

Team Yao at Factify 2022: Utilizing Pre-trained Models and Co-attention Networks for Multi-Modal Fact Verification

Wei-Yao Wang¹, Wen-Chih Peng¹

¹Department of Computer Science, National Yang Ming Chiao Tung University, Hsinchu, Taiwan

Abstract

In recent years, social media has enabled users to get exposed to a myriad of misinformation and disinformation; thus, misinformation has attracted a great deal of attention in research fields and as a social issue. To address the problem, we propose a framework, Pre-CoFact, composed of two pre-trained models for extracting features from text and images, and multiple co-attention networks for fusing the same modality but different sources and different modalities. Besides, we adopt the ensemble method by using different pre-trained models in Pre-CoFact to achieve better performance. We further illustrate the effectiveness from the ablation study and examine different pre-trained models for comparison. Our team, Yao, won the fifth prize (F1-score: 74.585%) in the Factify challenge hosted by De-Factify @ AACL 2022, which demonstrates that our model achieved competitive performance without using auxiliary tasks or extra information. The source code of our work is publicly available¹.

Keywords

Multi-modal fact verification, Transformer, Co-attention, De-Factify

1. Introduction

Fake news has become easier to spread due to the growing number of users of social media. For example, about 59% of social media consumers expect that news spread via social media may be inaccurate [1]. To influence social thoughts, there are many fake news stories that mislead readers about the news content by replacing some true content with false details. Besides, fake news with textual and visual content can better attract readers and it is hard to judge than only using textual content. Therefore, it is essential to detect multi-modal fake news to eliminate its negative impacts.

Fake checkers aim to check the worthiness, evidence or verified claim retrieval [2]. Recent works have presented a number of approaches for tackling fake news detection automatically. In uni-modal detection, Shu et al. [3] exploited a tri-relationship (publishers, news pieces, and users) to model the relations and interactions for detecting news disinformation. Przybyła [4] utilized the style the news articles are written in to estimate their credibility. In multi-modal detection, Jin et al. [5] proposed an att-RNN that combines a recurrent neural network with

¹<https://github.com/wywyWang/Multi-Modal-Fact-Verification-2021>

De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection, co-located with AACL 2022. 2022 Vancouver, Canada

✉ sf1638.cs05@nctu.edu.tw (W. Wang); wcpeng@nctu.edu.tw (W. Peng)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

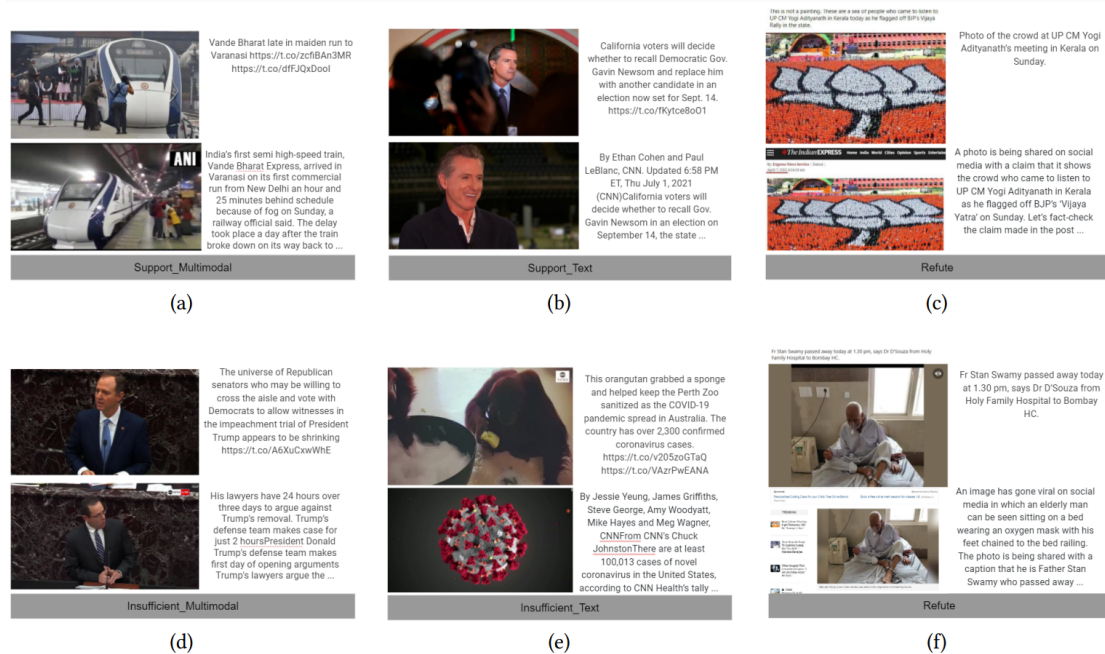


Figure 1: A screenshot from [8], which illustrates sample examples of five categories from the Factify dataset.

an attention mechanism to fuse textual content and visual images. MCAN is proposed by extracting spatial-domain features and textual features by pre-trained models [6]. Further, to address the fact that fake images are often re-compressed images or tampered images, which shows periodicity in the frequency domain, they used discrete cosine transform as in [7], then designed a CNN-based network for capturing frequency-domain features from images.

A real-world problem, identifying if the claim entails the document, is the challenge called Factify [8, 9] hosted by De-Factify¹. Figure 1 shows some examples for all five categories: Support_Multimodal, Support_Text, Insufficient_Multimodal, Insufficient_Text, and Refute. To tackle the problem, in this paper, we propose Pre-CoFact with pre-trained models and co-attention networks to perform the shared task, which first extracts features from both text and images, then fuses this information through the co-attention module. Specifically, two powerful Transformer-based pre-trained models, DeBERTa [10] and DeiT [11], are adopted for extracting features from both claims and documents' text and images, respectively. Afterwards, several co-attention modules are designed for fusing the contexts of the text and images. Finally, these embeddings are aggregated as corresponding embeddings to classify the category of the news.

The main results of this paper can be summarized as follows:

- Using text and images directly can achieve expressive results without any auxiliary tasks, preprocessing methods, or extra information (e.g., optical character recognition (OCR))

¹<https://aiisc.ai/defactify/factify.html>

- from images).
- Adopting pre-trained models helps improve the performance of the shared task, and using co-attention networks can learn the correlation from the same modality (text or images from claims and documents) and the dependencies between different modalities (text and images).
 - Our ensemble model outperforms the machine learning models [8] at least 48% and 40% in terms of validation score and testing score. Besides, extensive experiments were further conducted to examine the capability of the proposed model.

2. Dataset

Factify is a dataset for multi-modal fact verification, which contains images, textual claims, reference textual documents and images. Each sample includes `claim_image`, `claim`, `claim_ocr`, `document_image`, `document`, `document_ocr`, and `category`. The detail of each field is described as follows:

- `claim_image`: the image of the given claim.
- `claim`: the text of the given claim.
- `claim_ocr`: the text from the `claim_image` detected by the host.
- `document_image`: the image of the given reference.
- `document`: the text of the given reference.
- `document_ocr`: the text from the `document_image` detected by the host.
- `category`: the category of the data sample from a list of five classes.

The category is composed of 1) `Support_Multimodal`: both the claim text and image are similar to that of the document, 2) `Support_Text`: the claim text is similar or entailed, but images of the document and claim are not similar, 3) `Insufficient_Multimodal`: the claim text is neither supported nor refuted by the document but images are similar to the document, 4) `Insufficient_Text`: both text and images of the claim are neither supported nor refuted by the document, although it is possible that the text claim has common words with the document text, and 5) `Refute`: the images and/or text from the claim and document are completely contradictory.

The training set contains 35,000 samples, which has 5,000 samples of each class, and the validation set contains 7,500 samples, which has 1,500 samples of each class. The test set, which is used to evaluate the private score, also contains 7,500 samples. For more details, we refer readers to [8].

3. Related Works

3.1. Fake News Detection

There have been a series of studies combating fake news detection to mitigate a societal crisis [12]. Vo and Lee [13] proposed a novel neural ranking model which jointly utilizes textual and visual matching signals. This is the first work using multi-modal data in social media posts to search for verified information, which can increase users' awareness of fact-checked information

when they are exposed to fake news. Lee et al. [14] adopted a perplexity-based approach in the few-shot learning, which assumes that the given claim may be fake if the corresponding perplexity score from evidence-conditioned language models is high. BertGCN [15] is proposed by integrating the advantages of large-scale pre-trained models and graph neural networks for fake news detection, which is able to learn the representations from the massive amount of pre-trained data and the label influence through the propagation. MCAN [6] adopts a large-scale pre-trained NLP model and a pre-trained computer vision (CV) model for extracting features from text and images, respectively. Besides, MCAN also extracts frequency domain features from images, and then uses multiple co-attention layers to fuse this information.

These approaches demonstrate the effectiveness of using pre-trained models for fake news detection, which motivated us to use pre-trained models as well. Besides, MCAN inspires us to fuse the contexts of different modalities or the same modality (e.g., text from claims and documents).

3.2. Large-Scale Pre-trained Models

Transformer [16] has been used for machine translation and has inspired many competitive approaches in natural language processing (NLP) tasks. Transformer-based pre-trained language models (PLMs) have significantly improved the performance of various NLP tasks due to the ability to understand contextualized information from the pre-trained dataset. Since BERT [17] was presented, we have seen the rise of a set of large-scale PLMs such as GPT-3 [18], RoBERTa [19], XLNet [20], ELECTRA [21], and DeBERTa [10]. These PLMs have been fine-tuned using task-specific labels and have created a new state of the art in many downstream tasks.

Recently, vision Transformer (ViT) [22] is a Transformer encoder architecture directly applied to image classification with patching raw images as input to NLP, which achieves competitive results compared to state-of-the-art convolutional networks by pre-training a large private image dataset JFT-300M [23]. ViT demonstrates that convolution-free networks can still learn the relation in the images. To reduce the pre-trained dataset size and training efficiency, several follow-up studies have been conducted. DINO was proposed by [24] to improve the standard ViT model through self-supervised learning. [11] proposed DeiT, which used a novel distillation procedure based on a distillation token to ensure the student learns from the teacher through attention.

These pre-trained models demonstrate the generalization of various domains. Further, using pre-trained models benefits capturing rich information of downstream tasks, which can also reduce the burden of training from scratch. These advantages motivated us to adopt state-of-the-art pre-trained models for transforming images and text into contextual embeddings. Besides, we focused on using Transformer-based pre-trained models for feature extraction.

4. Method

4.1. Problem Formulation

Let $C = \{CT_i, CI_i, DT_i, DI_i\}_{i=1}^{|C|}$ denote the corpus of the dataset, where the i -th sample is composed of the claim text $CT_i = w_1^{CT_i} w_2^{CT_i} \dots$, the claim image CI_i , the document text $DT_i = w_1^{DT_i} w_2^{DT_i} \dots$,

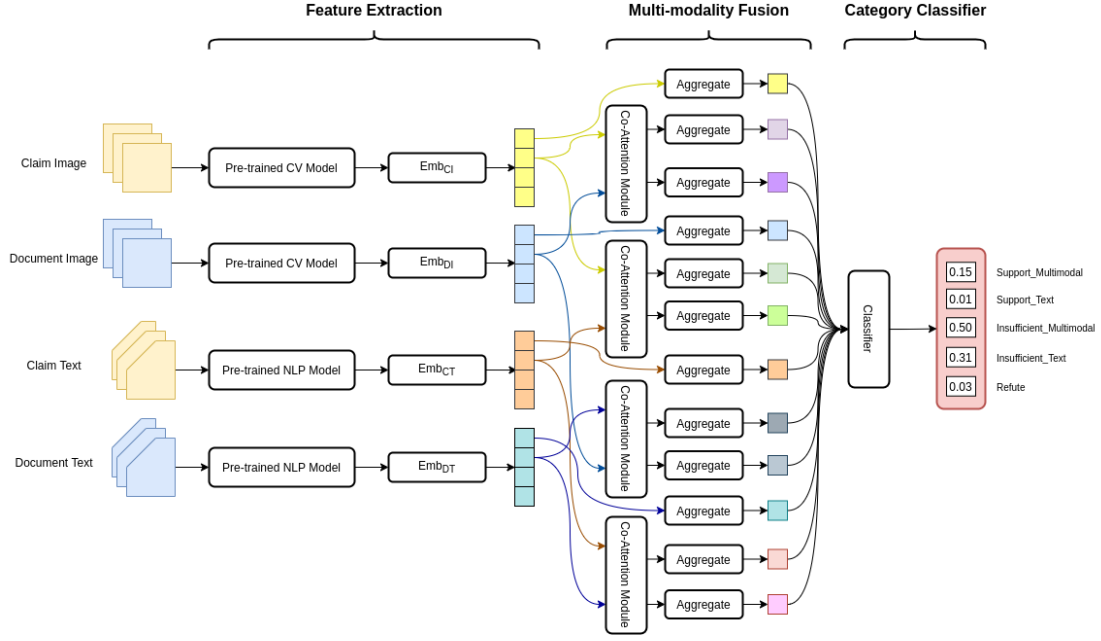


Figure 2: Illustration of the Pre-CoFact framework. Each square can be seen as a token with a d dimension vector. The feature extraction part aims to transform text and images into corresponding embeddings. The multi-modality fusion part fuses this information from the same modality (images/text from the claim and document) and different modalities (images and text from the claim/document) to obtain contexts. Finally, the category classifier predicts the possible classes based on the embeddings from feature extraction and the embeddings from multi-modality fusion.

and the document image DI_i . The i -th target $y_i \in \{Support_Multimodal, Support_Text, Insufficient_Multimodal, Insufficient_Text, Refute\}$. The goal is to find out support, insufficient-evidence and refute between given claims and documents.

4.2. Pre-CoFact Overview

Figure 2 illustrates the overview of the proposed Pre-CoFact framework. The input contains the claim image, the claim text, the document image, and the document text. The feature extraction part adopts DeiT [11] as the pre-trained CV model and DeBERTa [10] as the pre-trained NLP model, and feeds the outputs of pre-trained models to the image embedding layer and text embedding layer for transforming images and texts into corresponding embeddings. The multi-modality fusion part fuses this information from the same modality (images/text from the claim and document) and different modalities (images and text from the claim/document) based on multiple co-attention layers. Finally, the category classifier predicts the possible classes based on the embeddings from feature extraction and the embeddings from multi-modality fusion.

4.3. Feature Extraction

The enrichment of pre-trained models enables us to have rich information without training from scratch. Moreover, Transformer-based pre-trained models demonstrate the success on both NLP and CV tasks. However, it is essential to fine-tune for fitting in our task. To this end, we first use DeBERTa as our pre-trained NLP model and DeiT as our pre-trained CV model, and then we use the embedding layer for transforming pre-trained embeddings to embeddings in our task. Specifically, the i -th output of the embedding layer is calculated as follows:

$$E_{CI_i} = Emb_{CI}(DeiT(CI_i)), E_{DI_i} = Emb_{DI}(DeiT(DI_i)), \quad (1)$$

$$E_{CT_i} = Emb_{CT}(DeBERTa(CT_i)), E_{DT_i} = Emb_{DT}(DeBERTa(DT_i)), \quad (2)$$

where the output dimensions of DeiT and DeBERTa are 768, the Emb is composed of a MLP and an activation function, and $E_{CI_i}, E_{CT_i}, E_{DI_i}, E_{DT_i}$ are d dimension vectors. It is noted that the activation functions in Emb we used are ReLU and Mish [25] for testing the results.

4.4. Multi-Modality Fusion

After generating embeddings of text and images, we adopt multiple co-attention layers as in [6] to fuse the embeddings. To check the relation between claim and document, we use the co-attention layer to separately fuse 1) images of claims and documents and 2) text of claims and documents. Besides, the relation between text and images from the claims or document can be viewed as checking whether they are relative or not. Therefore, we also adopt the co-attention layer for fusing 3) images and text of claims and 4) images and text of documents.

Specifically, each co-attention layer takes two inputs E_A and E_B and produces two outputs H_A, H_B . Here we use a single head to derive as the following equations:

$$Q_A = E_A W^{Q_A}, K_A = E_A W^{K_A}, V_A = E_A W^{V_A}, Q_B = E_B W^{Q_B}, K_B = E_B W^{K_B}, V_B = E_B W^{V_B}, \quad (3)$$

$$\tilde{H}_A = Norm(E_A + softmax(\frac{Q_A K_B^T}{\sqrt{d}}) V_B), \tilde{H}_B = Norm(E_B + softmax(\frac{Q_B K_A^T}{\sqrt{d}}) V_A), \quad (4)$$

$$H_A = Norm(\tilde{H}_A + FFN(\tilde{H}_A)), H_B = Norm(\tilde{H}_B + FFN(\tilde{H}_B)), \quad (5)$$

where $W^{Q_A}, W^{K_A}, W^{V_A}, W^{Q_B}, W^{K_B}, W^{V_B} \in \mathbb{R}^{d \times d}$, and $Norm$ and FFN is the same normalization method and feed forward network as in [16]. Co-attention block has been widely used in VQA tasks [26], as it can capture dependencies of different inputs. Therefore, we use the co-attention layer for fusing:

$$H_{CIDI_i}, H_{DICI_i} = CoAtt(E_{CI_i}, E_{DI_i}), H_{CTDT_i}, H_{DTCT_i} = CoAtt(E_{CT_i}, E_{DT_i}), \quad (6)$$

$$H_{CIDT_i}, H_{DTCI_i} = CoAtt(E_{CI_i}, E_{DT_i}), H_{CTDI_i}, H_{DICT_i} = CoAtt(E_{CT_i}, E_{DI_i}), \quad (7)$$

where $CoAtt$ denotes the co-attention layer.

After applying the co-attention mechanism, the aggregation function is adopted to aggregate fused tokens into a representative token. That is, given a fused embedding with $\mathbb{R}^{N \times d}$, where N is the sequence length, we use mean aggregation to output $\mathbb{R}^{1 \times d}$. Besides, we also feed $E_{CI_i}, E_{CT_i}, E_{DI_i}, E_{DT_i}$ into the aggregation function for classification.

4.5. Category Classifier

To predict the label of the given claims and documents, we first concatenate 8 aggregated outputs $H_{CIDI_i}, H_{DICI_i}, H_{CTDI_i}, H_{DTCT_i}, H_{CIDT_i}, H_{DTCI_i}, H_{CTDI_i}, H_{DICT_i}$ from the co-attention layers and original 4 aggregated embeddings $E_{CI_i}, E_{CT_i}, E_{DI_i}, E_{DT_i}$ to obtain the input of the classifier Z_i . It is worth noting that the outputs of embeddings are also used since the original information can provide some clues for classifying the news. Afterwards, the i -th output of the classifier is the probability as follows:

$$Z_i^{M_1} = \sigma(Z_i W^Z), Z_i^{M_2} = \sigma(Z_i^{M_1} W^{M_1}), \quad (8)$$

$$\hat{y}_i = \text{softmax}(Z_i^{M_2} W^{M_2}), \quad (9)$$

where $W^Z \in \mathbb{R}^{12d \times d}$, $W^{M_1} \in \mathbb{R}^{d \times d_{M_1}}$, and $W^{M_2} \in \mathbb{R}^{d_{M_1} \times 5}$. Note that σ is the same as in *Emb*, which uses both ReLU and Mish for testing the results.

We trained our model by minimizing cross-entropy loss \mathbb{L} to learn the prediction of the categories:

$$\mathbb{L} = - \sum_{i=1}^{|C|} y_i \log(\hat{y}_i). \quad (10)$$

4.6. Ensemble Method

Each classifier may have its strengths and weakness, and ensemble methods have been widely used to enhance the performance. Therefore, we follow [27] to use the power weighted sum to enhance the performance of the model. The formula is derived as follows:

$$p = p_1^N \times w_1 + p_2^N \times w_2 + \dots + p_k^N \times w_k, \quad (11)$$

where p_1, \dots, p_k are the predicted probability from the corresponding model, w_1, \dots, w_k are weights with respect to the corresponding model, k is the number of trained models, and N is the weight of power. It is noted that these parameters are tuned by hand.

5. Results and Analysis

5.1. Experimental Setup

5.1.1. Implementation Details

The dimension d was set to 512, the inner dimension of the feed-forward layer was 1024, and the number of heads was set to 4. The dropout rate was 0.1, and the max sequence length was 512. The batch size was 32, the learning rates were set to $3e-5$ and $2e-5$, the training epochs were set to 30, and the seeds were tested with 41 and 42. The power N was set to 0.5, and the weights were set to 0.6, 0.2, 0.1, 0.2, 0.3, which were manually tuned by validation score. The pre-trained DeBERTa was deberta-base², and the DeiT was deit-base-patch16-224³. The parameters of the two pre-trained models were frozen. All images were transformed by resizing to 256, center

²<https://huggingface.co/microsoft/deberta-base>

³<https://huggingface.co/facebook/deit-base-patch16-224>

Model	w/o CoAtt	w/o CoAtt(text, image)	Pre-CoFact (Ours)	Ensemble (Ours)
Weighted F1 (%)	75.68 (-4.34)	76.43 (-3.59)	78.46	80.02 (+1.56)

Table 1

Ablation study of our model in terms of validation score. w/o CoAtt denotes using four embeddings for classification and w/o CoAtt(text, image) denotes using only the same modality (Equ. 6).

Model	DINO [24]	XLM-RoBERTa [28]	RoBERTa [19]	Pre-CoFact (Ours)
Weighted F1 (%)	73.94 (-4.52)	74.11 (-4.35)	77.53 (-0.93)	78.46

Table 2

Variant pre-trained models in terms of validation score. Pre-CoFact uses DeiT and DeBERTa as pre-trained models. DINO is replaced DeiT by DINO [24]. XLM-RoBERTa and RoBERTa are replaced DeBERTa by XLM-RoBERTa [28] and RoBERTa [19], respectively.

cropping to 224, and normalizing. We preprocessed only for transforming images, and then we stored the text and processed images in corresponding pickle files for training and evaluating. All the training and evaluation phases were conducted on a machine with Intel Xeon 4110 CPU @ 2.10GHz, Nvidia GeForce RTX 2080 Ti, and 252GB RAM. The source code is available at <https://github.com/wywyWang/Multi-Modal-Fact-Verification-2021>.

5.1.2. Evaluation Metric

To evaluate the performance of the task, the weighted average F1 score was used across the 5 categories.

5.2. Quantitative Results

5.2.1. Ablation Study

We first conducted an ablation study to ensure the effective design of our proposed Pre-CoFact. As shown in Table 1, it is evident that without co-attention networks (w/o CoAtt), the performance is degraded. Further, applying co-attention only on the same modality (w/o CoAtt(text, image)) is insufficient, which demonstrates the need for modeling dependencies between different modalities. It is noted that our ensemble method slightly improves the performance compared to Pre-CoFact. Our ensemble method includes Pre-CoFact, Pre-CoFact with replacing DeBERTa with XLM-RoBERTa, Pre-CoFact with replacing DeBERTa with RoBERTa, Pre-CoFact with replacing DeBERTa with RoBERTa and replacing ReLU with Mish, and Pre-CoFact with replacing ReLU with Mish.

We also use different pre-trained models to examine the module influence as shown in Table 2. It can be seen that DeiT is more suitable than DINO for this task. Besides, XLM-RoBERTa also degrades the performance, while RoBERTa is slightly worse than Pre-CoFact with DeBERTa.

Rank	Team	Support _Text (%)	Support _Multimodal (%)	Insufficient _Text (%)	Insufficient _Multimodal (%)	Refute (%)	Final (%)
5	Yao	68.881	81.610	84.836	88.309	100.00	74.585
-	Baseline	82.675	75.466	74.424	69.678	42.354	53.098

Table 3

Performance of our model in terms of testing score. Our method achieved fifth prize with only about a 2.2% gap, while we outperformed the baseline by 40.5%.

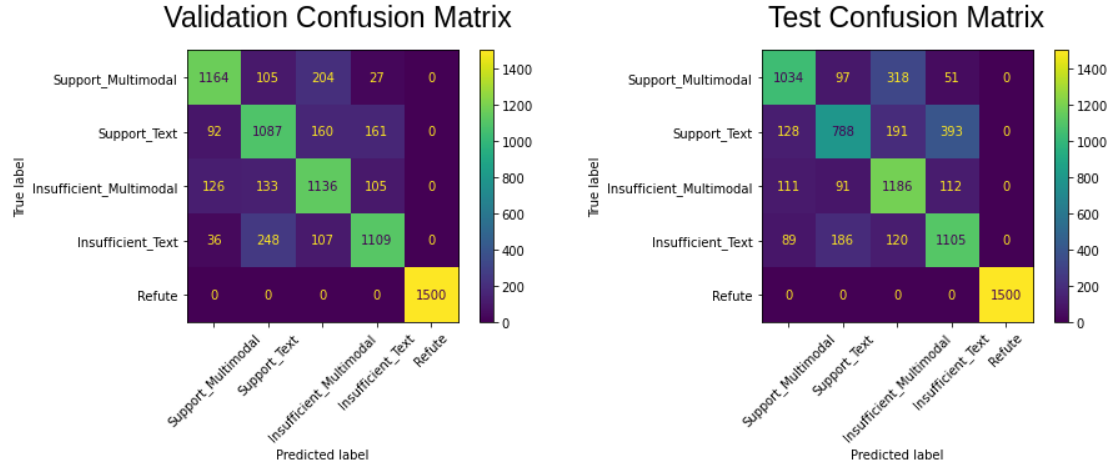


Figure 3: Confusion matrix of the validation set and testing set.

5.2.2. Testing Performance

Table 3 shows the performance of the testing set. Our approach achieved 74.585% of the F1-score, winning the fifth prize in detecting fake news. This result outperformed the baseline by 40.5%, while it still has only a 2.2% gap compared to the first prize. Despite the result, our approach still demonstrates that using only text and images can achieve competitive performance.

5.2.3. Confusion Matrix

Figure 3 shows the confusion matrix of the validation set and testing set. It can be observed that our model can precisely classify refute on both the validation set and testing set, while our model misjudged whether the text is entailed when the image is not entailed.

6. Conclusion

In this paper, we proposed Pre-CoFact utilizing pre-trained models and multiple co-attention networks to alleviate the effect of fake news for the Factify task. To achieve better performance, we adopted an ensemble method by weighting several models. The ablation study demonstrates

the effectiveness of our proposed approach. From the testing score, our method illustrates that using only text and images without extra information can also achieve competitive performance.

References

- [1] E. Shearer, A. Mitchell, News use across social media platforms in 2020, 2021. URL: <https://www.pewresearch.org/journalism/2021/01/12/news-use-across-social-media-platforms-in-2020/>.
- [2] P. Nakov, D. P. A. Corney, M. Hasanain, F. Alam, T. Elsayed, A. Barrón-Cedeño, P. Papotti, S. Shaar, G. D. S. Martino, Automated fact-checking for assisting human fact-checkers, in: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, 2021, pp. 4551–4558.
- [3] K. Shu, S. Wang, H. Liu, Beyond news contents: The role of social context for fake news detection, in: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 2019, pp. 312–320.
- [4] P. Przybyla, Capturing the style of fake news, in: The Thirty-Fourth AAAI Conference on Artificial Intelligence, 2020, pp. 490–497.
- [5] Z. Jin, J. Cao, H. Guo, Y. Zhang, J. Luo, Multimodal fusion with recurrent neural networks for rumor detection on microblogs, in: Proceedings of the 2017 ACM on Multimedia Conference, 2017, pp. 795–816.
- [6] Y. Wu, P. Zhan, Y. Zhang, L. Wang, Z. Xu, Multimodal fusion with co-attention networks for fake news detection, in: Findings of the Association for Computational Linguistics, volume ACL/IJCNLP 2021 of *Findings of ACL*, 2021, pp. 2560–2569.
- [7] P. Qi, J. Cao, T. Yang, J. Guo, J. Li, Exploiting multi-domain visual information for fake news detection, in: 2019 IEEE International Conference on Data Mining, 2019, pp. 518–527.
- [8] S. Mishra, S. Suryavardan, A. Bhaskar, P. Chopra, A. Reganti, P. Patwa, A. Das, T. Chakraborty, A. Sheth, A. Ekbal, C. Ahuja, Factify: A multi-modal fact verification dataset, in: Proceedings of De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection, CEUR, 2022.
- [9] P. Patwa, S. Mishra, S. Suryavardan, A. Bhaskar, P. Chopra, A. Reganti, A. Das, T. Chakraborty, A. Sheth, A. Ekbal, C. Ahuja, Benchmarking multi-modal entailment for fact verification, in: Proceedings of De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection, CEUR, 2022.
- [10] P. He, X. Liu, J. Gao, W. Chen, Deberta: decoding-enhanced bert with disentangled attention, in: 9th International Conference on Learning Representations, 2021.
- [11] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, H. Jégou, Training data-efficient image transformers & distillation through attention, in: Proceedings of the 38th International Conference on Machine Learning, volume 139 of *Proceedings of Machine Learning Research*, 2021, pp. 10347–10357.
- [12] P. Nakov, G. D. S. Martino, Fake news, disinformation, propaganda, media bias, and flattening the curve of the COVID-19 infodemic, in: KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2021, pp. 4054–4055.
- [13] N. Vo, K. Lee, Where are the facts? searching for fact-checked information to alleviate

- the spread of fake news, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, 2020, pp. 7717–7731.
- [14] N. Lee, Y. Bang, A. Madotto, P. Fung, Towards few-shot fact-checking via perplexity, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2021, pp. 1971–1981.
 - [15] Y. Lin, Y. Meng, X. Sun, Q. Han, K. Kuang, J. Li, F. Wu, Bertgcn: Transductive text classification by combining GCN and BERT, CoRR abs/2105.05727 (2021).
 - [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 2017, pp. 5998–6008.
 - [17] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, 2019, pp. 4171–4186.
 - [18] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, 2020.
 - [19] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized BERT pretraining approach, CoRR abs/1907.11692 (2019).
 - [20] Z. Yang, Z. Dai, Y. Yang, J. G. Carbonell, R. Salakhutdinov, Q. V. Le, Xlnet: Generalized autoregressive pretraining for language understanding, in: Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, 2019, pp. 5754–5764.
 - [21] K. Clark, M. Luong, Q. V. Le, C. D. Manning, ELECTRA: pre-training text encoders as discriminators rather than generators, in: 8th International Conference on Learning Representations, 2020.
 - [22] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale, in: 9th International Conference on Learning Representations, 2021.
 - [23] C. Sun, A. Shrivastava, S. Singh, A. Gupta, Revisiting unreasonable effectiveness of data in deep learning era, in: IEEE International Conference on Computer Vision, 2017, pp. 843–852.
 - [24] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, A. Joulin, Emerging properties in self-supervised vision transformers, CoRR abs/2104.14294 (2021).
 - [25] D. Misra, Mish: A self regularized non-monotonic neural activation function, CoRR abs/1908.08681 (2019).
 - [26] P. Gao, Z. Jiang, H. You, P. Lu, S. C. H. Hoi, X. Wang, H. Li, Dynamic fusion with intra-

and inter-modality attention flow for visual question answering, in: IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 6639–6648.

- [27] W. Wang, K. Chang, Y. Tang, Emotiongif-yankee: A sentiment classifier with robust model based ensemble methods, CoRR abs/2007.02259 (2020).
- [28] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 2020, pp. 8440–8451.