



Improving Generalization for Multimodal Fake News Detection

Sahar Tahmasebi
sahar.tahmasebi@tib.eu
TIB – Leibniz Information Centre for
Science and Technology
Hannover, Germany

Sherzod Hakimov
sherzod.hakimov@uni-potsdam.de
University of Potsdam
Potsdam, Germany

Ralph Ewerth*
Eric Müller-Budack*
ralph.ewerth@tib.eu
eric.mueller@tib.eu
L3S Research Center, Leibniz
University Hannover
Hannover, Germany

ABSTRACT

The increasing proliferation of misinformation and its alarming impact have motivated both industry and academia to develop approaches for fake news detection. However, state-of-the-art approaches are usually trained on datasets of smaller size or with a limited set of specific topics. As a consequence, these models lack generalization capabilities and are not applicable to real-world data. In this paper, we propose three models that adopt and fine-tune state-of-the-art multimodal transformers for multimodal fake news detection. We conduct an in-depth analysis by manipulating the input data aimed to explore models performance in realistic use cases on social media. Our study across multiple models demonstrates that these systems suffer significant performance drops against manipulated data. To reduce the bias and improve model generalization, we suggest training data augmentation to conduct more meaningful experiments for fake news detection on social media. The proposed data augmentation techniques enable models to generalize better and yield improved state-of-the-art results.

CCS CONCEPTS

• Information systems → Multimedia information systems;

KEYWORDS

Multimodal fake news detection, social media, news analytics

ACM Reference Format:

Sahar Tahmasebi, Sherzod Hakimov, Ralph Ewerth, and Eric Müller-Budack. 2023. Improving Generalization for Multimodal Fake News Detection. In *International Conference on Multimedia Retrieval (ICMR '23)*, June 12–15, 2023, Thessaloniki, Greece. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3591106.3592230>

1 INTRODUCTION

Misinformation and fake news that contain false information to deliberately deceive readers [1], have become a pressing challenge for society. Typically, fake news use different modalities such as images,

*Also affiliated with TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICMR '23, June 12–15, 2023, Thessaloniki, Greece
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0178-8/23/06.
<https://doi.org/10.1145/3591106.3592230>

text, and videos to attract attention. Thus, automated *multimodal* solutions that incorporate information from multiple modalities are essential to detect misinformation on social media or in news.

Multimodal research on fake news detection has been growing rapidly, and current approaches [9, 18–20] have mainly used deep learning to extract features from text and image. For example, Wang et al. [19] introduced an *Event Adversarial Neural Network (EANN)*, which extracts event-invariant features to detect fake news on unforeseen events. *Spotfake* [18] demonstrated the effectiveness of combining pre-trained language models [5] and computer vision approaches [15]. Zhang et al. [20] added a domain classifier and proposed a *BERT*-based (Bidirectional Encoder Representations from Transformers) domain adaptation neural network (*BDANN*) for multimodal fake news detection. These approaches achieve impressive results but are trained on rather small datasets with similar characteristics between train and test data. However, due to the dynamic nature of news and realistic use cases for social media, where the news could be affected by human-produced false content, they mostly fail in real-world scenarios. This failure mostly arises from overfitting to input data resulting in a lack of generalization.

In this paper, we address the aforementioned issues and present the following contributions. (1) We propose three multimodal models that adapt state-of-the-art transformers for fake news detection in social media. Experimental results on the *MediaEval Benchmark* for the *Verifying Multimedia Use* task [2, 3] demonstrate the effectiveness of these models compared to the state of the art. (2) We suggest several content manipulation techniques to provide new settings for training and evaluation and conduct an in-depth analysis of model performance in more meaningful test scenarios that confirm poor generalization of previous approaches. (3) We suggest training data augmentation using the large-scale *Visual News* dataset [12] that contain news from a wide range of topics to improve generalization. (4) Furthermore, we introduce an ensemble approach that combines models trained with various manipulation techniques to further improve generalization. Experimental results have demonstrated that it provides the best overall results across various test sets. Source code, models, and dataset are publicly made available on *GitHub*¹ to allow for fair comparisons.

The remainder of this paper is organized as follows. The multimodal models and data augmentation for fake news detection are proposed in Section 2. The experimental setup, results and a revised training setup to improve generalization are presented in Section 3. Section 4 concludes this paper and provides future work directions.

¹*GitHub* repository: <https://github.com/TIBHannover/MM-FakeNews-Detection>

Table 1: Distribution of training, validation, and test samples.

Dataset	Category	Train	Validation	Test	All
MediaEval 2015 [2]	Fake	5,245	1,479	717	13,480
	Real	3,426	1,398	1,215	
MediaEval 2016 [3]	Fake	5,962	1,479	630	14,578
	Real	4,641	1,398	468	

2 MULTIMODAL FAKE NEWS DETECTION

In this section, we propose data augmentation strategies to improve model generalization as well as three models based on recent transformers for multimodal fake news detection.

2.1 Dataset Augmentation

As mentioned before, datasets for multimodal fake news detection [2, 3] are typically rather small and contain limited topics. As a consequence, supervised models tend to overfit which limits their performance on unseen data. While large-scale news dataset that cover a huge variety of topics are available, they usually contain only *real* articles and are not labeled for multimodal fake news detection. We propose to use various content manipulation techniques, such as the replacement of entities (e.g., events) or images, to automatically create *fake* samples to train deep learning approaches for multimodal fake news detection with better generalization capabilities. Details on the manipulations techniques and training strategies including an ensemble model are provided in Section 3.3.

2.2 Multimodal Transformers Models

We propose three multimodal approaches for fake news detection to predict whether a post is *fake* ($y = 0$) or *real* ($y = 1$). They consist of three modules: (1) A modal-specific encoder that extracts multimodal features from both text \mathbf{x}_T and image \mathbf{x}_I , (2) a fusion module that combines the features from both modalities, and (3) a classification module that outputs class probabilities. The goal is to find the best model $\mathcal{F}(\mathbf{x}_T, \mathbf{x}_I) \rightarrow y \in \{0, 1\}$.

2.2.1 BERT-ResNet Model. This model adopts pre-training techniques to encode each modality. For the textual part T , it uses BERT [5] to obtain text embedding \mathbf{x}_T as it achieves good results for various natural language processing tasks. We use a residual network with 50 layers (*ResNet-50*) [6] pre-trained on *ImageNet* [14] to extract image representations \mathbf{x}_I . By using two fully-connected (FC) layers of size 256 and *GELU* (*Gaussian Error Linear Unit*) [7] activation, we transform the embeddings of both modalities to a joint space ($\mathbf{x}_T, \mathbf{x}_I \in \mathbb{R}^{256}$). The embeddings are concatenated and passed to a FC layer with size 128 and *ReLU* (*Rectified Linear Unit*) activation. Finally, a FC classification layer with two neurons and *sigmoid* activation function is used to output probabilities for *fake* and *real*.

2.2.2 MLP-CLIP Model. This model extracts multimodal features using the *ViT-B/32* variant of *CLIP* (contrastive language-image pre-training) [13], which has proven to be powerful for many downstream tasks. The *CLIP* model outputs two embeddings of the same size for image \mathbf{x}_I and text \mathbf{x}_T . Each embedding is fed into a multi-layer perceptron (*MLP*) comprising two FC layers with sizes 256

Table 2: Precision (P), recall (R), and F1-score (F1) of our models and the state of the art. Results in gray serve as reference but are not directly comparable due to different experimental setups (e.g. train and test splits). Reproduced results using official implementations are denoted with * and are directly comparable. Results in bold are the best on our test splits.

	MediaEval (ME) 2015			MediaEval (ME) 2016		
	P	R	F1	P	R	F1
SpotFake [18]	-	-	-	0.791	0.753	0.760
MVAE [10]	-	-	-	0.745	0.748	0.744
BDANN [20]	-	-	-	0.820	0.780	0.795
Singhal et al. [17]	-	-	-	0.836	0.832	0.830
EANN [19]	0.847	0.812	0.829	-	-	-
Singh et al [16]	0.867	0.844	0.848	-	-	-
CAFE [4]	0.806	0.806	0.806	-	-	-
Spofake*	0.831	0.824	0.827	0.751	0.723	0.727
BDANN*	0.780	0.711	0.723	0.700	0.652	0.599
BERT-ResNet	0.870	0.873	0.872	0.634	0.631	0.632
MLP-CLIP	0.926	0.943	0.933	0.807	0.806	0.806
CLIP-MMBT	0.860	0.673	0.677	0.689	0.644	0.640

and 128 and *ReLU* activation function. We concatenate the embeddings and feed them into a FC layer with *sigmoid* activation and two neurons to output class probabilities for *fake* and *real*.

2.2.3 CLIP-MMBT Model. This approach leverages a multimodal fusion method called *MultiModal BiTransformers (MMBT)* [11] for classification. Pre-trained *ResNet-50x4* variant of *CLIP* [13] and *BERT* [5] are used to create representation for image \mathbf{x}_I and text \mathbf{x}_T , respectively. Both representations are passed through *MMBT* for multimodal fusion that uses a FC layer with one output neuron and *sigmoid* activation as the last layer to output a probability $\rho \in [0, 1]$ that is used to predict the class *fake* ($\rho < 0.5$) or *real* ($\rho \geq 0.5$).

3 EXPERIMENTAL SETUP AND RESULTS

This section provides details on the experimental setting (Section 3.1), a comparison to the state of the art (Section 3.2), and impact of content manipulation including result analysis (Section 3.3).

3.1 Experimental Setting

3.1.1 Implementation details: We tried multiple hyperparameters for all models and provide the best setups and models on *GitHub*¹. All models are trained for 30 epochs with a batch size of 32 and the best model with the lowest validation loss is used for testing.

3.1.2 Datasets: We use the 2015 and 2016 *MediaEval Benchmark* for the *Verifying Multimedia Use* task [2, 3]. To the best of our knowledge, they are the only publicly available datasets for fake news detection containing both image and text of the tweets written in English. Both datasets have a development and test set, each with its own events. Each tweet contains text with corresponding images or video. Given our focus on image content, we remove tweets with videos and tweets for which associated images are not available anymore. Table 1 shows the dataset statistics. Data splits are provided on *GitHub*¹ to allow for fair comparisons.



Figure 1: Manipulation techniques and results of (a) Spotfake and (b) MLP-CLIP (VNME-Ens). Images are replaced with similar ones due to licensing issues. Links to the original images are provided on [GitHub](#)¹.

3.2 Comparison to the State of the Art

Table 2 presents a comparison of our proposed models with the state of the art. Although related work has been evaluated on the same dataset, different sizes of subsets for training, validation, and test were used. Reasons are that there are some samples whose images are no longer available or specific data filtering methods (e.g., based on text length). As a consequence, results are reported with different experimental setups and *not directly comparable*. Nevertheless, we include results from related work for reference.

To ensure a fair evaluation, we reproduce results for *BDANN* [20] and *Spotfake* [18] since they provide applicable official [GitHub](#) implementations. We choose the best hyperparameters reported in their paper. The *MLP-CLIP model* achieves the best results among our models and the reproduced models from the state of the art for both datasets. This suggests that the *CLIP model* learns expressive visual and textual representations, which can be used in multimodal fake news detection. Table 2 also shows that the reproduced results for *Spotfake* and *BDANN* are worse compared to the reported ones (gray background) that have been obtained with a different experimental setup (e.g., train and test split). Their lower performance implies that our experimental setup is more challenging.

3.3 Impact of Content Manipulation

Despite good test performances, model generalization is questionable due to the scarcity of real events and limited number of unique images in the *ME* datasets. This can lead classifiers to learn domain-specific features instead of making an actual distinction between *fake* and *real*. In this section, we analyze fake news detection models in various use cases with manipulated samples. As proposed in Section 2.1, we apply data augmentation to avoid overfitting during model training to improve generalization.

3.3.1 Manipulation Techniques. To test model generalization, we selected 100 posts from *ME 2015* labeled as *real* and applied two types of text and image manipulation techniques to generate new test sets. Figure 1 depicts an example of each manipulation.

Event Replacement. We randomly replace events that are already provided as meta information in the *ME* dataset with other events from the dataset. This induces false information, which changes the ground truth of all samples from *real* to *fake*.

Table 3: Number of train, validation (valid.), and test samples for different manipulation strategies (MS) used to compose different versions of the VNME dataset.

Data	Split	MS	Fake	Real	VNME		
					(Img)	(Evt)	(All)
Visual News (VN)	Train	-	0	258,488	✓	✓	✓
	Train	EvRep	245,384	0	×	✓	✓
	Train	FakeIm	258,487	0	✓	×	✓
	Valid.	-	0	98,374	✓	✓	✓
	Valid.	EvRep	80,537	0	×	✓	✓
	Valid.	FakeIm	98,373	0	✓	×	✓
Media-Eval 2015 (ME)	Train	-	5,245	3,426	✓	✓	✓
	Train	EvtRep	2,923	0	×	✓	✓
	Train	FakeIm	3,425	0	✓	×	✓
	Valid.	-	1,398	1,479	✓	✓	✓
	Valid.	EvtRep	0	1,267	×	✓	✓
	Valid.	FakeIm	1,479	0	✓	×	✓
	Test	-	717	1,215	✓	✓	✓

Event Removal: Starting from the same initial set, we automatically remove all events from the text. As the ground truth can be both *real* or *fake*, one expert with computer science background manually annotated the validity of all 100 samples with removed events and decided on the label *real* or *fake* as follows. The sample remains *real* if (1) the text is still in line with the image, and (2) the meaning of text has not changed. Otherwise, it is annotated as *fake*.

Replacement with Fake Image: We randomly replace images with other images depicting a different event in the test set. Since this produces semantically uncorrelated image-text pairs, the ground truth for these samples is *fake*.

Replacement with Real Image. We replace the input image with another similar image from the test set. We manually check the chosen image to ensure it is from same event with similar content. Thus, all samples remain *real* after the manipulation.

3.3.2 Training Strategies for Improved Generalization. The *ME* dataset contains only few real events and limited number of unique images. Thus, we extend the dataset using real images and their

Table 4: Accuracy (Acc) as well as number of samples predicted as fake (N_F) and real (N_R) for different models and test data manipulations (number of fake / real ground-truth samples). Models denoted with ‡ are solely trained on ME 2015. Note that models with * are reproduced and that VNME-Ens is an ensemble of MLP-CLIP models trained on VNME.

Method	ME 2015	Original			FakeIm			RealIm			EvtRep			EvtRem			Total
	717 / 1,215	0 / 100		100 / 0		0 / 100		100 / 0		6 / 94		6 / 94		206 / 294			
	Acc	N_F	N_R	Acc	N_F	N_R	Acc	N_F	N_R	Acc	N_F	N_R	Acc	N_F	N_R	Acc	Acc
BDANN*, ‡	0.76	16	84	0.84	12	88	0.12	15	85	0.85	19	81	0.19	17	83	0.77	0.55
Spotfake*, ‡	0.84	37	63	0.63	30	70	0.30	18	82	0.82	37	63	0.37	37	63	0.61	0.54
BERT-ResNet, ‡	0.87	28	72	0.72	25	75	0.25	21	79	0.79	28	72	0.28	28	72	0.68	0.54
CLIP-MMBT, ‡	0.75	3	97	0.97	10	90	0.10	2	98	0.98	4	96	0.04	4	96	0.90	0.59
MLP-CLIP, ‡	0.93	27	73	0.73	40	60	0.40	31	69	0.69	51	49	0.51	39	61	0.41	0.54
• VNME-Img	0.69	3	97	0.97	90	10	0.90	5	95	0.95	24	76	0.24	16	84	0.80	0.77
• VNME-Evt	0.70	6	94	0.94	20	80	0.20	19	81	0.81	75	25	0.75	47	53	0.51	0.64
• VNME-All	0.68	21	79	0.79	100	0	1.00	0	100	1.00	61	39	0.61	38	62	0.60	0.80
• VNME-Ens	0.70	6	94	0.94	100	0	1.00	3	97	0.97	62	38	0.62	35	65	0.63	0.83

associated captions from *VisualNews* (VN) [12] since it includes much more samples from many domains, topics, and events. To create *fake* samples, we apply the aforementioned data manipulation techniques as indicated in Table 3. Considering that events are not provided in VN, we apply *spaCy* [8] for named entity recognition to detect named entities and their types. Each entity is replaced randomly with another entity of the same type to realize *Event Replacement* (*EvtRep*). All samples without any detected named entities have been removed since they cannot be manipulated.

We train our best model (*MLP-CLIP*) using training data with different manipulations to evaluate their impact. As indicated in Table 3, the *MLP-CLIP* model (*VNME-All*) is trained with all training splits, while *MLP-CLIP* (*VNME-Evt*) and *MLP-CLIP* (*VNME-Img*) models are trained with the original data and splits manipulated with *Event Replacement* and *Fake Image Replacement*. Furthermore, considering that ensemble models tend to have more reliable and better outcome, we evaluate the *MLP-CLIP* (*VNME-Ens*) model that combines the outputs of the previous models by majority voting. The idea is to utilize all three models which are trained by particular manipulation techniques to reduce the overall error rate.

3.3.3 Results. Table 4 shows a comparison of the original and manipulated sets. The performance of the models trained exclusively on ME drastically decreases for all manipulated test variants except *RealIm* where the ground truth remains *real*. For the remaining test sets, the ground truth changes to *fake* in most cases, but models tend to predict the same labels for the original and manipulated samples indicating that they are not sensitive to these manipulations (Figure 1). *MLP-CLIP* is the only model that changed predictions towards the correct label. However, its performance still drops compared to original test set especially on manipulated sets (*FakeIm*, *EvtRep*) where the ground truth changes to *fake*.

In contrast, our proposed *MLP-CLIP* models that are trained with additional data show significantly better results. As expected, models trained with datasets including text (*VNME-Evt*) or image manipulations (*VNME-Img*) show improved results for the respective modifications. We conclude that a modality-specific data manipulation technique in the training process improves the model generalization and robustness for the respective modality. For example, for

MLP-CLIP (*VNME-Img*), a significant number of samples which were originally predicted as *real* are then predicted as *fake* for *FakeIm* test set. The model that is trained with all modifications (*VNME-All*) provides the best overall performance among single models. It is only outperformed by the ensemble model *MLP-CLIP* (*VNME-Ens*), which shows that combining models for certain manipulations can improve the performance compared to baselines using a single model, i.e., *Spotfake* [18] (Figure 1). While it does not have the best results for each manipulated set, it provides the best overall performance averaged over all test sets which makes it the most reliable model for fake news detection. Overall, the results confirm that more diverse data and the application of data manipulation techniques in the training process improves model generalization.

4 CONCLUSION

In this paper, we have studied the problem of multimodal fake news detection and proposed three multimodal models using state-of-the-art transformer architectures that outperformed baselines on the 2015 and 2016 *MediaEval Benchmark* for the *Verifying Multimedia Use* task. Furthermore, we have applied text and image manipulation techniques to create more diverse test scenarios, which demonstrated that current approaches tend to overfit and lack generalization due to the limited size and variations of the training data. We have proposed a solution for this issue by adding more training data from different topics and by applying data manipulation techniques to create diverse samples. The models trained on these datasets have shown better generalization capabilities. In future work, we will explore different kinds of manipulation techniques, fusion strategies to combine different models, and test our approach on larger annotated test sets with more diverse topics.

ACKNOWLEDGMENTS

This work was funded by European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement no. 812997 (CLEOPATRA ITN), and by the German Federal Ministry of Education and Research (BMBF, FakeNarratives project, no. 16KIS1517).

REFERENCES

- [1] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives* 31, 2 (2017), 211–36.
- [2] Christina Boididou, Katerina Andreadou, Symeon Papadopoulos, Duc-Tien Dang-Nguyen, Giulia Boato, Michael Riegler, and Yiannis Kompatsiaris. 2015. Verifying Multimedia Use at MediaEval 2015. 1436 (2015). <http://ceur-ws.org/Vol-1436/Paper4.pdf>
- [3] Christina Boididou, Symeon Papadopoulos, Duc-Tien Dang-Nguyen, Giulia Boato, Michael Riegler, Stuart E. Middleton, Andreas Petlund, and Yiannis Kompatsiaris. 2016. Verifying Multimedia Use at MediaEval 2016. 1739 (2016). http://ceur-ws.org/Vol-1739/MediaEval_2016_paper_3.pdf
- [4] Yixuan Chen, Dongsheng Li, Peng Zhang, Jie Sui, Qin Lv, Lu Tun, and Li Shang. 2022. Cross-modal Ambiguity Learning for Multimodal Fake News Detection. In *The ACM Web Conference 2022, WWW 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*. ACM, 2897–2905. <https://doi.org/10.1145/3485447.3511968>
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019*. Association for Computational Linguistics, 4171–4186. <https://doi.org/10.18653/v1/n19-1423>
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. IEEE Computer Society, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [7] Dan Hendrycks and Kevin Gimpel. 2016. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415* (2016).
- [8] Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. *To appear* 7, 1 (2017), 411–420.
- [9] Zhiwei Jin, Juan Cao, Han Guo, Yongdong Zhang, and Jiebo Luo. 2017. Multimodal Fusion with Recurrent Neural Networks for Rumor Detection on Microblogs. In *ACM on Multimedia Conference, MM 2017, Mountain View, CA, USA, October 23-27, 2017*. ACM, 795–816. <https://doi.org/10.1145/3123266.3123454>
- [10] Dhruv Khattar, Jaipal Singh Goud, Manish Gupta, and Vasudeva Varma. 2019. MVAE: Multimodal Variational Autoencoder for Fake News Detection. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*. ACM, 2915–2921. <https://doi.org/10.1145/3308558.3313552>
- [11] Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, and Davide Testuggine. 2019. Supervised Multimodal Bitransformers for Classifying Images and Text. In *Visually Grounded Interaction and Language (ViGIL), NeurIPS 2019 Workshop, Vancouver, Canada, December 13, 2019*. <https://vigilworkshop.github.io/static/papers/40.pdf>
- [12] Fuxiao Liu, Yinghan Wang, Tianlu Wang, and Vicente Ordonez. 2020. Visual News : Benchmark and Challenges in Entity-aware Image Captioning. *CoRR abs/2010.03743* (2020). [arXiv:2010.03743](https://arxiv.org/abs/2010.03743) <https://arxiv.org/abs/2010.03743>
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning, ICLR 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 8748–8763. <http://proceedings.mlr.press/v139/radford21a.html>
- [14] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal on Computer Vision* 115, 3 (2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [15] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. <http://arxiv.org/abs/1409.1556>
- [16] Bhuvanesh Singh and Dilip Kumar Sharma. 2022. Predicting image credibility in fake news over social media using multi-modal approach. *Neural Comput. Appl.* 34, 24 (2022), 21503–21517. <https://doi.org/10.1007/s00521-021-06086-4>
- [17] Shivangi Singhal, Tanisha Pandey, Saksham Mrig, Rajiv Ratn Shah, and Ponnuram Kumaraguru. 2022. Leveraging Intra and Inter Modality Relationship for Multimodal Fake News Detection. In *Companion of The Web Conference 2022, Virtual Event / Lyon, France, April 25 - 29, 2022*. ACM, 726–734. <https://doi.org/10.1145/3487553.3524650>
- [18] Shivangi Singhal, Rajiv Ratn Shah, Tanmoy Chakraborty, Ponnuram Kumaraguru, and Shin'ichi Satoh. 2019. SpotFake: A Multi-modal Framework for Fake News Detection. In *IEEE International Conference on Multimedia Big Data, BigMM 2019, Singapore, September 11-13, 2019*. IEEE, 39–47. <https://doi.org/10.1109/BigMM.2019.00-44>
- [19] Yaqing Wang, Fenglong Ma, Zhiwei Jin, Ye Yuan, Guangxu Xun, Kishlay Jha, Lu Su, and Jing Gao. 2018. EANN: Event Adversarial Neural Networks for Multi-Modal Fake News Detection. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*. ACM, 849–857. <https://doi.org/10.1145/3219819.3219903>
- [20] Tong Zhang, Di Wang, Huanhuan Chen, Zhiwei Zeng, Wei Guo, Chunyan Miao, and Lizhen Cui. 2020. BDANN: BERT-Based Domain Adaptation Neural Network for Multi-Modal Fake News Detection. In *International Joint Conference on Neural Networks, IJCNN 2020, Glasgow, United Kingdom, July 19-24, 2020*. IEEE, 1–8. <https://doi.org/10.1109/IJCNN48605.2020.9206973>