

Image Matters: A New Dataset and Empirical Study for Multimodal Hyperbole Detection

Huixuan Zhang, Xiaojun Wan

Wangxuan Institute of Computer Technology, Peking University
State Key Laboratory of Media Convergence Production Technology and Systems
zhanghuixuan@stu.pku.edu.cn, wanxiaojun@pku.edu.cn

Abstract

Hyperbole, or exaggeration, is a common linguistic phenomenon. The detection of hyperbole is an important part of understanding human expression. There have been several studies on hyperbole detection, but most of which focus on text modality only. However, with the development of social media, people can create hyperbolic expressions with various modalities, including text, images, videos, etc. In this paper, we focus on multimodal hyperbole detection. We create a multimodal detection dataset from *Weibo* (a Chinese social media) and carry out some studies on it. We treat the text and image from a piece of weibo as two modalities and explore the role of text and image for hyperbole detection. Different pre-trained multimodal encoders are also evaluated on this downstream task to show their performance. Besides, since this dataset is constructed from five different keywords, we also evaluate the cross-domain performance of different models. These studies can serve as a benchmark and point out the direction of further study on multimodal hyperbole detection.

Keywords: multimodal, hyperbole, datasets

1. Introduction

As defined by Merriam Webster, exaggeration means “an act or instance of exaggerating something : overstatement of the truth”¹ and hyperbole has the same meaning. Though hyperbolic expressions state something beyond fact or truth, they are not considered as lies. On the contrary, hyperbole is a common way of expressing one’s strong feeling or opinion. As the second most used rhetorical device (Kreuz et al., 1996), the detection of hyperbole bears great importance for understanding human language and expression.

Traditional studies focus more on text modality only. Their proposed datasets and methods considered mainly single short sentences (Troiano et al. (2018) Kong et al. (2020), Biddle et al. (2021)) However, the expression of hyperbole is not limited to text only. For example, we can hardly know whether the expression “*Winter comes early.*” is hyperbolic or not unless we can see whether it is indeed so cold (Figure 1a). Images, as widely seen on social media, serve an important role in expressing exaggerated meanings. They can serve as facts to reveal hyperbole contained in texts (Figure 1a) as well as express hyperbole themselves (Figure 1b) or together with texts (Figure 1c). Meanwhile, some images can also just serve as background information and does no help to the expression of hyperbole (Figure 1d), which adds additional difficulty to this task.

*Weibo*² is one of the most popular social media in

China and all over the world. With abundant publicly available posts with texts and images, it serves as the perfect resource for real-life multimodal hyperbole detection. We just need to automatically crawl enough posts from it and annotate them as hyperbole or not.

Our contributions in this paper can be summarized as follows: (1) We propose the first Chinese multimodal hyperbole detection dataset which can be used for further study.³ (2) We analyze the dataset and get some statistical observations. We also summarize how two modalities (text and image) together express hyperbole. (3) We employ several neural methods on hyperbole detection, including unimodal and multimodal encoders, and reach these conclusions: (3.1) Image modality is more useful rather than misleading. With gating and attention mechanism, we can achieve relatively fine performance. We also point out that common sense is the potential direction to achieve better results. (3.2) Compared with unimodal encoders and fusing methods, pre-trained multimodal methods perform badly on this downstream task. We analyzed its reason. (3.3) The dataset is built upon five different and disjoint keywords with generally disjoint contents. When applied on cross domains, the methods don’t work very well. The systematic bias between posts with different keywords is the bottleneck hindering the generalization ability, especially for methods requiring deep interactions between modalities.

¹<https://www.merriam-webster.com/dictionary/exaggeration>

²<https://weibo.com>

³Dataset can be found in https://github.com/1leoZhang/Multimodal_Hyperbole.




(a) 这个天气是提前过冬
(Winter comes early.)



(b) 天气太热了
(It is too hot.)



(c) 心情是想杀了全世界
(My mood is that I'd like to kill everyone in the world.)
(The words in the image says "I did not piss off anyone.")



(d) 这个景区周六周日来就是要命! 堵车能堵一小时
(It's driving us to death to go to this scenic spot on weekend. We spent hours on traffic jams.)

Figure 1: Examples of image aiding the detection of hyperbole or serving as background information(all the examples are hyperbole). (a). In the image we can see that it is not so cold. (This is the surface of a weather application which says it is $13^{\circ}C$ in Urumqi. It's much colder in winter in Urumqi.) So we must refer to both text and image to correctly determine the post as hyperbole.; (b).The image is necessary because the text just states some fact. But with the melting face we can confidently label it as hyperbole because however hot it is, people do not melt.; (c). The grey background and stiff face in the image increases the confidence of annotating it as hyperbole.;(d).The image is just a screenshot of a navigating software. The hyperbolic meaning here comes from that people do not die from just going to this scenic spot, which cannot be revealed by this image.

2. Related Works

2.1. Studies on Hyperbole and Hyperbole Detection

Not surprisingly, earlier studies on hyperbole focus more on its linguistic features. Mora (2009) made a detailed analysis of the semantic features of hyperbole and first introduced a taxonomy which generally categorizes all exaggerations along two dimensions: quantity(overstating the quantitative or objective property of something, such as time, mass, etc.) and quality(overstating the qualitative or subjective property of something, such as feeling, emotion, etc.). These two types can be subcategorized into smaller ones, but that is beyond our discussion.

Some multimodal analysis work has also been conducted. Ferré (2014) manually analyzed how hyperbole is expressed with text, video and audio modalities. He focused on gestures, facial expressions, tones, etc. His work directly pointed out that other modalities can help reveal hyperbole contained in texts or words.

For the automatic detection of hyperbole, Troiano et al. (2018) created a hyperbole detection dataset HYPO, which contains 709 hyperbolic sentences and their corresponding non-hyperbolic ones. They applied traditional machine learning methods and proposed some hand-crafted features (named QQ) to address this task.

Biddle et al. (2021) followed Troiano's work. They introduced triplet sampling and used triplet loss for forcing their model to differentiate hyperbolic and non-hyperbolic sentences. They used BERT as the

backbone text encoder instead of pre-trained word vectors(used by Troiano's work).

As for Chinese hyperbole detection, Kong et al. (2020) followed Troiano's idea and created the first Chinese hyperbole detection dataset HYPO-cn. They analyzed statistical features of the sentences in the dataset thoroughly. They also manually analyzed strategies humans used to express hyperbole. They compared several traditional and deep-learning methods on this dataset. They reached the conclusion that deep-learning methods outperform traditional methods on this task.

However, to the best of our knowledge, there hasn't been any study on multimodal hyperbole detection.

2.2. Other Relevant Multimodal Tasks

Multimodal sarcasm detection has become a hot topic in recent years. Cai et al. (2019) first proposed a dataset built on Twitter, which is still widely used by researchers. They also proposed a hierarchical fusion model to address this task. Xu et al. (2020) followed this work and proposed D&R Net. Liang et al. (2021) introduced in-modal and cross-modal graphs(InCrossMGs) to leverage the ironic expressions from both in- and cross-modal perspectives by extracting the dependency relations of fine-grained phrases or patches within the distinct modality and across multiple modalities. Liu et al. (2022) leverages graph attention networks and image captions as knowledge enhancement. On the other side, Castro et al. (2019) focused on video and text modalities and proposed a sarcasm detection dataset named MUSTARD. MUSTARD is compiled from popular TV shows and consists

of audiovisual utterances annotated with sarcasm labels along with its corresponding context. Wu et al. (2021) continued on this work and proposed IWAN to model the word-level incongruity between modalities via a scoring mechanism.

Multimodal sentiment analysis is also a widely-discussed topic. Wang et al. (2020) proposed an end-to-end fusion method based on transformers. Yang et al. (2021) introduced multi-channel graph neural networks. Xue et al. (2022), Du et al. (2022) focused on attention-based methods such as gated attention. Inspirations can be gained from these methods.

Multimodal hyperbole detection bears both similarities and differences with the tasks mentioned above. It requires the understanding and interaction of both modalities. However, the relation between the two modalities can be rather complicated. This will be discussed in Section 4.

3. Dataset Creation

3.1. Hyperbole Definition

To avoid ambiguity, we need to give a clear definition to hyperbole. Based on the definition given by Merriam Webster ⁴ and Troiano et al. (2018), we here define hyperbole as “an expression (words, images, etc) that goes significantly beyond fact or common sense but not taken as lies”. A post is considered hyperbolic if and only if it contains hyperbolic expression.

For example, the post “*I’m dying of heat*” is hyperbolic because normally most people do not easily die of heat. But if there is a post “*5 workers died of extreme heat*” with real pictures, the post is certainly not considered hyperbolic. However, if a post says “*100 billion workers died of extreme heat*”, it is then considered neither hyperbole or non-hyperbole. Instead, it is probably a lie. As for images, Figure 1b provides a perfect example. Since people do not melt however hot it is, a melting face representing the unbearably hot weather is typical hyperbole.

3.2. Data Collection and Preprocessing

Weibo is a popular Chinese social medium. We automatically crawl about 10000 posts from Weibo with 5 keywords: 堵车(traffic jam), 风景(scenery), 天气(weather), 心情(mood), 衣服(clothing). Each keyword contains roughly same number of posts (about 2000). We choose the five keywords because they are generally disjoint and close to daily lives. We carefully set criteria so that each post crawled is original and with at least one image. After collecting these posts, we first clean up the data. The urls in the texts are replaced with a specific token `<url>`. Some emojis are actually hyperlinks to images, so they are replaced with another

⁴<https://www.merriam-webster.com/dictionary/exaggeration>

Keyword	Total Count	Hyperbole
堵车(traffic jam)	496	252
风景(scenery)	402	204
天气(weather)	396	194
心情(mood)	451	207
衣服(clothing)	442	198
total	2160	1055

Table 1: Basic statistics of our dataset

`<emoji>` token. Hashtags (words or phrases led by #) are not removed because there is no obvious distinction between hashtags in hyperbolic and non-hyperbolic posts according to our observation. Some posts have more than one image, so we randomly select one image and drop the others. Since this is done before any human work, there is no concern that it will introduce any inconsistency or bias. Then we conduct a rough selection. Posts with too long or unreadable texts or unclear pictures are manually dropped to save some annotation work. (Here we define “too long” as posts longer than 200 characters.) We set the limitation because longer posts are harder for annotation. Annotators may miss the hyperbolic part of a too long post thus leading to unreliable results. We believe this setting makes our dataset cleaner. Advertisement posts that appear repeatedly are also dropped to avoid polluting the dataset.

3.3. Data Annotation

We employ several human annotators to annotate these posts as hyperbole or non-hyperbole. All of them are Chinese college students. They are required to follow an annotation specification given by us. We first ask them to label 200 randomly chosen posts, and select two of them whose annotation results fit the specification best. We communicate with them about their mistakes to help them understand our definition of hyperbole better. They then annotate the rest posts. We manually check each inconsistent case and discuss with them to determine whether each is hyperbole.

Since there are much more non-hyperbolic posts, we randomly choose from them and create the final dataset in which #hyperbole:#non-hyperbole is approximately 50%:50%. Some basic statistics of the final dataset are listed in Table 1.

4. Dataset Analysis

In this section, we make some analysis on our constructed dataset.

First, we consider the length of the text. In Chinese, the length is simply the number of characters in the text. We assume that hyperbolic posts and non-hyperbolic posts may have similar text length. But the result turns out that hyperbolic posts have significantly longer texts than non-hyperbolic ones.

What's more, compared with non-hyperbolic ones, hyperbolic posts of different keywords also have more varied text length. The average text length of different keywords are shown in Table 2. As can be seen, non-hyperbolic posts have texts with relatively similar lengths among different keywords, while hyperbolic posts have quite varied text length among keywords.

We guess that there are two main reasons leading to this phenomenon. First, hyperbolic expressions are usually related with strong emotions. People tend to use more words to fully express their strong feelings. The other reason is that some hyperbolic sentences with certain keywords bear some kind of similarity. Taking figure 2 as an example, there are several hyperbolic posts of keyword 风景(scenery) with this pattern to express their love towards a star. This pattern does not appear in other keywords or in most non-hyperbolic posts, which partly explains why hyperbolic posts contain longer texts and why hyperbolic posts have varied text length among different keywords.



<url> <url>鹿晗，晚安<url>#全世界最帅的老板# 鹿晗<url>#gucci品牌代言人鹿晗# <url> @M鹿M 你是我看不完的风景，唱不完的情歌，写不出的浪漫，永远都想捧在手心里的宝藏男孩 (Good Night, Lu Han. #The most handsome boss in the world# Lu Han. #Gucci brand spokesman Lu Han.# You are my endless scenery, endless love song, endless romance, the treasure boy I want to hold in my hands forever.)

Figure 2: An example of hyperbolic post expressing their love towards a star with the word 风景(scenery).(There are many much longer examples, for conciseness we choose a relatively short one here.)

Then we focus on lexical features of hyperbolic posts. There's no doubt that some words appear more often in hyperbolic posts than non-hyperbolic ones. To quantitatively examine this phenomenon, We calculate smoothed chi-score (Yang and Pedersen, 1997) $\tilde{K}(d)$ of each word d . The bigger $\tilde{K}(d)$ is, the higher correlation between d and hyperbole is. For details about smoothed chi-score, please refer to appendix A. Since there are no spaces

between words in Chinese, we use *Spacy*⁵ to separate words and remove stop words. We calculate $\tilde{K}(d)$ for all remaining words. The top-10 words are listed in Table 3. As can be seen, one can easily think of hyperbolic expressions with some words such as 死(death). Other words are corresponding to certain keywords, such as 冷(cold) is closely related to 天气(weather). However, when observing these words alone, most of them does not contain hyperbolic meaning themselves, which means that there is no reason to determine whether a post is hyperbole through some kind of “keywords”, thus further proving the difficulty of this task. It is also interesting that though we do not consider emotional exaggeration(such as the expression “Ah Ah it's so cold!”) as real hyperbole (According to our definition, this expression says nothing beyond fact), 啊啊(AhAh) is still closely related to hyperbole. It meets our assumption above that hyperbole is related to strong emotions.

Finally, we consider the roles that images play in expressing hyperbole. We ask the annotators to check all the images alone, trying to determine whether the post can be classified as hyperbole with image modality only. We find out that in all hyperbolic posts, only 15.8% images are themselves hyperbolic (like the one in Figure 1b). We also need to note that some images are considered hyperbolic because they contain hyperbolic texts (Figure 3). We count roughly and find there are only 4% of all hyperbolic posts that belong to this category, which is acceptable. This result reveals that images mostly serve as an assistant when expressing hyperbole rather than being hyperbolic themselves. As is discussed above, images can serve as ground truth to reveal hyperbole (Figure 1a) or give us higher confidence to determine one post as hyperbole (Figure 1c). According to this observation, we may question whether introducing image modality is helpful in recognizing hyperbole, which will be discussed in detail in 5.1.



完全没心情在社交平台活跃。物流走了2个半月了还没到，我的心好疲惫。(Completely no mood for staying active on social media. It's been two and a half months and my express still hasn't arrived. My heart feels so tired.)

Figure 3: An example of images with hyperbolic words in it. (The words in the images say “I hate the whole world before sleep.”)

⁵<https://github.com/explosion/spaCy>

	堵车(traffic jam)	风景(scenery)	天气(weather)	心情(mood)	衣服(clothing)
non-hyperbole	61.3	56.8	57.7	64.3	70.9
hyperbole	81.0	96.1	79.3	65.5	88.3

Table 2: Average text lengths of different keywords.

word(d)	$\tilde{K}(d)$
死(death)	0.00412
冷(cold)	0.00372
买到(bought)	0.00372
腿(leg)	0.00359
弄(cause)	0.00359
适合(fit)	0.00342
啊啊(AhAh)	0.00342
加(plus)	0.00342
头发(hair)	0.00342
开学(term begins)	0.00342

Table 3: Top 10 words with highest $\tilde{K}(d)$

In summary, image and text can express hyperbole in the following ways: (1) one modality expresses hyperbole itself and the other modality serves as background information (*case-I*); (2) two modalities together express hyperbolic meanings, one expressing hyperbole and the other serving as ground truth to reveal this hyperbole (*case-II*); (3) two modalities together express hyperbolic meaning, one expressing hyperbole and the other offering extra confidence (*case-III*).

5. Empirical Studies on Hyperbole Detection

We treat the task of multimodal hyperbole detection as a binary classification task with hyperbole as positive sample. In this section, we aim to investigate and answer the following three questions related to this task:

- **RQ1:** Is image modality more useful rather than misleading in multimodal hyperbole detection?
- **RQ2:** Are pre-trained multimodal encoders effective on this task?
- **RQ3:** Can typical methods achieve good performance on cross-domain circumstances?

In the next three subsections, we will conduct empirical studies to discuss the above three questions, respectively.

5.1. Experiments on Different Modalities (RQ1)

As discussed in Section 4, images mostly serve as assistants rather than express hyperbole themselves. So it is natural to ask: is image modality more useful than misleading in hyperbole detection? If the answer is yes, which way is more

suitable for utilizing image? For example, is deep modality fusion more helpful than shallow ones?

We describe our methods from shallower fusion to deeper fusion. First, we use two unimodal encoders: BERT (Devlin et al., 2019) and ResNet50 (He et al., 2016), to encode text and image modality individually. Specifically, we use BERT that is pre-trained on Chinese (Cui et al., 2020). We fine-tune them to achieve better performance.

After unimodal encoding, we apply different fusion methods on this task. Denote f_{text} as the feature vector of text (outputted by BERT) and f_{image} as the feature vector of image (outputted by ResNet), the most trivial idea is simply concatenating f_{text} and f_{image} (We name this method *concat*). A slight improvement on this idea is implementing a gating method to select more significant features. We name this method *gate*. To be more detailed, consider two feature vectors f_{image} and f_{text} , the gating method we use can be summarized as:

$$g_{text} = \sigma(W_{text}f_{image} + b_{text})$$

$$g_{image} = \sigma(W_{image}f_{text} + b_{image})$$

$$output = concatenate(g_{text} \otimes f_{text}, g_{image} \otimes f_{image})$$

where $W_{text}, W_{image}, b_{text}, b_{image}$ are learnable parameters, σ is the *Sigmoid* function ($\sigma(x) = \frac{1}{1+e^{-x}}$). Attention is a widely used method which can utilize fine-grained features. We utilize both text-level and token-level features outputted by BERT as well as both image-level and patch-level features outputted by ResNet. We implement a multi-head attention mechanism to extract cross-modality fine-grained features. To be more detailed, we use the fine-grained features (i.e., token-level or patch-level features) of one modality as Q , the fine-grained features of the other modality as K and V . Thus we obtain two cross-modality fine-grained feature representations. We use another two multi-head attention modules to further aggregate these fine-grained feature representations respectively (using image- or text-level feature representation as Q and the corresponding cross-modality fine-grained feature representation as K and V) and get the encoding vector of either modality. We then concatenate these two encodings and implement a gating mechanism similarly (We name this method *attn-gate*.) The main architecture of this method is shown in Figure 4.

After fusing with either of the above methods, we feed the feature vector into several fully connected

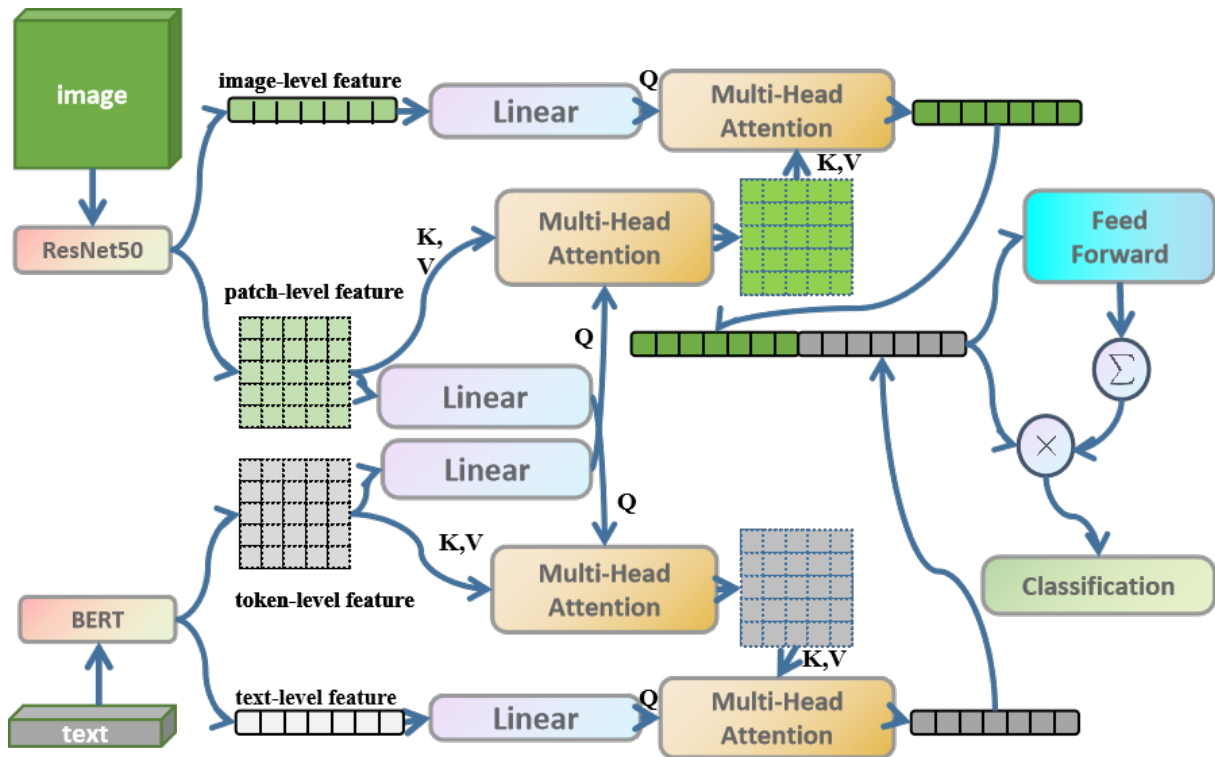


Figure 4: The architecture of the *attn-gate* fusion model.

layers to finish the final classification. We use cross-entropy as loss function and Adam (Kingma and Ba, 2014) as optimizer. All these models share the same hyper-parameters. We believe the above methods are enough to answer our question on this relatively small dataset.

We report 10-fold cross validation results of the models to ensure the reliability of our results. Our hyper-parameters are listed in Table 4. The average results are shown in Table 5.

Batchsize	32
Learning rate	1e-4
Epochs	100
Dropout	0.5
Random seed	1008
Number of heads	4

Table 4: Hyper parameters used.

Some conclusions can be drawn from the results. First, with single modality, text-only model (BERT) performs much better than image-only model (ResNet). This is compatible with our analysis that most images are not hyperbolic themselves. Therefore although we’ve already discussed the importance of image modality for humans to determine hyperbole, it’s still questionable that whether image modality is more useful rather than misleading under this circumstance.

Our following result proves that image modality is indeed more useful rather than misleading. Though

model	F1-score	Accuracy
BERT	0.702 \pm 0.174	0.713 \pm 0.027
ResNet	0.553 \pm 0.078	0.495 \pm 0.028
Concat	0.734 \pm 0.199	0.743 \pm 0.017
Gate	0.742 \pm 0.193	0.749 \pm 0.029
Attn-Gate	0.750 \pm 0.178	0.758 \pm 0.018

Table 5: Average results with 10-fold cross-validation (BERT and ResNet use only single-modality information.). The subscripts show the standard deviation of each experiment.

with text modality only can achieve a fine result, introducing image modality, even simply giving it the same significance as text modality (*concat*), still apparently increases the performance. Comparing with *BERT*, this improvement is *statistically significant*. This result clearly shows that introducing image modality together with text modality is useful for determining hyperbole.

Further, gating mechanism may be useful for increasing model performance. This insight can be drawn from our analysis of the dataset. Since under most circumstances, there are probably some features serving as background information, we should give these features less focus. As can be seen from the results, even with trivial gating mechanism, the result is slightly improved compared with simply concatenating the two modalities’ outputs. However, this trivial gating method still has its own problems. First, using BERT only can achieve relatively



(a) 这个天气是提前过冬 (Winter comes early.)
 (b) 想看个天气参考一下买什么衣服。看完后，算了，不买了，这天还穿什么衣服<emoji> (Referring to weather forecast to see which clothes to buy. Now I think there's no need wearing any clothes in such weather.)

Figure 5: Two examples of hyperbole cases.

fine results, which shows that text modality contains more information for detecting hyperbole. Thus we observe that the trivial gating method tends to give text modality a higher priority. This systematic bias leads to insufficient utility of image modality. Another problem is that *gate*, as well as *concat*, is unable to model the fine-grained relationship between text and image. Figure 5a gives an example. All other methods except *attn-gate* failed on this case. Since they just simply use high-level features of two modalities, they are unable to notice the connection between certain tokens (*Winter*) and a certain feature (the temperature in the image). In contrast, *attn-gate* is able to extract fine-grained features between two modalities. In fact, comparing with the trivial *concat* method, *attn-gate* is *statistically significantly* better. Since we performed 10-fold cross validation, we believe that this improvement clearly shows that even though most hyperbole are not directly expresses through images, introducing images, especially using deep fusion, is still helpful in recognizing hyperbole. There are also cases where all these methods failed. Figure 5b is a typical failure case. The reason is that these models do not know even under such hot weather, people still need to wear clothes. The expression “there’s no need wearing any clothes” is an overstatement of how hot it is. Since models do not have any common sense knowledge, they are unable to recognize this kind of hyperbole, which leads to the failure. In summary, our answer to this question (RQ1) is that: images are more useful rather than misleading. With gating and attention mechanism, we can more accurately utilize fine-grained fused features between two modalities, which can help identify hyperboles of *case-II* and *case-III*. However, without common sense, the model can never fully understand hyperbole. To further improve the perfor-

mance, the key is introducing common sense into the models.

5.2. Pre-trained Multimodal Encoders Evaluation (RQ2)

Pre-trained multimodal models are proved effective in many downstream tasks including zero-shot learning, image-text retrieval, VQA, etc. However, when speaking of multimodal classification tasks, few discussions are made before. In this subsection, we apply two pre-trained models CLIP (Radford et al., 2021) and BriVL (Huo et al., 2021) on this downstream task to evaluate their performance. CLIP (Radford et al., 2021) is a popular multimodal encoder and BriVL (Huo et al., 2021) is a similar one that is pre-trained on Chinese. Besides, BriVL claims to use a more advanced algorithm (inspired by MoCo (He et al., 2020)) to obtain better encodings of both modalities.

We still report 10-fold cross-validation results. Specifically, we use CLIP that is pre-trained on Chinese texts (Yang et al., 2022). For BriVL, we directly use the open-source code and tools⁶.

We try two ways to use these pre-trained models. One is giving prompts (*prompt*) to the text (We add prompts “这是夸张”(This is hyperbole.) and “这不是夸张”(This is not hyperbole) at the beginning of each text. So we get two texts with different prompts and same content.). To be more detailed, for each case we encode one image and the two texts with different prompts, and try to maximize the cosine similarity between the image and the text with correct prompt. This is a common method used in image classification tasks.

Another way is using these models simply as encoders and treat their output as feature vectors. Then we concatenate them (*concat*) or use gating methods (*gate*) just like what we have done above. In this way no prompts are added to texts. Both the two ways are seldom discussed in similar tasks like multimodal sarcasm detection previously. Details of both ways are shown in figure 6.

We fine-tune these models to achieve better results. However, the result turns out to be very disappointing. As can be seen in Table 6, they perform much worse than most methods mentioned in Section 5.1.

model	F1-score	Accuracy
CLIP+prompt	0.632 \pm 0.159	0.642 \pm 0.042
CLIP+concat	0.584 \pm 0.165	0.644 \pm 0.035
CLIP+gate	0.580 \pm 0.174	0.642 \pm 0.025
BriVL+prompt	0.587 \pm 0.170	0.628 \pm 0.025
BriVL+concat	0.665 \pm 0.165	0.667 \pm 0.038
BriVL+gate	0.644 \pm 0.162	0.637 \pm 0.031

Table 6: Results of multimodal models.

⁶<https://github.com/BAAI-WuDao/BriVL>

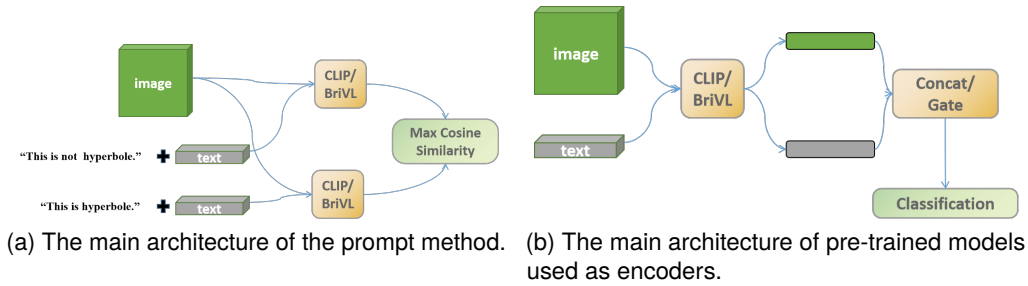


Figure 6: The architecture of the two methods using pre-trained models.

We analyze why the results are much worse. Radford et al. (2021) has already pointed out that CLIP performs badly on classifying abstract concepts like determining whether a scene is normal or not. Hyperbole detection is a typical task involving abstract concept, so it is not surprising that *prompt* does not work well on this task. What’s more, since this kind of task involves classifying image and text together rather than image itself, the prompt method is even more unsuitable. We believe the same reason works for BriVL.

Besides, the text in this kind of task is usually not a simple description of the image. It is a common case that there is no strong semantic correlation between corresponding text and image, which is quite different from the dataset CLIP is trained on. As is discussed in Radford et al. (2021) and Huo et al. (2021), CLIP performs badly under this circumstance. Huo et al. (2021) claims to have tried to avoid assuming this strong correlation in BriVL. As can be seen, it turns out to be effective to some extent when compared with CLIP. However, some texts in this task have very little or even no explicit semantic correlation with corresponding images. We believe that this bias probably leads to poor encoding of either modality and causes the bad performance.

5.3. Cross-domain Experiments (RQ3)

In this part, we briefly evaluate the methods mentioned in Section 5.1 on cross-domain hyperbole detection. As is shown in Section 3.2, the posts are of five different keywords and their contents are generally disjoint. We use posts from three keywords for training, one for validating and one for testing. For each method, we repeat this procedure five times to test on all five keywords and report average performance. The results are shown in Table 7. Since pre-trained multimodal models perform badly, we do not discuss them in this part.

As can be seen, all methods perform worse than in Section 5.1, which meets our expectation because cross-domain hyperbole detection is certainly a tougher task. The results are not that terrible, which shows the generalization ability of the methods pro-

model	F1-score	Accuracy
BERT	0.652 \pm 0.162	0.667 \pm 0.045
ResNet	0.513 \pm 0.079	0.500 \pm 0.020
Concat	0.700 \pm 0.168	0.711 \pm 0.028
Gate	0.683 \pm 0.161	0.704 \pm 0.037
Attn-Gate	0.700 \pm 0.167	0.714 \pm 0.043

Table 7: Results of cross-domain experiments. The subscripts show the standard deviation of each experiment.



(a) 骑上我心爱的小自行车，它不会堵车，但会被罚款
(Riding on my lovely little bike. It won't be stuck in a traffic jam, but I may be penalized.)



(b) 让我晚回家半个小时的堵车<emoji>
(The traffic jam that makes me arrive half an hour later than usual.)

Figure 7: Two examples of test cases (Both of them are not hyperbole).

posed above. Introducing image modality still apparently increases the performance as can be seen in *concat*. However, we are surprised to find that gating and attention mechanism do not bring higher performance. We take two examples (Figure 7) from the keyword 堵车(traffic jam)(traffic jam) to illustrate the reason.

As can be seen in Figure 7a, the sentence contains an expression *won't be stuck in a traffic jam* which is seemingly too absolute here. However, in this case, it is probably a fact that bikes will not be stuck in traffic jams, so it is not hyperbole. When trained on other keywords, the models will not have such knowledge (“bikes don’t get stuck in traffic jams”), so they naturally fail on this case. This explains why all methods perform worse on this task.

Similarly, the screenshot of a navigating software only appears in posts with the certain keyword 堵车(traffic jam), so the models have no idea of what it is. Therefore, models requiring deep interaction between text modality and image modality may be confused more seriously. We believe that this systematic inconsistency is the reason why deep fusion methods perform much worse on this task. And again, common sense knowledge is still helpful in easing this problem.

6. Conclusion and Future Work

In this paper, we first introduce a new task of multimodal hyperbole detection and build a dataset for our study. We point out that different modalities can express hyperbole in mainly three ways. We conduct several experiments and prove that image modality is useful in hyperbole detection. We point out the importance of deep fusion of two modalities. Pre-trained multimodal models are proved ineffective on this task. Cross-domain results show that the methods still work relatively fine while the inconsistency between posts with different keywords is certainly the bottleneck of this task. For future work, we plan to make use of common sense knowledge in our model to achieve a better result.

7. Ethics Statement

There is not any discrimination in this work. All the examples shown in this paper are just for academic use and do not represent our attitude. We respect different human races and different aesthetic judgments. However, the posts contained in our dataset are automatically crawled from the Internet, so there may be potential controversial statements and dirty words. We sincerely apologize to those who may feel offended.

Acknowledgement

This work was supported by Beijing Science and Technology Program (Z231100007423011), National Key R&D Program of China (2021YFF0901502), National Science Foundation of China (No. 62161160339) and Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology). We appreciate the anonymous reviewers for their helpful comments. Xiaojun Wan is the corresponding author.

8. Bibliographical References

Rhys Biddle, Maciek Rybinski, Qian Li, Cecile Paris, and Guandong Xu. 2021. [Harnessing privileged information for hyperbole detection](#). In *Proceedings of the The 19th Annual Workshop of the Australasian Language Technology Association*,

pages 58–67, Online. Australasian Language Technology Association.

Yitao Cai, Huiyu Cai, and Xiaojun Wan. 2019. [Multimodal sarcasm detection in twitter with hierarchical fusion model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, page 2506–2515. Association for Computational Linguistics.

Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. 2019. Towards multimodal sarcasm detection (an _obviously_ perfect paper). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Florence, Italy. Association for Computational Linguistics.

Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. [Revisiting pre-trained models for Chinese natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 657–668, Online. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Human Language Technology: Conference of the North American Chapter of the Association of Computational Linguistics*, page 4171–4186.

Yongping Du, Yang Liu, Zhi Peng, and Xingnan Jin. 2022. Gated attention fusion network for multimodal sentiment classification. *Knowledge-Based Systems*, 240.

Gaëlle Ferré. 2014. [Multimodal hyperbole](#). *Multimodal Communication*, 3(1):25–50.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer.

Yuqi Huo, Manli Zhang, Guangzhen Liu, Haoyu Lu, Yizhao Gao, Guoxing Yang, Jingyuan Wen, Heng Zhang, Baogui Xu, Weihao Zheng, Zongzheng Xi, Yueqian Yang, Anwen Hu, Jinming Zhao, Ruichen Li, Yida Zhao, Liang Zhang, Yuqing Song, Xin Hong, Wanqing Cui, Dan Yang Hou, Yingyan Li, Junyi Li, Peiyu Liu, Zheng

- Gong, Chuhao Jin, Yuchong Sun, Shizhe Chen, Zhiwu Lu, Zhicheng Dou, Qin Jin, Yanyan Lan, Wayne Xin Zhao, Ruihua Song, and Ji-Rong Wen. 2021. [Wenlan: Bridging vision and language by large-scale multi-modal pre-training](#). *CoRR*, abs/2103.06561.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li Kong, Chuanyi Li, Jidong Ge, Bin Luo, and Vincent Ng. 2020. Identifying exaggerated language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, page 7024–7034. Association for Computational Linguistics.
- Roger J. Kreuz, Richard M. Roberts, Brenda K. Johnson, and Eugenie L. Bertus. 1996. Figurative language occurrence and co-occurrence in contemporary literature. In *Empirical Approaches to Literature and Aesthetics*, pages 83–87, Norwood, NJ. Ablex Publishing Corporation.
- 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.
- Hui Liu, Wenya Wang, and Haoliang Li. 2022. Towards multi-modal sarcasm detection via hierarchical congruity modeling with knowledge enhancement. *arXiv preprint arXiv:2210.03501*.
- Laura Cano Mora. 2009. All or nothing: A semantic analysis of hyperbole. *Journal of Language and Applied Language*, 4:25–35.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Enrica Troiano, Carlo Strapparava, Gözde Özbal, and Serra Sinem Tekiroglu. 2018. A computational exploration of exaggeration. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, page 3296–3304. Association for Computational Linguistics.
- Zilong Wang, Zhaohong Wan, and Xiaojun Wan. 2020. Transmodality: An end2end fusion method with transformer for multimodal sentiment analysis. In *Proceedings of The Web Conference 2020*, pages 2514–2520.
- Yang Wu, Yanyan Zhao, Xin Lu, Bing Qin, Yin Wu, Jian Sheng, and Jinlong Li. 2021. [Modeling incongruity between modalities for multimodal sarcasm detection](#). *IEEE MultiMedia*, 28(2):86–95.
- 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.
- Xiaojun Xue, Chunxia Zhang, Zhendong Niu, and Xindong Wu. 2022. Multi-level attention map network for multimodal sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering*.
- An Yang, Junshu Pan, Junyang Lin, Rui Men, Yichang Zhang, Jingren Zhou, and Chang Zhou. 2022. Chinese clip: Contrastive vision-language pretraining in chinese. *arXiv preprint arXiv:2211.01335*.
- Xiaocui Yang, Shi Feng, Yifei Zhang, and Daling Wang. 2021. Multimodal sentiment detection based on multi-channel graph neural networks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 328–339. Association for Computational Linguistics.
- Yiming Yang and Jan O Pedersen. 1997. A comparative study on feature selection in text categorization. In *Icml*, volume 97, page 35. Citeseer.

A. Smoothed Chi-Score

Consider a certain word d , the chi-score of the word $K(d)$ is

$$K(d) = \frac{(A + B + C + D)(AD - BC)}{(A + C)(B + D)(A + B)(C + D)}$$

, where A is the number of hyperbolic posts that contain d , B is the number of non-hyperbolic posts that contain d , C is the number of hyperbolic posts that do not contain d and D is the number of non-hyperbolic posts that do not contain d . To smooth it, we use

$$\tilde{A} = A + 1, \tilde{B} = B + 1, \tilde{C} = C + 1, \tilde{D} = D + 1$$

. So we have

$$\tilde{K}(d) = \frac{(\tilde{A} + \tilde{B} + \tilde{C} + \tilde{D})(\tilde{A}\tilde{D} - \tilde{B}\tilde{C})}{(\tilde{A} + \tilde{C})(\tilde{B} + \tilde{D})(\tilde{A} + \tilde{B})(\tilde{C} + \tilde{D})}$$