FNR: A Similarity and Transformer-Based Approach to Detect Multi-Modal Fake News in Social Media

Faeze Ghorbanpour

Department of Computer Engineering, Sharif University of Technology

Maryam Ramezani

Department of Computer Engineering, Sharif University of Technology

MohammadAmin Fazli

Department of Computer Engineering, Sharif University of Technology

Hamid R. Rabiee

Department of Computer Engineering, Sharif University of Technology

Abstract—The availability and interactive nature of social media have made them the primary source of news around the globe. The popularity of social media tempts criminals to pursue their immoral intentions by producing and disseminating fake news using seductive text and misleading images. Therefore, verifying social media news and spotting fakes is crucial. This work aims to analyze multi-modal features from texts and images in social media for detecting fake news. We propose a Fake News Revealer (FNR) method that utilizes transform learning to extract contextual and semantic features and contrastive loss to determine the similarity between image and text. We applied FNR on two real social media datasets. The results show the proposed method achieves higher accuracies in detecting fake news compared to the previous works.

■ **THE INTRODUCTION** The proliferation of social media has made it the first source of news for many people. News is published on social networks every day, and people, willingly or unwillingly, spread it from one channel to

another, from person to person, and react to it. According to the study in [1], the source of most Americans' news is through the Internet rather than newspapers or radio, and three-quarters of people said they receive news via email or social

Figure 1: Twitter posts about April 15, 2013 Boston terrorist attack

media. Additionally, 37 percent of people share news on Facebook and Twitter, suggesting that these two media have created a new way of sharing news. Social media news differs from other news sources like news agencies or microblogs. Usually, social media news is written by ordinary people, using informal language, being brief, and having poor-quality images. Our focus is on the news published on social media.

The two characteristics of fake news are: intentionally written and provable to be false, which distinguishes it from other kinds of news like rumors, satires and spams [2]. Ease of use, sharing, and disseminating news by the public in social media can lead criminals to produce and publish fake news with malicious intent. These types of fraud can be perpetrated to mislead the public, harm an institution, person, government, or harm public and private stock markets. Meanwhile, the lack of awareness of the falsity of the news causes ordinary people to publish without knowledge. The scope of the news becomes wider and causes more significant damage, leading to distrust and disregard for the correct news and warnings work. It makes it difficult for reporters and journalists to cover correct and important news.

Visual misinformation tends to involve much more straightforward deceptions. It is prevalent for old photos and videos to be repurposed as evidence of recent events. Using out-of-context images is another way to validate the news and gain trust. This means that they use the photo of the correct news in another event or context for fake news. Images play a pivotal role in influencing public opinion and creating false perceptions. According to psychological research

[3], when individuals see an image alongside a trivial statement such as turtles are deaf, they are more likely to believe it. In a simulated social media environment, a post that incorporates photos get more likes and shares, as well as people's perception that it is factual [4]. Hence, considering the visual features of news and the similarity relations between its image and text is essential for recognizing the truth of the news.

In social media, we do not always know the source of the news, but we can find the reaction of users against the news post. On Twitter, for example, people retweet or comment on tweets with or without intent. We only see retweets and comments, so determining their reliability is critical. These reactions can create fake news and divert attention from the main news, which prevents people from paying attention to the real story. For example, Figure 1 shows some tweets related to a terrorist attack in Boston in 2013. As you can see, fake tweets are used to mislead readers, divert news, or further their agendas by exploiting the public's attention. In fake tweet Figure 1c by resembling a terrorist's backpack to another person, he tried to introduce another person as a terrorist. In Figure 1d, by using the image of the actual news next to another news about a missing person, the missing person was introduced as a terrorist.

As a means to deal with these challenges, we propose an end-to-end framework referred to as Fake News Revealers (FNR): A similarity and transformer-based approach to detect multi-modal fake news in social media. In this approach, we featurize text using BERT, a language representation model, which stands for Bidirectional

Encoder Representations from Transformers [5]. In addition, we use the ViT, a vision transformer model that makes use of patches of images as tokens in an NLP application and provides the sequence of linear embeddings of these patches as inputs to the transformer [6]. After these two embedding modules, we use two projection modules that project extracted features into a similar-sized array and tune weights based on the task and dataset. To optimize the model, we have used two loss functions: contrastive loss, a supervised loss between image and text, and classification loss, a cross-entropy loss between the predicted and actual label of each news item.

Our main contributions are:

- Utilizing the transformer models for both image and text, which yields better results than other text or image classification models in fake news detection.
- Using contrastive loss between the image and text to explore the relations between them for each news item.
- Outperforming the state-of-the-art multi-modal fake news detection methods on two public social media datasets.

RELATED WORKS

With the expansion of using social networks, automatic detection of fake news has become essential. The intentional nature of fake news and its adverse effects and implications have encouraged more researchers to focus on this issue.

In this section, we categorize the related works based on the modes of their inputs. Then, we focus on the recent multi-modals fake news methods.

Single Modality

In early works, only one mode of data was used to detect fake news, with textual data receiving the most attention due to its prevalence in the news. Linguistic features were utilized in [2] to validate news on Twitter, and structural and cognitive features were extracted to detect fake news on social networks in [7]. These methods, particularly those that utilize linguistic features, are topic-specific and cannot be generalized to all topics. Additionally, the methods in [2], and [7], do not extract features automatically, resulting

in insufficient and out of proportion solutions. The authors in [8] and [9], implemented machine learning models on social media images to detect fake news.

Using recurrent deep networks, the authors in [10], pioneered the use of deep models for fake news detection, demonstrating improved results.

Multiple Modality

The authors in [11] employed a visual system to answer questions via deep networks for fake news detection using multi-modal data. Alternatively, [12] used subtitle texts in addition to images for detection of fake news.

EANN [13] is an adversarial neural network that uses image and text in the news. It tries to solve the independent identification of the news events challenge by reducing the event's impact, and the news occurrence with an adversarial mechanism. The goal is to generalize the solution to unexpected events. This model uses the pre-trained VGG19 network models for the image, and a deep canonization network for textual properties. MVAE [14] presents a variational auto-encoder and an encoder-decoder network to detect fake news using the learned hidden vectors. A deep BiLSTM network is used to extract textual features, and a pre-trained model is used to extract image features [13].

SpotFake [15] detects fake news by embedding text and images in vectors and then fusing these vectors. CARMN [16] utilizes a cross model attention mechanism to calculate the relationship between image and text. Then, using a self-attention mechanism it obtains the feature vectors and determine fake news using a concatenation of these feature vectors. Finally, AMFB [17] uses attention-based BiLSTM to capture textual features and attention-based CNN-RNN blocks to capture visual features. It then uses a multi-layer perceptron to classify the calculated features.

Our work focuses on fake news detection for social media by proposing a transformer and similarity-based method. Compared with the existing multi-modal fake news detection methods for the general scenario, our approach extracts contextual features from images and considers more interactions between the multi-modal data. Table 1 provides a comparison between the previous state of the art and the proposed method.

Figure 2: The proposed Fake News Revealer (FNR) architecture

Table 1: Related work comparison

Method	Text Encoder	Image Encoder	Fusion Type
EANN [13]	Text-CNN	VGG19	Concat
MVAE [14]	BiLSTM	VGG19	auto-encoder
SpotFake [15]	BERT	VGG19	Concat
CARMN [16]	Word level sentence embeddings	VGG19	Concat + Crossmodal attention
AMFB [17]	Attention based BiLSTM	Attention based CNN+BiRNN	Element wise multiplication
FNR (ours)	BERT	Visual transformer (ViT)	Concat + Similarity

PROPOSED METHOD

The Fake News Revealer (FNR) architecture is shown in Figure 2. The structure of our approach is composed of four parts: 1) text, 2) image, 3) similarity, and 4) classifier. These modules are then combined to provide a label for each news. BERT and VGG19 are used as pre-trained models for extracting text and image feature vectors, respectively. These feature vectors are then combined and classified by a fully connected classification layer.

We also utilize a projector module to resize the input vector for both text and image structures by applying a linear function. The projector’s weights, unlike frozen encoders, are learned end-to-end during training and make the representa-

tion vector of image and text variables. Refer to Figure 3a for more details.

The input to our algorithm is n news (N) each consists of text, image, and label (fake or real). Let N_i represent news i , and T_i , I_i and L_i be its text, image and label, respectively. Then:

$$N_i = (T_i, I_i, L_i) \quad (1)$$

Our goal is to predict L_i using T_i and I_i . We process the news in a batch mode through the deep model, so in each run, the number of the input data is b .

Textual Part

The purpose of this module is to featurize text and embed it into a vector. This part consists of two main sub-modules; the first sub-module is an encoder that extracts representative features extracted from a pre-trained model. The second component is a projector.

BERT [5] has recently been proven to be the most advanced language representation model for a variety of NLP tasks by pretraining on large corpora. During pretraining, BERT can handle short texts like tweets due to their predefined maximum length. Accordingly, the proposed model uses the BERT’s textual feature extraction capabilities.

This part takes T with size (b, t) (t is the maximum text length size). After applying BERT, we take a vector B with size $(b, 768)$, where 768 is the size of the last hidden layer of BERT, and apply the projector to obtain vector F_T with size (b, k) :

$$F_T = w_2 \times (\text{gelu}(w_1 \times B + b_1)) + (w_1 \times B + b_1) + b_2 \quad (2)$$

where w_1 , w_2 , b_1 and b_2 are weights and biases of linear layers inside the text projector.

Visual Part

This module also consists of two sub-modules. The first sub-module is an image encoder that encodes images into a representation vector. We used the ViT Vision Transformer [6] as encoder, and a transformer encoder model (BERT-like) is pre-trained on an extensive collection of images in a supervised fashion. The model is fed with fixed-size patches (resolution 16×16) linearly embedded in images. Using this pre-trained model, we are able to optimize it to our model as the second sub-module. The second sub-module is a projector, as described before.

Subsequently, this module takes I with parameters $(b, \text{width}, \text{height}, \text{depth})$ representing width, height and depth of images, and ViT project it into the vector space V with size $(b, 768)$, where 768 is the size of last hidden vector of ViT model. After applying the projector, we obtain F_i with size (b, k) with the same length as F_t that is obtain from the textual part. The projector works according to the following formula:

$$F_I = w_4 \times (\text{gelu}(w_3 \times V + b_3)) + (w_3 \times V + b_3) + b_4 \quad (3)$$

where w_3 , w_4 , b_3 and b_4 are weights and biases of linear layers inside the image projector.

Similarity Part

In this module, we calculate the similarity of texts and images via a supervised contrastive loss [18]. We have image and text features that are matrices with size (b, k) , and to determine if they are similar, their inner products are often calculated. The similarities between text and image of b tweets are represented in the (b, b) matrix that we call the predicted matrix (P):

$$P = F_T F_I^T \quad (4)$$

This loss function considers an image and a text to be the most similar to itself. Thus, we consider the expected matrix as the average

Figure 3: The sub-modules of projector

similarity of the text to text and the image to image according to the following formula [19]:

$$E = \text{softmax}\left(\frac{F_T F_T^T + F_I F_I^T}{2}\right) \quad (5)$$

After calculating the expected matrix (E), we use cross-entropy to find the actual loss. The contrastive loss is the average of the text similarity loss l_T , and the image similarity loss l_I [18]:

$$l_T = -(E \times \log(P) + (1 - E) \times \log(1 - P)) \quad (6)$$

$$l_I = -(E^T \times \log(P^T) + (1 - E^T) \times \log(1 - P^T)) \quad (7)$$

$$l_s = \frac{l_T + l_I}{2} \quad (8)$$

Classifier Part

In this module, the text and image feature vectors are concatenated to obtain the desired news representation:

$$F_c = \text{concat}(F_T, F_I) \quad (9)$$

This news representation is then passed through two linear layers for fake news classification. After this linear mapping, we consider a vector of size $(b, 2)$ with two classes. According

to Figure 3b and assuming w_6 , w_5 , b_6 and b_5 are weights and biases of linear layers inside the classifier:

$$F = \text{softmax}(w_6 \times (\text{gelu}(w_5 \times F_c + b_5)) + b_6) \quad (10)$$

After passing the vector through the softmax function, we optimize the model by calculating the predicted and actual labels' cross-entropy (α is the balancing factor for the weighted loss function):

$$l_c = -(\alpha L \times \log(F) + (1 - L) \times \log(1 - F)) \quad (11)$$

The loss for the whole network is obtained from the two losses with λ as a tradeoff parameter:

$$l = l_c + \lambda \cdot l_s. \quad (12)$$

EXPERIMENTS

This section presents the implementation issues and experimental results of applying FNR to real datasets and its comparison with the state of the art methods.

Datasets

We used two publicly available datasets gathered from social media as described below.

Twitter This dataset was introduced in [20] for automatic verification of multimedia tasks to distinguish fake/real news on Twitter¹. Data consists of a development set and a test set, each with its own events. The rows of data contain text, images or video, and additional information about the user's profile. The tweets in this dataset are mainly in English (other languages were translated to English before running the experiments).

Weibo This dataset [10] was collected from Weibo² social media from 2012 to 2016 and is written in the Chinese language. Each line of data in this dataset also contains text, user information, and an image. The app's authentication system has tagged texts. This database is divided by [13]

¹<https://twitter.com/>

²<https://weibo.com/>

Table 2: The distribution of datasets

		Twitter	Weibo
Train	Fake	6649	3748
	Real	4599	3758
Test	Fake	545	999
	Real	444	995
All data		12237	9500

into two sets of test data and the train, as the news events of each set are different and. We utilized the same test and train sets.

Only the data with text and images are used in both datasets, and only the first image was utilized. Table 2 lists the number of train/test data and fake/real data that have been used in our experiments.

Implementation Details

The Pytorch framework is used to build our architecture with Python 3.6. We optimized the learning rate using AdamW and calculated different learning rates and weight decays for each part of our architecture to make model converge faster. We have used the Optuna³ library to find the best parameters. The best parameters are as follows: classification learning rate of 0.005, classification weight decay of 0.07, projection vector size (k) of 64, dropout of 0.3, λ was set to 1, and α was the ratio of class data with a large number versus class data with a smaller number.

We set the batch size to 256, our epoch number was set to 100, and the maximum text length (t) for Twitter was 32 words and for Weibo was 200 characters. All input images resized to height=224, width=224 and depth=3. We used a learning rate scheduler and an early stopping checkpoint to avoid overfitting. For the encoder models, the Hugging Face⁴ library was used. The implementation is available in our repository⁵.

Because the input text is raw text gathered from social media, it is non-standard and noisy, and it needs to be cleaned up using normalization techniques. We performed preprocessing which includes converting abbreviations to complete the forms, removing unnecessary punctuations, and deleting non-standard characters. Texts in var-

³<https://optuna.org/>

⁴<https://huggingface.co/>

⁵<http://git.dml.ir/fghorbanpoor/FakeNewsRevealer>

ious languages were also translated to match the dataset’s language. Our text is subsequently tokenized and made ready to be encoded after the preprocessing step. The images were also preprocessed by deleting low-quality images, re-sizing them to (224x224), and converting them into appropriate input for the encoder.

Baselines

The following are some benchmark algorithms and models that we chose for a comparative study.

Single Modality We just considered news text and examined some algorithms to achieve a better feature vector in the textual module. These parameters were tuned by Optuna, and the best results are reported. Our text is subsequently tokenized and ready to be encoded after the preprocessing step.

For the visual module, we used images of news and tuned the parameters with Optuna. The CNN architecture is three parallel convolutional layers with different filter sizes were utilized in this step.

Multi-Modality For multi-modality, we have chosen recent works that are state of the art. These works are listed in Table 1.

We tested two versions of the proposed model in this section; one without considering the similarity measurements (FNR-WS), and the other with considering similarity measurements (FNR-S).

Results

To compare the proposed method in solving multi-modal fake news detection with the previous works, we considered the following criteria for evaluation; accuracy, recall, precision, and F1-score. When dealing with classification problems, these criteria are always taken into consideration.

We are trying to solve a binary classification with balanced data, since the number of fake news and real news are almost equal. We can also consider the AUC criterion and the receiver operating characteristic (ROC) curve for classification evaluation. We included their measurement in the comparison table and plotted the curves for three previous works for comparison purposes.

Figure 4: Ablation study on Twitter dataset

Figure 5: Ablation study on Weibo dataset

Ablation Study An ablation study was performed to explore which data mode is more beneficial or why a multi-modal approach should be used and why we should consider similarity measures. First, we used and tested each mode separately, and then we fused the modes and investigated the effectiveness of a multi-modal approach for fake news detection. Finally, we added the contrastive loss measurement to investigate its effectiveness in enhancing the results.

As demonstrated by Figure 4 on Twitter, because the language of the tweets is brief, imprecise, and filthy, it is less accurate on its own, and the image mode performed better. However, the outcome improves by merging these two modalities, demonstrating that these two modes cover each other’s shortcomings. Adding the similarity part, which considers the relationship between the text and the image, results in more improvements.

As illustrated in Figure 5, the images on Weibo are not very expressive, and the actual news images are almost certainly being exploited

Table 3: Comparison for the Twitter dataset

	Model Name	Accuracy	AUC	F1-score Micro	Fake			Real		
					Precision	Recall	F1-score	Precision	Recall	F1-score
Text	Logistic Regression	0.62	0.46	0.62	0.69	0.55	0.61	0.56	0.70	0.62
	SVM	0.61	0.46	0.61	0.68	0.55	0.61	0.55	0.68	0.62
	BiLSTM	0.61	0.59	0.60	0.62	0.73	0.67	0.58	0.45	0.51
	BERT	0.69	0.69	0.64	0.67	0.68	0.68	0.60	0.59	0.59
Image	CNN	0.62	0.46	0.62	0.69	0.55	0.61	0.56	0.55	0.61
	VGG19	0.68	0.46	0.68	0.74	0.64	0.69	0.62	0.72	0.67
	ViT	0.72	0.71	0.76	0.74	0.86	0.80	0.79	0.63	0.70
multi-modal	EANN	0.69	0.72	0.69	0.75	0.58	0.65	0.62	0.76	0.69
	MVAE	0.67	0.66	0.67	0.70	0.69	0.69	0.63	0.64	0.63
	SpotFake	0.77	0.74	0.76	0.72	0.92	0.81	0.85	0.56	0.68
	CARMN	0.73	0.69	0.73	0.70	0.88	0.78	0.78	0.54	0.64
	AMFD	0.75	0.74	0.75	0.76	0.79	0.78	0.73	0.70	0.71
	FNR-WS	0.77	0.76	0.77	0.74	0.90	0.81	0.83	0.62	0.71
	FNR-S	0.79	0.79	0.79	0.78	0.85	0.82	0.79	0.71	0.75

Figure 6: ROC curve on Twitter dataset

Figure 7: ROC curve on Weibo dataset

for false news, which does not help to detect fake news on its own. Nonetheless, the text mode outperforms the visual mode. However, when these two modalities are included concurrently, the model’s performance significantly improves, as these two type of modalities complement each other. When the relationship between text and image is considered, the accuracy also increases.

Performance Comparison Using the mentioned criteria, the following comparison table for the Twitter dataset is presented in Table 3 and for the Weibo dataset in Table 4 and Figure 6 compares the ROC curves of three relevant studies and our work on the Twitter dataset, and Figure 7 compares them on the Weibo dataset.

Based on the results in these tables, BERT has

performed better than the other text algorithms. It has properly extracted conceptual and lexical features, and has a more accurate approach in detecting fake news with text only.

Moreover, ViT outperforms other fake news detection algorithms that rely solely on images, and with a transformer-based network, it extracts valuable and efficient features.

According to the results, some current methods in multi-modal fake news detection perform more favorably compare to single mode methods, indicating that the combination of data modes and the extraction of their intermodal properties are beneficial. Spotfake outperformed the competition because it pulls more useful characteristics from text and images. In addition, both CARMN and AMFD have presented a more performant

Table 4: Comparison for the Weibo dataset

					Fake			Real		
	Model Name	Accuracy	AUC	F1-score Micro	Precision	Recall	F1-score	Precision	Recall	F1-score
Text	Logistic Regression	0.71	0.67	0.71	0.71	0.80	0.75	0.71	0.59	0.65
	SVM	0.70	0.70	0.70	0.72	0.74	0.73	0.67	0.65	0.66
	BiLSTM	0.66	0.44	0.66	0.62	0.78	0.69	0.73	0.55	0.63
	BERT	0.81	0.83	0.81	0.81	0.81	0.81	0.81	0.82	0.81
Image	CNN	0.52	0.39	0.50	0.79	0.24	0.38	0.58	0.87	0.63
	VGG19	0.60	0.47	0.60	0.60	0.61	0.60	0.60	0.59	0.59
	ViT	0.68	0.72	0.68	0.67	0.69	0.68	0.68	0.68	0.67
multi-modal	EANN	0.81	0.86	0.81	0.89	0.66	0.76	0.77	0.93	0.85
	MVAE	0.79	0.79	0.79	0.89	0.65	0.75	0.74	0.93	0.82
	SpotFake	0.86	0.90	0.86	0.87	0.92	0.90	0.81	0.70	0.75
	CARMN	0.84	0.90	0.85	0.86	0.93	0.89	0.81	0.66	0.73
	AMFD	0.83	0.89	0.83	0.86	0.90	0.88	0.75	0.68	0.71
	FNR-WS	0.87	0.90	0.87	0.89	0.85	0.87	0.85	0.89	0.87
	FNR-S	0.88	0.94	0.88	0.89	0.87	0.88	0.87	0.89	0.88

combination of attention-based mechanisms.

The Fake News Revealer (FNR) framework consistently outperforms the competition on a variety of performance metrics. Each modality can retain its distinct traits with our methodology while seamlessly combining similarity and complementary information from the other modalities.

CONCLUSION AND FUTURE WORKS

A new multi-modal framework for detecting fake news employing text and image transformers, as well as similarity measures with a contrastive loss function, is presented in this work. To extract the textual information, we employed BERT, which is a transformer-based encoder. The visual characteristics were retrieved using ViT, a BERT-like transformer that was explicitly created for image processing. The connection between an image and a sentence is extracted using a contrastive loss and then fed into a two-layer linear classifier to assess if the news is real or fake. We tested the proposed method on two freely available multi-modal datasets. Extensive tests have been carried out to compare to existing models. The suggested framework (FNR) outperforms the state of the art methods in detecting fake news. Since a large amount of news is being published with different modes, one can use other data modes such as video, audio, and news graphs.

REFERENCES

1. D. Gross, "Survey: More americans get news from internet than newspapers or radio," accessed:2020-01-16. [Online]. Available: <http://www.cnn.com/2010/TECH/03/01/social.network.news/index.html>
2. K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *SIGKDD Explor. Newsl.*, vol. 19, no. 1, p. 22–36, Sep 2017.
3. E. J. Newman, M. Garry, D. M. Bernstein, J. Kantner, and D. S. Lindsay, "Nonprobative photographs (or words) inflate truthiness," *Psychonomic Bulletin & Review*, vol. 19, no. 5, pp. 969–974, Oct 2012.
4. E. Fenn, N. Ramsay, J. Kantner, K. Pezdek, and E. Abed, "Nonprobative photos increase truth, like, and share judgments in a simulated social media environment," *JARMAC*, vol. 8, no. 2, pp. 131–138, 2019.
5. J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proc. North Amer. Ch. Assoc. Comput. Linguistics: Human Lang. Technol.*, 2019, pp. 4171–4186.
6. A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021. [Online]. Available: <https://arxiv.org/abs/2010.11929>
7. S. Kwon, M. Cha, K. Jung, W. Chen, and Y. Wang, "Prominent features of rumor propagation in online so-

- cial media,” in *Proc. 13th IEEE Conf. Data Min.*, 2013, pp. 1103–1108.
8. M. Gupta, P. Zhao, and J. Han, “Evaluating event credibility on twitter,” in *Proc. SIAM Conf. Data Min.*, 2012, pp. 153–164.
9. D. Tian, “A review on image feature extraction and representation techniques,” *Int. J.*, vol. 8, pp. 385–395, Jan 2013.
10. Z. Jin, J. Cao, H. Guo, Y. Zhang, and J. Luo, “Multimodal fusion with recurrent neural networks for rumor detection on microblogs,” in *Proc. 25th ACM Conf. Multimedia*, 2017, pp. 795–816.
11. S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, “Vqa: Visual question answering,” in *Proc. IEEE Conf. Comput. Vis.*, 2015, pp. 2425–2433.
12. M. Farajtabar, J. Yang, X. Ye, H. Xu, R. Trivedi, E. Khalil, S. Li, L. Song, and H. Zha, “Fake news mitigation via point process based intervention,” in *Proc. 34th Conf. Mach. Learn. Res.*, 2017, p. 1097–1106.
13. 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 *Proc. 24th ACM SIGKDD Conf. Data Min. Knowl. Discov.*, 2018, p. 849–857.
14. D. Khattar, J. S. Goud, M. Gupta, and V. Varma, “Mvae: Multimodal variational autoencoder for fake news detection,” in *Proc. 19th WWW Conf.*, 2019, p. 2915–2921.
15. S. Singhal, R. R. Shah, T. Chakraborty, P. Kumaraguru, and S. Satoh, “Spotfake: A multi-modal framework for fake news detection,” in *Proc. IEEE 5th Conf. Multimedia BigMM*, 2019, pp. 39–47.
16. C. Song, N. Ning, Y. Zhang, and B. Wu, “A multimodal fake news detection model based on crossmodal attention residual and multichannel convolutional neural networks,” *Inf. Process. Manag.*, vol. 58, no. 1, p. 102437, 2021.
17. R. Kumari and A. Ekbal, “Amfb: Attention based multimodal factorized bilinear pooling for multimodal fake news detection,” *Expert Syst. Appl.*, vol. 184, p. 115412, 2021.
18. A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” 2021. [Online]. Available: <https://arxiv.org/abs/2103.00020>
19. K. Salama, “Keras documentation: Natural language image search with a dual encoder,” accessed:2021-11-08. [Online]. Available: https://keras.io/examples/nlp/nl_image_search/
20. C. Boididou, K. Andreadou, S. Papadopoulos, D. T.

Dang Nguyen, G. Boato, M. Riegler, M. Larson, and I. Kompatsiaris, “Verifying multimedia use at mediaeval 2015,” in *MediaEval*, Sep 2015.

Faeze Ghorbanpour is currently a research fellow at the Digital Media Laboratory, Sharif University of Technology. Her research interests include computational social science and social networks. She received her B.Sc. and M.Sc. degrees in computer engineering from the Sharif University of Technology in Tehran, Iran, in 2018 and 2021, respectively.

Maryam Ramezani received her B.Sc. and M.Sc. degrees in Information Technology Engineering from the Sharif University of Technology, Tehran, Iran. She is currently a PhD candidate in the Department of Computer Engineering at the Sharif University of Technology. Her current research interests include machine learning with application in social and complex networks.

MohammadAmin Fazli received his B.Sc. degree in hardware engineering and the M.Sc. and PhD degrees in software engineering from the Sharif University of Technology in 2009, 2011, and 2015. He is currently an Assistant Professor at the Sharif University of Technology and the R&D supervisor of the Intelligent Information Center. His research interests include game theory, combinatorial optimization, computational economics, graphs and combinatorics, complex networks, and dynamical systems.

Hamid R. Rabiee (SM '07) received his BS and MS degrees (with Great Distinction) in Electrical Engineering from CSULB, Long Beach, CA (1987, 1989), his EEE degree in Electrical and Computer Engineering from USC, Los Angeles, CA (1993), and his Ph.D. in Electrical and Computer Engineering from Purdue University, West Lafayette, IN, in 1996. From 1993 to 1996 he was a Member of Technical Staff at AT&T Bell Laboratories. From 1996 to 1999 he worked as a Senior Software Engineer at Intel Corporation. He was also with PSU, OGI and OSU universities as an adjunct professor of Electrical and Computer Engineering from 1996-2000. Since September 2000, he has joined Sharif University of Technology, Tehran, Iran. He was also a visiting professor at the Imperial College of London for the 2017-2018 academic year. He is the founder of Sharif University Advanced Information and Communication Technology Research Institute (AICT), ICT Innovation Center, Advanced Technologies Incubator (SATI), Digital Media Laboratory (DML), Mobile Value Added Services Laboratory (VASL), Bioinformatics and Computational Biology

Laboratory (BCB) and Cognitive Neuroengineering Research Center. He is also a consultant and member of AI in Health Expert Group at WHO. He has been the founder of many successful High-Tech start-up companies in the field of ICT as an entrepreneur. He is currently a Professor of Computer Engineering at Sharif University of Technology, and Director of AICT, DML, and VASL. He has received numerous awards and honors for his Industrial, scientific and academic contributions, and holds three patents. His research interests include statistical machine learning, Bayesian statistics, data analytics and complex networks with applications in social networks, multimedia systems, cloud and IoT privacy, bioinformatics, and brain networks.