

MMSD2.0: Towards a Reliable Multi-modal Sarcasm Detection System

Libo Qin^{1*} Shijue Huang^{2,4*} Qiguang Chen⁵ Chenran Cai^{2,4}
Yudi Zhang Bin Liang^{2,4} Wanxiang Che⁵ Ruifeng Xu^{2,3,4} †

¹School of Computer Science and Engineering, Central South University, China

²Harbin Institute of Technology, Shenzhen, China ³Peng Cheng Laboratory, Shenzhen, China

⁴Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies

⁵Research Center for SCIR, Harbin Institute of Technology, Harbin, China

lbqin@csu.edu.cn, {joehsj310, crcai1023}@gmail.com, yvdi.zhang@outlook.com
bin.liang@stu.hit.edu.cn, {qgchen, car}@ir.hit.edu.cn, xurui Feng@hit.edu.cn

Abstract

Multi-modal sarcasm detection has attracted much recent attention. Nevertheless, the existing benchmark (MMSD) has some shortcomings that hinder the development of reliable multi-modal sarcasm detection system: (1) There are some spurious cues in MMSD, leading to the model bias learning; (2) The negative samples in MMSD are not always reasonable. To solve the aforementioned issues, we introduce MMSD2.0, a correction dataset that fixes the shortcomings of MMSD, by removing the spurious cues and re-annotating the unreasonable samples. Meanwhile, we present a novel framework called multi-view CLIP that is capable of leveraging multi-grained cues from multiple perspectives (i.e., text, image, and text-image interaction view) for multi-modal sarcasm detection. Extensive experiments show that MMSD2.0 is a valuable benchmark for building reliable multi-modal sarcasm detection systems and multi-view CLIP can significantly outperform the previous best baselines.

1 Introduction

Sarcasm detection is used to identify the real sentiment of user, which is beneficial for sentiment analysis and opinion mining task (Pang and Lee, 2008). Recently, due to the rapid progress of social media platform, *multi-modal sarcasm detection*, which aims to recognize the sarcastic sentiment in multi-modal scenario (e.g., text and image modalities), has attracted increasing research attention. Specifically, as illustrated in Figure 1, given the text-image pair, multi-modal sarcasm detection system predicts sarcasm label because the image shows a traffic jam that contradicts the text “*love the traffic*”. Unlike traditional sarcasm detection task, the uniqueness of multi-modal sarcasm detection lies in effectively modeling the consistency and sarcasm relationship among different modalities.

*Equally contributed authors.

† Corresponding author.



I love the traffic in this city so much ... #traffic_jam

Figure 1: Multi-modal sarcasm example. Box and words in red color denote the correlated sarcastic cues. Word #traffic_jam with hash symbol # is a hashtag.

Thanks to the rapid development of deep neural networks, remarkable success has been witnessed in multi-modal sarcasm detection. Specifically, Schifanella et al. (2016) makes the first attempt to explicitly concatenate the textual and visual features for multi-modal sarcasm detection. A series of works have employed attention mechanism to implicitly incorporate features from different modalities (Cai et al., 2019; Xu et al., 2020; Pan et al., 2020). More recently, graph-based approaches have emerged to identify significant cues in sarcasm detection (Liang et al., 2021, 2022; Liu et al., 2022), which is capable of better capturing relationship across different modalities, thereby dominating the performance in the literature.

While current multi-modal sarcasm detection systems have achieved promising results, it is unclear whether these results faithfully reflect the multi-modal understanding ability of models. In fact, when a text-modality only model RoBERTa is applied to multimodal sarcasm detection, its performance significantly surpasses the state-of-the-art multi-modal model HKE (Liu et al., 2022) with 6.6% improvement (See Detailed Analysis § 4.3). Such observation suggests that the performance of current models may heavily depend on spurious cues in textual data, rather than truly capturing the

relationship among different modalities, resulting in low reliability. Further exploration reveals that the characteristics of the MMSD benchmark (Cai et al., 2019) may be the cause of this phenomenon: (1) Spurious Cues: MMSD benchmark has some spurious cues (e.g., hashtag and emoji word) occurring in an unbalance distribution of positive and negative examples, which leads to the model bias learning; (2) Unreasonable Annotation: MMSD benchmark simply assigns the text without special hashtag (e.g., #sarcasm, etc.) as negative examples (i.e., not sarcastic label). We argue that this construction operation is unreasonable because sentence without #sarcasm tag can also express the sarcastic intention. Take utterance in Figure 2 as an example, the utterance without #sarcasm tag still belong to the sarcastic sample. Therefore, further chasing performance on current MMSD benchmark may hinder the development of reliable multi-modal sarcasm detection system.

Motivated by the above observation, we shift our eyes from traditional complex network design work (Cai et al., 2019; Liang et al., 2021, 2022) to the establishment of reasonable multi-modal sarcasm detection benchmark. Specifically, we introduce MMSD2.0 to address these problems. To solve the first drawback, MMSD2.0 removes the spurious cues (e.g., sarcasm word) from text in the MMSD, which encourages model to truly capture the relationship across different modalities rather than just memorize the spurious correlation. Such operation can benefit the development of bias mitigation in multi-modal sarcasm detection studies. To address the second problem, we directly make our efforts to re-annotate the unreasonable data. Specifically, given each utterance labeled with “not sarcastic” in the MMSD, we use crowdsourced workers to check and re-annotate the label. This correction process results in changes to over 50% of samples in original MMSD.

In addition to the dataset contribution, we propose a novel framework called multi-view CLIP, which can naturally inherit multi-modal knowledge from pre-trained CLIP model. Specifically, multi-view CLIP utilizes different sarcasm cues captured from multiple perspectives (i.e., text, image and text-image interaction view), and aggregates multi-view information for final sarcasm detection. Compared to previous superior graph-based approaches, multi-view CLIP has the following advantages: (1) multi-view CLIP does not

require any image pre-processing step for graph building (e.g., object detection); (2) multi-view CLIP does not require any complex network architecture and can naturally make full use of knowledge in VL pre-trained model for multi-modal sarcasm detection.

Contributions of this work can be summarized as follows:

- To the best of our knowledge, we make the first attempt to point out the behind issues in current multi-modal sarcasm benchmark, which motivates researchers to rethink the progress of multi-modal sarcasm detection;
- We introduce MMSD2.0, a correction dataset removing the spurious cues and fixing unreasonable annotation, for multi-modal sarcasm detection, which takes a meaningful step to build a reliable multi-modal sarcasm system;
- We propose a novel multi-view CLIP framework to capture different perspectives of image, text and image-text interaction, which attains state-of-the-art performance.

To facilitate future research, all datasets, code are publicly available at <https://github.com/JoeYing1019/MMSD2.0>.

2 MMSD2.0 Dataset Construction

MMSD2.0 aims to solve two main shortcomings of MMSD, which consists of (1) Spurious Cues Removal (§2.1) and (2) Unreasonable Samples Re-Annotation via Crowdsourcing (§2.2). This section provides MMSD2.0 dataset construction in detail and the overall process is illustrated in Figure 2.

2.1 Spurious Cues Removal

In our in-depth analysis, we observe that the spurious cues come from two resources: (1) hashtag words and (2) emoji words. Therefore, we will remove the two resources, respectively.

Hashtag Word Removal. In MMSD, as shown in Figure 3(a), we observe that distribution of hashtag word number in positive sample and negative sample is obviously unbalanced. As seen, the number of hashtag words in positive sample is on average more than 1 while less than 1 in negative samples in train, validation and test set. In other word, model only need to learn spurious correlation (hashtag word number) to make a correction

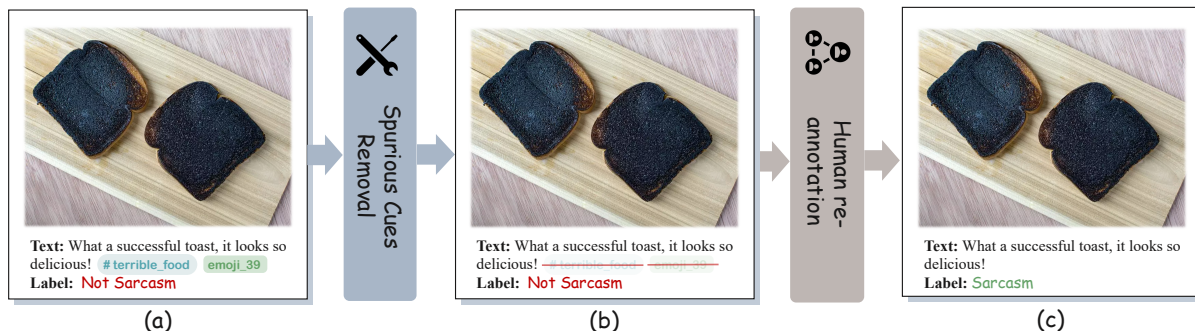




Figure 2: Overall process of construction MMSD2.0 dataset. Given the example in (a), Spurious Cues Removal stage first remove the spurious cues in text including hashtag word (#terrible_food) and emoji word (emoji_39) to acquire (b), then unreasonable samples re-annotation via crowdsourcing (human re-annotation) stage re-annotates the unreasonable samples to get final reasonable example (c).



(a) Average number of hashtag words per sample in MMSD dataset.



(b) Percentage of emoji words in MMSD dataset.

Figure 3: Text analysis of MMSD dataset.

prediction rather than truly understand multi-modal correlation in sarcasm detection.

To address the issue, we remove hashtag words from text in the MMSD dataset. This allows the model to capture image features and using them to guide the final prediction, rather than relying on the hashtag word number as a spurious cue.

Emoji Word Removal. Similarly, we find that the distribution of emoji words between positive and negative samples is also unbalanced. Specifically, as shown in Figure 3(b), only 19.3% of which exist in both positive and negative samples while the rest 80.7% of emoji words only appear in one type of sample (e.g., positive sample or negative sample). This indicates that model can simply use emoji word distribution as a priority for predicting

rather than truly capturing multi-modal cues.

To tackle this issue, we remove all the emoji words in text to force the model learning truly multi-modal sarcasm features rather than relying on spurious textual cues.

2.2 Unreasonable Samples Re-Annotation via Crowdsourcing

This section consists of two stages for re-annotation: (1) sample selection stage to choose the unreasonable samples and (2) re-annotation via crowdsourcing stage to fix the unreasonable samples according to the sample selection stage.

Sample Selection Stage. MMSD simply considers the samples without special hashtags like “#sarcasm” as negative sample (i.e., not sarcasm). In this work, we argue that the process is unreasonable because samples without #sarcasm tag can also express sarcasm intention. Therefore, we select all negative samples in the MMSD dataset (over 50%) as potentially unreasonable samples for further processing.

Re-annotation via Crowdsourcing. For all selected samples in sample selection stage, we directly re-annotate them via crowdsourcing by hiring human experts. Given a sample, each annotator is required to annotate with labels among: (i) Sarcasm: if the sample express sarcasm intention; (ii) Not Sarcasm: if the sample not express sarcasm intention; (iii) Undecided: if the sample is hard to decide by the annotator.

After the whole annotation is done, the samples with label Undecided is re-annotated again by three experts to decide the final annotation.

| MMSD/MMSD2.0 | Train | Validation | Test |
|--------------|---------------|-------------|-------------|
| Sentences | 19,816/19,816 | 2,410/2,410 | 2,409/2,409 |
| Positive | 8,642/9,572 | 959/1,042 | 959/1,037 |
| Negative | 11,174/10,240 | 1,451/1,368 | 1,450/1,372 |

Table 1: Comparison between MMSD and MMSD2.0.

2.3 Quality Control and Data Statistics

2.3.1 Quality Control

To control quality of the annotated dataset, we conduct two verification methods.

Onboarding Test. Before the whole annotation work, we require all the annotators to annotate 100 reasonable samples selected by model and the annotation result will be checked by 3 experts. Only those who achieve 85% annotation performance can join the annotation process.

Double Check. We randomly sample 1,000 annotated samples and ask another new annotator to re-annotate sarcasm label. Then, we calculated the Cohen’s Kappa (McHugh, 2012) between the previous labels and new labels. Finally, we get a kappa of 0.811 indicating an almost perfect agreement (Landis and Koch, 1977).

2.3.2 Data Statistics

Table 1 provides a detailed comparison of the statistics of the MMSD and MMSD2.0 datasets. We can observe that the distribution of positive and negative samples in MMSD2.0 is more balanced.

3 Approach

This section first describes the basic architecture for CLIP (§3.1) and then illustrates the proposed multi-view CLIP framework (§3.2).

3.1 CLIP

CLIP (Contrastive Language-Image Pretraining) (Radford et al., 2021), a powerful vision-and-language pre-trained model, has shown remarkable success on various VL downstream tasks (Conde and Turgutlu, 2021; Fu et al., 2022; Zhou et al., 2022). As shown in Figure 4, CLIP contains a visual encoder \mathbb{V} and a text encoder \mathbb{T} where former usually adopts ResNet (He et al., 2015) or ViT (Dosovitskiy et al., 2021) and the latter employs transformer (Vaswani et al., 2017) as backbone. After representation acquired by \mathbb{T} and \mathbb{V} , a dot-product function ($\mathbb{T}(\cdot) \cdot \mathbb{V}(\cdot)$) is further used to calculate the similarity of the given image and text

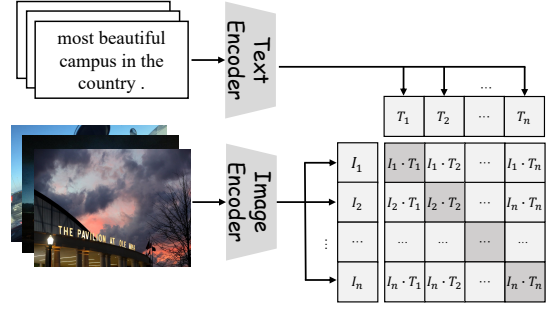


Figure 4: CLIP model.

and the highest score indicates the corresponding text and image is matched.

3.2 Multi-View CLIP

We introduce a novel multi-view CLIP framework for multi-modal sarcasm detection. Our framework consists of text view (§3.2.1), image view (§3.2.2) and image-text interaction view (§3.2.3), which explicitly utilize different cues from different views in CLIP model, thereby better capturing rich sarcasm cues for multi-modal sarcasm detection. The overall framework is shown in Figure 5.

3.2.1 Text View

A series of works (Xiong et al., 2019; Babanejad et al., 2020) have shown textual information can be directly used for performing sarcasm detection. Therefore, we introduce a text view module to judge the sarcasm from text perspective.

Given samples $\mathcal{D} = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^{N_D}$ where $\{\mathbf{x}, \mathbf{y}\}$ is a text-image pair, CLIP text encoder \mathbb{T} is used to output encoding representations \mathbf{T} :

$$\mathbf{T} = (t_1, \dots, t_n, t_{CLS}) = \mathbb{T}(\mathbf{x}), \quad (1)$$

where n stands for the sequence length of \mathbf{x} .

Then, text-view decoder directly employ t_{CLS} for multi-modal sarcasm detection:

$$\mathbf{y}^t = \text{softmax}(\mathbf{W}t_{CLS} + \mathbf{b}), \quad (2)$$

where \mathbf{y}^t is output distribution.

3.2.2 Image View

Image information can be considered as a textual supplement feature for sarcasm detection (Schifanella et al., 2016), which motivates us to propose an image CLIP view module to detect the sarcasm from image perspective.

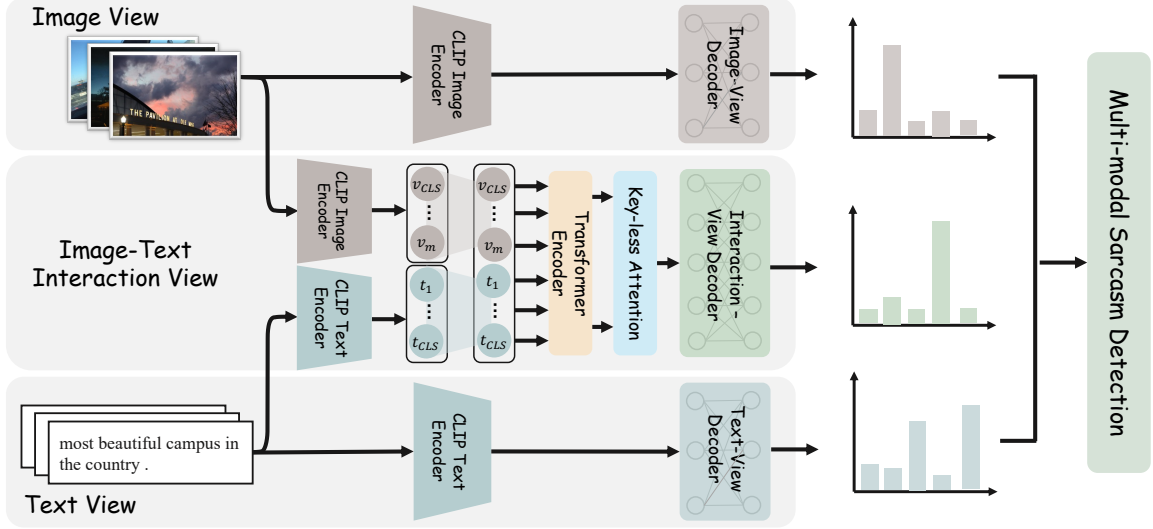


Figure 5: The overall framework of our multi-view CLIP. A pre-trained CLIP model encodes the inputs texts and images. Image view and text view utilize the information of image-only and text-only to capture sarcasm cues. Image-text interaction view fuses the cross-modal information. The three views is aggregated for final prediction.

Specifically, image view CLIP first leverage visual encoder \mathbb{V} to generate image representation:

$$\mathbf{I} = (\mathbf{v}_{CLS}, \mathbf{v}_1, \dots, \mathbf{v}_m) = \mathbb{V}(\mathbf{y}), \quad (3)$$

where m denotes the number of image patches.

Similarly, \mathbf{v}_{CLS} is used for prediction:

$$\mathbf{y}^v = \text{softmax}(\mathbf{W}\mathbf{v}_{CLS} + \mathbf{b}), \quad (4)$$

where \mathbf{y}^v is output distribution.

3.2.3 Image-text Interaction View

Modeling relationship across text and image modality is the key step to multi-modal sarcasm detection. We follow Fu et al. (2022) to use a transformer encoder (Vaswani et al., 2017) for sufficiently capturing interaction across different modalities.

Specifically, given the yielded text representation \mathbf{T} and image representation \mathbf{I} , we first concat them using $\mathbf{F} = (\mathbf{v}_{CLS}, \mathbf{v}_1, \dots, \mathbf{v}_m, \mathbf{t}_1, \dots, \mathbf{t}_n, \mathbf{t}_{CLS}) = \text{Concat}(\mathbf{T}, \mathbf{I})$. Then, we apply different linear functions to obtain the corresponding queries \mathbf{Q} , keys \mathbf{K} and values \mathbf{V} , and the updated representation $\hat{\mathbf{F}}$ can be denoted as $\hat{\mathbf{F}} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V}$, where d_k is the mapped dimension.

Given the updated image-text representation $\hat{\mathbf{F}} = (\hat{\mathbf{v}}_{CLS}, \hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_m, \hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_n, \hat{\mathbf{t}}_{CLS})$, we use key-less attention mechanism (Long et al., 2018) to further fuse the image-text interaction feature \mathbf{f} by calculating:

$$\mathbf{f} = \mathbf{p}_t \hat{\mathbf{t}}_{CLS} + \mathbf{p}_v \hat{\mathbf{v}}_{CLS}, \quad (5)$$

$$\mathbf{p}_t, \mathbf{p}_v = \text{softmax}(\mathbf{W}(\hat{\mathbf{t}}_{CLS}, \hat{\mathbf{v}}_{CLS}) + \mathbf{b}). \quad (6)$$

Finally, \mathbf{f} is utilized for prediction:

$$\mathbf{y}^f = \text{softmax}(\mathbf{W}\mathbf{f} + \mathbf{b}), \quad (7)$$

where \mathbf{y}^f is output distribution.

3.3 Multi-view Aggregation

Given the obtained $\mathbf{y}^t, \mathbf{y}^v, \mathbf{y}^f$, we adopt a late fusion (Baltrusaitis et al., 2019) to yield the final prediction \mathbf{y}^o :

$$\mathbf{y}^o = \mathbf{y}^t + \mathbf{y}^v + \mathbf{y}^f, \quad (8)$$

where \mathbf{y}^o can be regarded as leveraging rich features from different perspectives of text view, image view and image-text interaction view.

3.4 Model Training

We use a standard binary entropy loss for image, text, and image-text interaction view, and a joint optimization for training the entire framework:

$$\mathcal{L} = \sum_{i \in \{t, v, f\}} (\hat{\mathbf{y}}^i \cdot \log(\mathbf{y}^i) + (1 - \hat{\mathbf{y}}^i) \log(1 - (\mathbf{y}^i))), \quad (9)$$

where $\hat{\mathbf{y}}^i$ is the gold label. It is worth noting that we can directly use the final gold label for training image-view CLIP and text-view CLIP, which does not bring any annotation burden.

4 Experiments

4.1 Experimental Settings

We conduct experiments on MMSD (Cai et al., 2019) and MMSD2.0. We implement our model

| Model | MMSD | | | | MMSD2.0 | | | |
|---|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | Acc.(%) | P (%) | R (%) | F1 (%) | Acc. (%) | P (%) | R (%) | F1 (%) |
| <i>Text-modality Methods</i> | | | | | | | | |
| TextCNN (Kim, 2014) | 80.03 | 74.29 | 76.39 | 75.32 | 71.61 | 64.62 | 75.22 | 69.52 |
| Bi-LSTM (Graves and Schmidhuber, 2005) | 81.90 | 76.66 | 78.42 | 77.53 | 72.48 | 68.02 | 68.08 | 68.05 |
| SMSD (Xiong et al., 2019) | 80.90 | 76.46 | 75.18 | 75.82 | 73.56 | 68.45 | 71.55 | 69.97 |
| RoBERTa [†] (Liu et al., 2019) | 93.97 | 90.39 | 94.59 | 92.45 | 79.66 | 76.74 | 75.70 | 76.21 |
| <i>Image-modality Methods</i> | | | | | | | | |
| ResNet (He et al., 2015) | 64.76 | 54.41 | 70.80 | 61.53 | 65.50 | 61.17 | 54.39 | 57.58 |
| ViT (Dosovitskiy et al., 2021) | 67.83 | 57.93 | 70.07 | 63.40 | 72.02 | 65.26 | 74.83 | 69.72 |
| <i>Multi-Modality Methods</i> | | | | | | | | |
| HFM (Cai et al., 2019) | 83.44 | 76.57 | 84.15 | 80.18 | 70.57 | 64.84 | 69.05 | 66.88 |
| D&R Net (Xu et al., 2020) | 84.02 | 77.97 | 83.42 | 80.60 | – | – | – | – |
| Att-BERT (Pan et al., 2020) | 86.05 | 80.87 | 85.08 | 82.92 | 80.03 | 76.28 | 77.82 | 77.04 |
| InCrossMGs (Liang et al., 2021) | 86.10 | 81.38 | 84.36 | 82.84 | – | – | – | – |
| CMGCN (Liang et al., 2022) | 86.54 | – | – | 82.73 | 79.83 | 75.82 | 78.01 | 76.90 |
| HKE (Liu et al., 2022) | 87.36 | 81.84 | 86.48 | 84.09 | 76.50 | 73.48 | 71.07 | 72.25 |
| Multi-view CLIP | 88.33 | 82.66 | 88.65 | 85.55 | 85.64 | 80.33 | 88.24 | 84.10 |

Table 2: Main Results. Baseline results on dataset MMSD are taken from Liu et al. (2022). Results with - denote the code is not released. Results with † stand for that we re-implement the model.

based on the Huggingface Library (Wolf et al., 2020) and adopt clip-vit-base-patch32 as backbone. We use AdamW (Loshchilov and Hutter, 2019) as optimizer to optimize the parameters in our model. The batch size and epoch for training are 32 and 10, respectively. The learning rate for CLIP is $1e^{-6}$ and for the other parts is $5e^{-4}$. We select all hyper-parameters from the validation set. All experiments are conducted at Tesla V100s.

4.2 Baselines

Following Liu et al. (2022), we explore three type of baselines: (i) *text-modality methods*, (ii) *image-modality methods* and (iii) *multi-modality methods*.

(i) For *text-modality methods*, we employ (1) TextCNN (Kim, 2014) (2) Bi-LSTM (Graves and Schmidhuber, 2005) (3) SMSD (Xiong et al., 2019) and (4) RoBERTa (Liu et al., 2019).

(ii) For *image-modality methods*, our exploration including: (1) ResNet (He et al., 2015) and (2) ViT (Dosovitskiy et al., 2021), which are widely used backbone in CLIP.

(iii) For *multi-modality methods*, we compare multi-view CLIP with the following state-of-the-art baselines: (1) HFM (Cai et al., 2019) is a hierarchical fusion model for multi-modal sarcasm detection; (2) D&R Net (Xu et al., 2020) proposes a decomposition and relation network to model cross-modality feature; (3) Att-BERT (Pan et al., 2020) applies two attention mechanisms to model the text-only and cross-modal incongruity, respectively; (4) InCrossMGs (Liang et al., 2021) is a graph-based model using in-modal and cross-modal graphs to

capture sarcasm cues; (5) CMGCN (Liang et al., 2022) proposes a fine-grained cross-modal graph architecture to model the cross modality information; (6) HKE (Liu et al., 2022) is a hierarchical graph-based framework to model atomic-level and composition-level congruity. For a fair comparison, CMGCN and HKE are the versions without external knowledge.

4.3 Performance on MMSD

Following Liu et al. (2022), we adopt accuracy (Acc.), F1 score (F1), precision (P) and recall (R) to evaluate the model performance.

Table 2 (left part) illustrates the results on MMSD. We have the following observations: (1) *Text-modality methods* achieve promising performance and RoBERTa even surpasses multi-modality approaches, which indicates the sarcasm detection model can only rely on the text features to make a correct prediction, supporting the motivation of our work; (2) In *multi-modality approaches*, multi-view CLIP attains the best results, demonstrating the effectiveness of integrating features from different modality views.

4.4 Performance on MMSD2.0

Table 2 (right part) presents the results on MMSD2.0. As seen, we have the following observations: (1) Compared with the MMSD dataset, performance of all the baselines has decreased in different degrees on MMSD2.0 dataset. Especially the single-modality model RoBERTa’s performance decrease by 16.24% on F1 score. This indicates that MMSD2.0 successfully removes the spurious

| Model | Acc.(%) | P(%) | R(%) | F1(%) |
|---------------------------------|--------------|-------|-------|--------------|
| Ours | 85.64 | 80.33 | 88.24 | 84.10 |
| w/o $\mathcal{L}_{\mathcal{T}}$ | 84.18 | 80.60 | 83.32 | 81.93 |
| w/o $\mathcal{L}_{\mathcal{V}}$ | 83.69 | 76.97 | 88.62 | 82.38 |
| w/o $\mathcal{L}_{\mathcal{F}}$ | 82.44 | 73.80 | 91.80 | 81.82 |

Table 3: Ablation study of multi-view CLIP on MMSD2.0.

cues and can be used for a more reliable benchmark; (2) Additionally, the previous graph-based baselines such as CMGCN and HKE do not make good performance on MMSD2.0, even worse than the Att-BERT. We attribute it to the fact that such graph-based approaches highly rely on the spurious cues (e.g., hashtag and emoji word) when constructing the text semantic dependency graph, which cannot be accessed in MMSD2.0; (3) Lastly, multi-view CLIP not only attains the best results in multi-modality approaches but also in both text and image modality approaches, which further verifies the effectiveness of our framework.

4.5 Analysis

To understand the multi-view CLIP in more depth, we answer the following research questions: (1) Does each modality view contribute to the overall performance of the model? (2) What effect do the training strategies of CLIP have on the model performance? (3) What impact do different interaction approaches have on the model performance? (4) Does the multi-view CLIP approach remain effective in low-resource scenarios? (5) Why does the multi-view CLIP work?

4.5.1 Answer1: All Views Make Contribute To Final Performance

To analyze the effect of different modality views of our framework, we remove the training object of text view $\mathcal{L}_{\mathcal{T}}$, image view $\mathcal{L}_{\mathcal{V}}$ and image-text interaction view $\mathcal{L}_{\mathcal{F}}$ separately.

Table 3 illustrates the results. We can observe the accuracy decrease by 1.46%, 1.95% and 3.20% when removing $\mathcal{L}_{\mathcal{T}}$, $\mathcal{L}_{\mathcal{V}}$ and $\mathcal{L}_{\mathcal{F}}$, respectively. As seen, our framework attains the best performance when combining all these views. This suggests that all views in our framework make contribute to the final performance. It is worth noting that the removal of $\mathcal{L}_{\mathcal{F}}$ leads to a significant performance drop, which suggests that modeling the multi-modality interaction features is crucial for multi-modal sarcasm detection.




Figure 6: Different fine-tuning methods. Freeze All stands for both \mathbb{V} and \mathbb{T} are frozen, Freeze VE and Freeze TE indicate \mathbb{V} and \mathbb{T} are frozen respectively, and Full Finetuned means the whole CLIP is trainable.




Figure 7: Different interaction methods. Cross attention is based on self-attention but obtains query from one modality and key, value from another modality. MLP denotes applying a feed-forward layer after concatenating the information from different modalities.

4.5.2 Answer2: Full Finetuned Gains The Best Performance

To investigate the influence of training methods of backbone network CLIP, we conduct experiments on different training combination of \mathbb{V} and \mathbb{T} .

The results are shown in Figure 6. It reveals that full finetuning of CLIP leads to the best performance. An interesting observation is that freezing all of CLIP leads to almost the same performance as freezing only \mathbb{T} or \mathbb{V} . This can be attributed to the fact that the text and image representations in CLIP are aligned, and only training a single part can break this alignment property.

4.5.3 Answer3: Transformer Interaction Fuse Cross-modal Information Deeper

To further verify the effectiveness of our framework, we explore different interaction methods in text-image interaction view, including cross attention (Pan et al., 2020) and the MLP approach which concatenates information from different modalities and uses feedforward layer to fuse them.

Figure 7 illustrates the results. We can observe that our transformer interaction is more efficient than other two interactions. We attribute it to the fact that the transformer interaction is able to fuse the information from different modals in a more




Figure 8: Low-Resource Performance.

depth way, hence attaining the best performance.

4.5.4 Answer4: Multi-view CLIP can Transfer to Low-resource Scenario

To explore the effectiveness of multi-view CLIP in low-resource scenario, we experiment with different training sizes including 10%, 20% and 50%.

The results are shown in Figure 8, indicating that the multi-view CLIP approach still outperforms other baselines in low-resource settings. In particular, when the data size is 10%, the multi-view CLIP approach outperforms the baselines by a large margin in the same setting, and even surpasses Att-BERT when trained on the full dataset. We attribute this to the fact that the knowledge learned during the CLIP pre-training stage can be transferred to low-resource settings, indicating that our framework is able to extract sarcasm cues even when the training corpus is limited in size.

4.5.5 Answer5: Multi-view CLIP Utilize Correct Cues

To further explain why our framework works, we visualize the attention distribution of visual encoder \mathbb{V} to show why our framework is effective for multi-modal sarcasm detection.

As shown in Figure 9, our model can successfully focus on appropriate parts of the image with sarcasm cues. For example, our framework pays more attention on the no-fresh beans in Figure 9(a), which as the important cues to contradicted with “fresh” in text. Meanwhile, our framework focus on the bad weather regions in Figure 9(b), which also an important cues incongruent with “amazing” in text. This demonstrates that our framework can successfully focus on the correct cues.

5 Related Work

Sarcasm detection identifies the incongruity of sentiment from the context, which first attracts attention in text modality. Early studies focused on using feature engineering approaches to detect in-




Figure 9: Visualization.

congruity in text (Lunando and Purwarianti, 2013; Bamman and Smith, 2015). A series of works (Porria et al., 2016; Zhang et al., 2016; Xiong et al., 2019) explored deep learning network (e.g., CNN, LSTM and self-attention) for sarcasm detection. Babanejad et al. (2020) extended BERT by incorporating into it effective features for sarcasm detection. Compared to their work, we aim to solve sarcasm detection in the multi-modal scenario while their work focus on text sarcasm detection.

As the rapid popularization of social medial platform, multi-modal sarcasm detection attracts increasing research attention in recent years. Schifanella et al. (2016) explored the multi-modal sarcasm detection task for the first time and tackled this task by concatenating the features from text and image modalities respectively. Cai et al. (2019) proposed a hierarchical fusion model to fuse the information among different modalities and resealed a new public dataset. Xu et al. (2020) suggested to represent the commonality and discrepancy between image and text by a decomposition and relation network. Pan et al. (2020) implied a BERT architecture-based model to consider the incongruity character of sarcasm. Liang et al. (2021) explored constructing interactive in-modal and cross-modal to learn sarcastic features. Liang et al. (2022) proposed a cross-modal graph architecture to model the fine-grained relationship between text and image modalities. Liu et al. (2022) proposed a hierarchical framework to model both atomic-level congruity and composition-level congruity. In contrast to their work, we make the first attempt to address the bias issues in traditional MMSD dataset, towards to building a reliable multi-modal sarcasm detection system. In addition, we introduce MMSD2.0 to this end, aiming to facilitate the research. To the best of our knowledge, this is the first work to reveal the spurious cues in current multi-modal sarcasm detection dataset.

6 Conclusion

This paper first analyzed the underlying issues in current multi-modal sarcasm detection dataset and then introduced a MMSD2.0 benchmark, which takes the first meaningful step towards building a reliable multi-modal sarcasm detection system. Furthermore, we proposed a novel framework multi-view CLIP to capture sarcasm cues from different perspectives, including image view, text view, and image-text interactions view. Experimental results show that multi-view CLIP attains state-of-the-art performance.

Limitations

This work contributes a debias benchmark MMSD2.0 for building reliable multi-modal sarcasm detection system. While appealing, MMSD2.0 is built on the available MMSD benchmark. In the future, we can consider annotating more data to break through the scale and diversity of the original MMSD.

Ethics Statement

All data of MMSD2.0 come from the MMSD dataset [Cai et al. \(2019\)](#), which is an open-source dataset available for academic research. Our annotation process was carried out by annotators who were postgraduate students at Chinese universities and they were paid properly.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (62006062, 62176076, 62236004, 61976072), National Key R&D Program of China via grant 2020AAA0106501, Natural Science Foundation of Guangdong 2023A1515012922, the Shenzhen Foundational Research Funding (JCYJ20220818102415032), the Major Key Project of PCL2021A06, Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies 2022B1212010005.

References

- Nastaran Babanejad, Heidar Davoudi, Aijun An, and Manos Papangelis. 2020. [Affective and contextual embedding for sarcasm detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 225–243, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2019. [Multimodal machine learning: A survey and taxonomy](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(2):423–443.
- David Bamman and Noah A. Smith. 2015. Contextualized sarcasm detection on twitter. In *International Conference on Web and Social Media*.
- Yitao Cai, Huiyu Cai, and Xiaojun Wan. 2019. [Multi-modal sarcasm detection in Twitter with hierarchical fusion model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2506–2515, Florence, Italy. Association for Computational Linguistics.
- Marcos V. Conde and Kerem Turgutlu. 2021. [Clip-art: Contrastive pre-training for fine-grained art classification](#). In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3951–3955.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Jinmiao Fu, Shaoyuan Xu, Huidong Liu, Yang Liu, Ning Xie, Chien-Chih Wang, Jia Liu, Yi Sun, and Bryan Wang. 2022. [Cma-clip: Cross-modality attention clip for text-image classification](#). In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 2846–2850.
- Alex Graves and Jürgen Schmidhuber. 2005. [Framewise phoneme classification with bidirectional lstm and other neural network architectures](#). *Neural Networks*, 18(5):602–610. IJCNN 2005.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- J. Richard Landis and Gary G. Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Bin Liang, Chenwei Lou, Xiang Li, Lin Gui, Min Yang, and Ruifeng Xu. 2021. [Multi-modal sarcasm detection with interactive in-modal and cross-modal graphs](#). In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 4707–4715, New York, NY, USA. Association for Computing Machinery.

- Bin Liang, Chenwei Lou, Xiang Li, Min Yang, Lin Gui, Yulan He, Wenjie Pei, and Ruifeng Xu. 2022. [Multi-modal sarcasm detection via cross-modal graph convolutional network](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1767–1777, Dublin, Ireland. Association for Computational Linguistics.
- Hui Liu, Wenya Wang, and Haoliang Li. 2022. [Towards multi-modal sarcasm detection via hierarchical congruity modeling with knowledge enhancement](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xiang Long, Chuang Gan, Gerard Melo, Xiao Liu, Yandong Li, Fu Li, and Shilei Wen. 2018. [Multimodal keyless attention fusion for video classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Edwin Lunando and Ayu Purwarianti. 2013. [Indonesian social media sentiment analysis with sarcasm detection](#). In *2013 International Conference on Advanced Computer Science and Information Systems (ICACSIS)*, pages 195–198.
- Mary L. McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia Medica*, 22:276 – 282.
- Hongliang Pan, Zheng Lin, Peng Fu, Yatao Qi, and Weiping Wang. 2020. [Modeling intra and inter-modality incongruity for multi-modal sarcasm detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1383–1392, Online. Association for Computational Linguistics.
- Bo Pang and Lillian Lee. 2008. [Opinion mining and sentiment analysis](#). *Found. Trends Inf. Retr.*, 2(1–2):1–135.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, and Prateek Vij. 2016. [A deeper look into sarcastic tweets using deep convolutional neural networks](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1601–1612, Osaka, Japan. The COLING 2016 Organizing Committee.
- 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 *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Rossano Schifanella, Paloma de Juan, Joel Tetreault, and LiangLiang Cao. 2016. [Detecting sarcasm in multimodal social platforms](#). In *Proceedings of the 24th ACM International Conference on Multimedia, MM '16*, page 1136–1145, New York, NY, USA. Association for Computing Machinery.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface's transformers: State-of-the-art natural language processing](#).
- Tao Xiong, Peiran Zhang, Hongbo Zhu, and Yihui Yang. 2019. [Sarcasm detection with self-matching networks and low-rank bilinear pooling](#). In *The World Wide Web Conference, WWW '19*, page 2115–2124, New York, NY, USA. Association for Computing Machinery.
- Nan Xu, Zhixiong Zeng, and Wenji Mao. 2020. [Reasoning with multimodal sarcastic tweets via modeling cross-modality contrast and semantic association](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3777–3786, Online. Association for Computational Linguistics.
- Meishan Zhang, Yue Zhang, and Guohong Fu. 2016. [Tweet sarcasm detection using deep neural network](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2449–2460, Osaka, Japan. The COLING 2016 Organizing Committee.
- Yangming Zhou, Qichao Ying, Zhenxing Qian, Sheng Li, and Xinpeng Zhang. 2022. [Multimodal fake news detection via clip-guided learning](#).

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section Limitations
- A2. Did you discuss any potential risks of your work?
This work introduces a correction dataset to fix behind issues in current multi-modal sarcasm benchmark and proposes a novel framework multi-modal sarcasm detection, which may not contain potential risks.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section Abstract and Section 1 Introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 2 MMSD2.0 Dataset Construction and Section 3 Approach

- B1. Did you cite the creators of artifacts you used?
Section 1 Introduction
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section Ethics Statement
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
All the data are based on publicly available and used safely by previous works.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We will provide the documentation of our code and models in our code repo.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 2.3 Quality control and Data Statistics

C Did you run computational experiments?

Section 4 Experiments

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4 Experiments

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 4 Experiments

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4 Experiments

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Not applicable. Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Section Ethics Statement

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Section Ethics Statement

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Section Ethics Statement

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Section Ethics Statement

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Section Ethics Statement