoS ONE* **12**(1), e0168344 (2017)
19. Shu, K., Wang, S., Liu, H.: Beyond News Contents: The Role of Social Context for Fake News Detection, pp. 312–320. ACM (2019)
20. Alrubaian, M., et al.: A credibility analysis system for assessing information on Twitter. *IEEE Trans. Dependab. Secure Comput.* **15**(4), 661–674 (2016)
21. Jacob, D., et al.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Association for Computational Linguistics, pp. 4171–4186 (2019)
22. Cui, Y., et al.: Revisiting Pre-trained Models for Chinese Natural Language Processing. In: Association for Computational Linguistics, pp. 657–668 (2020)
23. Zadeh, L.A.: Fuzzy sets. *Inf. Control* **8**(3), 338–353 (1965)
24. Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, G.E.: Adaptive mixtures of local experts. *Neural Comput.* **3**(1), 79–87 (1991)
25. Ma, J., Zhao, Z., Yi, X., et al.: Modeling Task Relationships in Multi-task Learning With Multi-gate Mixture-of-Experts, pp. 1930–1939. ACM (2018)
26. Zhu, Y., Liu, Y., Xie, R., et al.: Learning to Expand Audience via Meta Hybrid Experts and Critics for Recommendation and Advertising, pp. 4005–4013. ACM (2021)
27. Kim, Y.: Convolutional neural networks for sentence classification. In: Empirical Methods in Natural Language Processing (EMNLP), pp. 1746–1751 (2014)

28. Liu, Z.Y., Zhang, L., Cunchao, T.U., et al.: Statistical and semantic analysis of rumors in Chinese social media. *Sci. Sin. Inf.* **45**(12), 1536–1546 (2015)
29. Song, C., Tu, C., Yang, C., et al.: CED: credible early detection of social media rumors. *IEEE Trans. Knowl. Data Eng.* **33**(8), 3035–3047 (2018)
30. Kingma, D., Ba, J.: Adam: a method for stochastic optimization. *Comput. Sci. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)* (2014)