

# Improvements and Extensions on Metaphor Detection

Weicheng Ma<sup>1</sup>, Ruibo Liu<sup>2</sup>, Lili Wang<sup>3</sup>, and Soroush Vosoughi<sup>4</sup>

Minds, Machines, and Society Group  
Department of Computer Science, Dartmouth College  
<sup>1,2,3</sup>{first.last.gr}@dartmouth.edu  
<sup>4</sup>soroush.vosoughi@dartmouth.edu

## Abstract

Metaphors are ubiquitous in human language. The metaphor detection task (MD) aims at detecting and interpreting metaphors from written language, which is crucial in natural language understanding (NLU) research. In this paper, we introduce a pre-trained Transformer-based model into MD. Our model outperforms the previous state-of-the-art models by large margins in our evaluations, with relative improvements on the F-1 score from 5.33% to 28.39%. Second, we extend MD to a classification task about the metaphoricality of an entire piece of text to make MD applicable in more general NLU scenes. Finally, we clean up the improper or outdated annotations in one of the MD benchmark datasets and re-benchmark it with our Transformer-based model. This approach could be applied to other existing MD datasets as well, since the metaphoricality annotations in these benchmark datasets may be outdated. Future research efforts are also necessary to build an up-to-date and well-annotated dataset consisting of longer and more complex texts.

## 1 Introduction

*Today we are drowning in a sea of social media posts.* Metaphors serve as strong modifiers to the intentions and meanings of written texts. In the header sentence, the metaphorical use of the word “drown” in the sentence well expresses the worries of the speaker towards the large number of messages in social media, compared to the narrative version of the sentence, e.g. “There are a lot of messages on social media”. As defined by Lakoff and Johnson (1980), metaphors involve words used outside their familiar domains. For example, the word “sea” in the leading sentence literally means a large body of water, but it is used metaphorically as a modifier to the phrase “social media posts” to emphasize the abundance of

messages in social media. Similarly, people can “drown” in water, but not in messages. As shown in this example, metaphors are expressed by the context but not the aspect words themselves, and there are no limits to the number of the metaphorical parts of speech.

Metaphor detection (MD) serves as a strong component in the natural language understanding (NLU) pipeline, since NLU models cannot correctly process the meaning of written text without understanding the metaphors in the content. MD serves to aid the NLU models by figuring out the metaphorical parts of speech in each sentence. However, this is a difficult task since metaphors are carried out by long spans of text, not by the appearance of single words or phrases. Existing algorithms and neural models are not able to encode long contexts without losing critical information related to metaphors. Moreover, the lack of labeled data and the difficulties in labeling metaphorical texts are obstacles to MD research as well. Due to these issues, the research on MD is still in an early stage and has not seen the improvements observed in other NLP tasks in recent years.

To reduce the annotation difficulties, researchers have been simplifying MD to a classification problem on the metaphoricality of one word or a word pair in each sentence. Existing MD benchmark datasets are almost all labeled in this manner. While the VUA dataset (Steen, 2010) extends MD into a sequential labeling problem, it still limits the metaphorical parts-of-speech to be one per sentence. This setting alleviates the pressure of early MD models which are based on handcrafted features (Strzalkowski et al., 2013; Hovy et al., 2013; Tsvetkov et al., 2013; Gedigian et al., 2006; Beigman Klebanov et al., 2016; Bracewell et al., 2014). Nonetheless, the limitation overly simplifies MD and makes existing MD

Sentence	Her husband often <b>abuses</b> alcohol.
Explanation	To use excessively
Example	Abuse alcohol

Table 1: One example sentence from the MOH dataset that is wrongly labeled as metaphorical. The explanation of the word in bold and the example come from the Merriam-Webster dictionary.

models inapplicable in NLU pipelines. Since [Rei et al. \(2017\)](#) first introduced deep learning to MD, recent models based on deep neural networks are already approaching the performance ceilings for the simplified version of MD. Given the growing power of deep neural networks, it is time to re-define the task beyond the simplistic settings.

To verify our hypothesis, we fine-tune and evaluate a pre-trained BERT ([Devlin et al., 2019](#)) model on all the MD benchmark datasets. Our model outperforms the previous state-of-the-art models with large margins, as expected. The evaluation results almost all exceed 90% in F-1 scores, suggesting that the existing MD settings and datasets are too easy for deep Transformer networks to solve. We also extend MD to a classification task on the sentence level by removing the labels about the candidate metaphorical words. While the results slightly drop on two MD datasets (0.32% and 3.44% in F-1 scores), they are still high, especially on trivial sentences. We believe it is time to expand MD to include sentence-level metaphoricity labeling and to be evaluated on longer, more complex texts.

In the evaluations, we uncover flaws in the MD benchmark datasets by analyzing the prediction errors our model makes. One example of the annotation errors is displayed in Table 1. While the word “abuse” in this context literally means “to use excessively”, it is annotated as metaphorical in the MOH dataset. The problematic annotations might result from recent updates to the dictionaries or changes in people’s habits in using English. This situation makes it difficult to label the benchmark datasets on the sentence level with the existing word-level annotations. To validate our concerns about the quality of the annotations, we clean up one of the MD benchmark datasets and have the new annotations checked by two native English speakers. We also benchmark the re-annotated dataset with our model. The same strategy can and should be applied to other MD datasets to keep the annotations up to date. We provide more details

regarding the data analysis and re-annotation process in Section 7.

The contributions of this paper are three-fold. First, we report new state-of-the-art performances on three MD benchmark datasets to display the power of pre-trained deep Transformer networks on MD. Second, we identify and clean up the annotation errors in one of the MD benchmark datasets through manual analysis and validation, which will be made publicly available. Third, we believe that the current settings of MD are overly simplistic for deep neural network models to solve, based on the evaluation performances of our model. Thus, we extend MD to a sentence-level classification task and provide benchmark results on the three MD datasets. Our future research efforts will involve the construction of an MD dataset with sentence-level annotations and longer and more complex texts.

## 2 The Metaphor Detection Task

Following [Gao et al. \(2018\)](#), we apply both the sequential labeling and word-level classification settings of MD in the experiments. Also, we generalize the classification setting of MD to the sentence level, disregarding the aspect labels. We describe the three settings of MD as follows. For clarity, we use  $s = \{w_1, w_2, \dots, w_k\}$  to denote a sentence with  $k$  words.

**Sequential labeling:** Given a sentence  $s$ , predict one label  $l_i$  for each word  $w_i$  indicating whether  $w_i$  is metaphorical in the context.

**Word-level classification:** Given a sentence  $s$  and an aspect word  $w_i \in s$  (usually verbs, with exceptions), predict the metaphoricity label  $l_i$  associated with the aspect word.

**Sentence-level classification:** Given a sentence  $s$ , predict whether  $s$  is metaphorical.

The first two settings of MD have been extensively studied in previous research. Since metaphors are expressed by the linguistic expressions, attributing the metaphoricity of a sentence to an aspect word overly simplifies MD. but annotating an MD dataset with complex sentences under the sequential labeling setting is too difficult and costly. We provide the sentence-level classification formulation of MD for higher annotation quality while avoiding annotating an MD dataset on the token level.

Dataset	Sentence	Label
MOH	He <b>absorbed</b> the knowledge or beliefs of his tribe.	Metaphorical
TroFi	To expect banks to <b>absorb</b> a cost without a commensurate charge defies logic ./.	Non-Literal
LCC	Thank Lyndon Johnson, his Great Society, and the <b>War</b> on <i>Poverty</i> .	3

Table 2: One example record in each of the three MD benchmark datasets. The bold words are the aspect words. In the LCC dataset, the target word (in italic) of the aspect word is also provided. The label sets are {Literal, Metaphorical} in the MOH dataset, {Literal, Non-Literal} in TroFi and {0, 1, 2, 3} in LCC.

### 3 Datasets

We base our evaluations and discussions on three MD benchmark datasets, namely MOH (Mohammad et al., 2016), TroFi (Birke and Sarkar, 2006, 2007), and LCC (Mohler et al., 2016).

The MOH dataset contains sentences from WordNet (Miller, 1995, 1998) examples and the other two corpora are collected from news articles. The average number of words in the MOH dataset (7.40) is much lower than the TroFi (29.65) and LCC (28.66) datasets. This makes the MOH dataset the simplest among the three benchmark datasets. All three datasets provide one aspect word and a metaphoricity label for each sentence. The label is associated with the aspect word. The LCC dataset additionally provides the annotation about the target word of the aspect word in each sentence. Different from the other two datasets, the LCC dataset annotates the metaphoricity scores of the aspect words in the set {-1, 0, 1, 2, 3}. In the experiments, we get rid of the -1 labels in the LCC dataset since it denotes uncertain annotations. We display one sample sentence from each dataset in Table 2.

The MOH dataset is constructed with 1640 sentences, 410 out of which are annotated as metaphorical. The TroFi dataset is made of 1592 literal sentences and 2145 non-literal ones. In the LCC dataset, 493 sentences are labeled as completely literal (0) while 1242, 1251 and 1838 sentences are annotated with metaphoricity scores of 1, 2, and 3, respectively. We perform 10-fold cross-validation on all the three benchmark datasets under the word-level classification, sentence-level classification and sequential labeling settings in the experiments for fairness.

Though there exist other benchmark MD datasets as well, we choose to use the above three datasets intentionally. The VUA dataset provides annotations for the sequential labeling setting of MD. However, it is not publicly avail-

able now so we cannot obtain the data. The TSV dataset (Tsvetkov et al., 2014) is also widely used, but its training set contains only a list of adjective-noun pairs without the context. Despite the important role the aspect words play in MD, the lack of context makes it improper to train or fine-tune deep Transformer-based models on the TSV dataset. Clues for the sentence-level metaphoricity prediction cannot be learned in the training process either. Thus we do not take these two datasets into our evaluation.

### 4 Related Work

Since MD is originally defined as a classification task, most early researchers solve it with logistic regression or SVM (Support Vector Machine) classifiers. To use the information in the context, researchers concern much about the interrelations between the aspect words and the words closely related to them. Thus POS (Part of Speech) tags and dependency paths are frequently used in MD research. Shutova and Sun (2013) and Shutova et al. (2010) group the grammatical relations between each pair of aspect word and its target word into clusters, and they use rules to find out metaphorical combinations. Topical information is also a crucial clue to the domain information of a sentence so it is widely used in MD. Jang et al. (2016) represent the domain distribution of a sentence with sentence LDA. They then base their metaphoricity predictions on the similarities, differences and transition patterns between adjacent sentence pairs.

It is interesting, though, that some words are regularly used metaphorically. The intrinsic characteristics of these words are often taken into account when solving MD. Strzalkowski et al. (2013) assume that highly imaginable words are promising metaphorical words. They lookup the imaginability scores of the aspect words in the MRCPD lexicon (Coltheart, 1981; Wilson,

1988) and label the words with high imaginability scores as metaphorical. Similarly, [Bracewell et al. \(2014\)](#) also consider imaginability in predicting the metaphoricity of words. [Tsvetkov et al. \(2013\)](#) and [Turney et al. \(2011\)](#) instead use abstractness of the aspect words or the entire sentences as features in detecting metaphors. Other word-based features include the WordNet features (e.g. synonyms and semantic categories) ([Strzalkowski et al., 2013](#); [Tsvetkov et al., 2013](#)), the VerbNet features (e.g. thematic roles) ([Beigman Klebanov et al., 2016](#)), the domains of the candidate-words’ arguments ([Gedigian et al., 2006](#)), and the named entity information ([Tsvetkov et al., 2013](#)). [Jang et al. \(2016\)](#) claim that metaphors reveal the emotional or cognitive features of the author, so they use the occurrence of words in the LIWC lexicon ([Tausczik and Pennebaker, 2010](#)) to model the sentences in their research.

Some researchers do not agree with the word-level classification setting of MD. Instead, [Hovy et al. \(2013\)](#) claim that every word in a sentence can be metaphorical or literal, making it unrealistic to list all the possible aspect words. They introduce the sequential labeling setting of MD and apply CRF (Conditional Random Field) to solve it. Researchers are actively studying MD as a sequential labeling task, but a well-annotated dataset under this setting is difficult to obtain at least for now.

Not until the year of 2016 did natural language processing (NLP) researchers start to use neural networks in MD. With the power of neural networks, more and more researchers begin to examine the use of longer-term context information in MD. Do [Do Dinh and Gurevych \(2016\)](#) encode the aspect words with an MLP (Multilayer Perceptron) taking the vectorized word embeddings, POS features and positional features as inputs. They predict the metaphoricity of each aspect word by feeding their encodings into a logistic regression classifier. [Bulat et al. \(2017\)](#) similarly use pre-trained word embeddings to represent the aspect words, and they use SVD (Singular Value Decomposition) to gain sentence representation for classification. [Shutova et al. \(2016\)](#) assume that the metaphoricity of a two-word phrase can be modeled with the cosine similarity between the aspect word embedding and the phrase embedding. To represent the aspect word and phrase, they slide a window of fixed size on the context and use

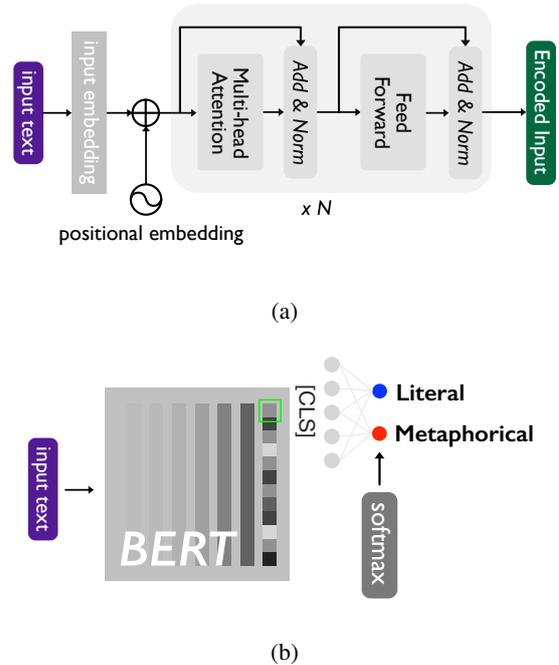


Figure 1: The architecture of Transformer networks (a) and our model (b).  $N$  denotes the number of self-attention layers in a Transformer model.

the information of all the words appearing in the window to encode the central word or the entire phrase. They also introduce visual embeddings of words into MD which, according to their experimental results, help improve the results of MD on two benchmark datasets. [Rei et al. \(2017\)](#) extend the idea of [Shutova et al. \(2016\)](#) by calculating a gated cosine similarity score between the two words’ embeddings in each phrase with neural networks. The research by [Gao et al. \(2018\)](#) consider the entire sentence as useful context information and use BiLSTM with the attention mechanism to extract the features from the sentence automatically. Most recently, [Dankers et al. \(2019\)](#) combine BERT with BiLSTM to jointly solve MD and the Emotion Regression task. Their model yields good results on MD, but it does not fully exploit the encoding ability of BERT. To go one step further, we design a BERT-based model and evaluate it on three standard evaluation datasets on MD in this paper.

## 5 Model Architecture

The Transformer networks have been overtaking the state-of-the-arts in the NLP field since their emergence. However, there have been very few works that have studied the usage of the Transformer networks in MD. To the best of our knowl-

Model	MOH			TroFi			LCC		
	WCLS	SCLS	SL	WCLS	SCLS	SL	WCLS	SCLS	SL
Dankers et al. (2019)	-	-	-	-	-	-	76.90	-	-
Gao et al. (2018)	79.10	-	75.60	72.00	-	71.10	-	-	-
Shutova et al. (2016)	75.00	-	-	-	-	-	-	-	-
Rei et al. (2017)	74.20	-	-	-	-	-	-	-	-
BERT	<b>85.52</b>	<b>86.32</b>	<b>89.18</b>	<b>92.44</b>	<b>92.12</b>	<b>94.45</b>	<b>81.00</b>	<b>77.56</b>	<b>91.48</b>

Table 3: Experimental results on the MOH, TroFi and LCC datasets with the word-level classification (WCLS), sentence-level classification (SCLS) and sequential labeling (SL) settings. All results are in terms of F-1 scores. BERT refers to our model.

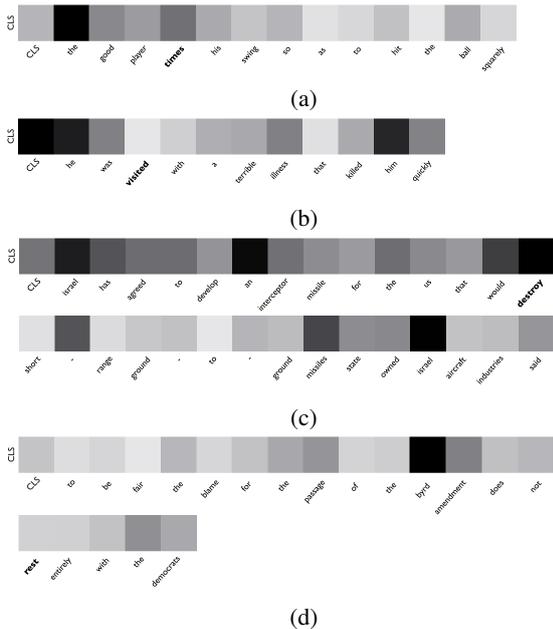


Figure 2: Attention heatmaps generated by our sentence-level classification model on the MOH and TroFi datasets. The words in bold are the aspect words.

edge, Dankers et al. (2019) made the first and only attempt in applying BERT (Devlin et al., 2019), one of the most prevalent pre-trained Transformer-based models, on MD. They build an MLP or additional attention layers on top of BERT to make metaphoricity predictions. In our point of view, however, combining BERT with complex neural network architectures is a waste of its strength. The additional layers co-trained with BERT are only exposed to the task-specific dataset which is much smaller than the BERT training data. This makes it difficult to adapt BERT to the classification layers. It is good enough to simply use a linear layer to resize the BERT output to the prediction space. We specify the neural network architecture underlying BERT in Figure 1a and show in Figure 1b the simple model architecture with

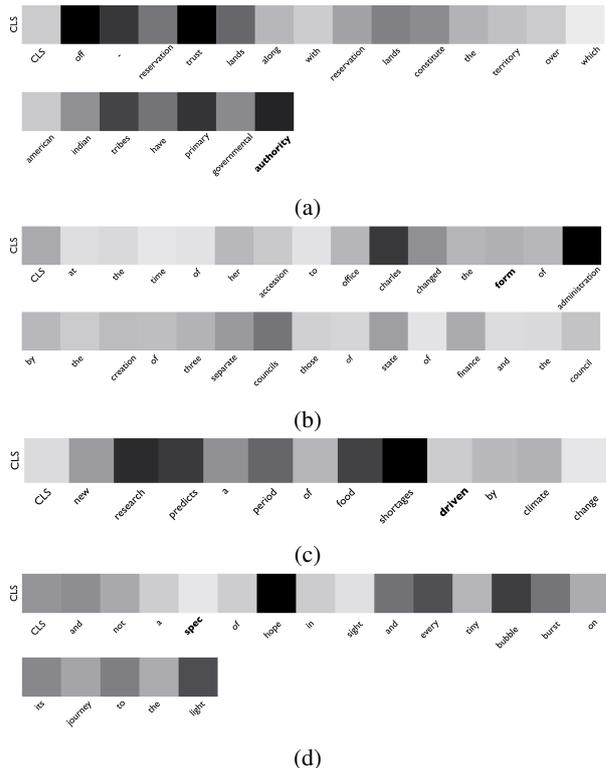


Figure 3: Attention heatmaps generated by our sentence-level classification model on the LCC dataset.

which we are able to achieve the state of the art on three MD benchmark datasets. Our experiments are based on the PyTorch implementation of the Transformer networks by Huggingface (Wolf et al., 2019).

## 6 Experiments

As is mentioned in previous sections, we fine-tune and evaluate BERT models for classification and for sequential labeling on the Trofi, MOH and LCC datasets with 10-fold cross validation. In the experiments, we use the pre-trained **bert-base-cased** model released by Google. The model architecture is a 12-layer Transformer model with 12

Dataset	ID	Sentence	Label	Pred
MOH	1	The house <b>looks</b> north.	metaphorical	literal
	2	The huge waves <b>swallowed</b> the small boat and it sank shortly thereafter.	metaphorical	literal
	3	You must <b>adhere</b> to the rules.	metaphorical	literal
	4	They <b>adhere</b> to their plan.	literal	literal
TroFi	5	At 9 p.m . , a doctor <b>examines</b> her and orders tests ./.	non-literal	literal
	6	The study , which <b>examined</b> 50 people who were wearing lap belts during auto accidents , concluded that 32 would have “ fared substantially better if they had been wearing a lap-shoulder belt . ”/”	non-literal	literal
	7	In order to focus federal resources on the SSC , its backers decided that Isabelle had to <b>die</b> ./.	non-literal	literal
LCC	8	From this calculation it is obvious that with any <b>form</b> of taxation per head the State is baling out the last coppers of the poor taxpayers in order to settle accounts with wealthy foreigners, from whom it has borrowed money instead of collecting these coppers for its own needs without the additional interest.	2	3
	9	The organism that causes gonorrhoea (gonococcus) is an <b>example</b> of a bacterial invader.	2	3
	10	Background Checks - Local Background Checks Can <b>Reduce</b> Deaths.	1	0

Table 4: Example prediction errors on the MOH, TroFi and LCC datasets. The source words are in bold.

attention heads on each layer. The hidden dimension of the model is 768. We limit the sentence lengths to 128 since it fits most of the sentences in the three datasets. In both the fine-tuning and evaluation process, we set the batch size to 128. As for training epochs, we use 5 for the aspect-based classification setting, 20 for the sentence-based setting and 20 for the sequential labeling formulation. We select the training epochs through manually monitoring the training process to avoid overfitting.

Our evaluation is performed under the three MD settings respectively. For the word-based classification setting, we mask out the aspect word in each sentence and concatenate the pair of sentences with and without the mask as input. In this way, we take advantage of BERT’s next sentence prediction mechanism. Since BERT infers the masked words with contextual information, it is highly probable that the masked word is used literally if the two sentences are predicted to be in the same context. In the sentence-level formulation, we directly feed into the model the original sentence without any change. The sequential labeling model takes the words and their indexes in

the sentence as input and predicts the metaphoricity label of each word. We label the aspect words with their annotated labels and regard all the other words as literally used in the evaluation.

Table 3 displays the 10-fold cross-validation results of our model and the baseline models on the three benchmark datasets. Our model outperforms the baseline models by large margins and constructs the new state of the art under all the three settings. The success of the models based on Elmo (Peters et al., 2018) and BERT demonstrates the importance of contextual information in MD. By comparing our model to that of Gao et al. (2018) which relies on Elmo embeddings, we demonstrate the outstanding encoding ability of BERT. Though both based on the BERT model, our model shows superior performance in MD than that of Dankers et al. (2019). This supports our assumption that overly complex classifiers built on top of BERT negatively affect the fine-tuning process.

The results show that in most cases, our model performs the best in the word-based classification setting. The more complex the sentences in the datasets are (LCC > TroFi > MOH), the more dif-

difficult the sentence-based classification setting of MD is than the word-based classification setting. This agrees with our expectation since there can be multiple metaphorical words in a sentence that influence the prediction of our model. Our model performs surprisingly well on the TroFi dataset, even better than on the MOH dataset. This might be due to the difficulty of training deep neural models on the overly simple sentences in the MOH dataset. Our model shows great potential under the sequential labeling setting as well. On all the three datasets, our model achieves F-1 scores close to or even above 90%. We are highly impressed by the power of the BERT model and we feel that the existing MD benchmark datasets are becoming too easy for deep Transformer-based models to solve. So it is time to construct new corpora containing longer and more complex text with multiple metaphorical components in each piece of text. By extending the MD research to more complex realistic scenes, the MD models can better aid the NLU research and benefit the NLP community.

## 7 Analysis and Discussions

We manually inspect the predictions our model makes to analyze the causes of the prediction errors. Table 4 displays typical prediction errors in our evaluation. The major problem with the MOH dataset is the unbalanced labels for each aspect word. Grouping by the aspect words, 194 out of the 438 word groups in the MOH dataset contain no metaphorical annotations and 11 groups have no literal annotations. The labels in the rest of the word groups are not balanced either. Models trained on unbalanced training data are likely to associate the label predictions with the appearance of the aspect words. Sentence 1 in Table 4 is the only metaphorical record with the aspect word “look” in the MOH dataset, for example. The model might have learned to classify all the sentences with the verb “look” into the literal class, generating this error case. On the other hand, most sentences in the MOH dataset are simple and the aspect words are often the only verbs. This increases the difficulty of our model to generalize the learned knowledge into predictions on longer and more complex sentences in the validation dataset. Sentence 2 features the metaphorical word “swallow” but our model is disturbed by the literal word “sink” and makes the wrong prediction. As all the sentences with “swallow” in

the MOH dataset are annotated as metaphorical, this prediction error proves that our model learns to classify not from the single aspect words, but a global view of the sentences. Some annotations in the MOH dataset are difficult for us to understand. For instance, “adhere to the rules” in Sentence 3 is labeled as metaphorical while “adhere to the plan” in Sentence 4 is literal. This leads to our hypothesis that the annotations may be wrong or outdated. With this idea in mind, we re-annotated the MOH dataset. In the resulted dataset, 402 out of the 1639 annotations (24.53%) are different from the original labels. To alleviate the problem caused by the subjectivity in the metaphoricity annotations, we sampled 100 from the records where our annotations do not agree with the original ones and had it validated with three native speakers. The agreement rate of the three independent annotators on the new annotations is 66%. This proves that our annotations are better in quality than the original labels. We use majority vote to re-label the MOH dataset and benchmark the revised dataset with our BERT-based model. The 10-fold cross-validation results are 94.21%, 94.21%, and 98.22% under the word-level classification, sentence-level classification and sequential labeling settings, respectively.

Our model performs much better on the TroFi dataset than on the MOH dataset, benefited from the abundant instances in each word group and the relatively balanced labels. However, we do not fully agree with the annotations either. The label for the word “examine” in Sentence 5, for example, is metaphorical, though the usage of “examine” in this sentence well aligns with its literal meaning “test or examine for the presence of disease or infection”. Similarly, the “examine” in Sentence 6 is used in its literal meaning “to question or examine thoroughly and closely”, but it is labeled as metaphorical. Since the TroFi dataset is collected from news articles, abbreviations sometimes cause trouble in the evaluation as well. The name “Isabelle” in Sentence 7 can well denote a person without preliminary knowledge about SSC (Superconducting Supercollider) in the context. It is then understandable why our model predicts the sentence as using the verb “die” literally. In the future, we suggest adding the surrounding sentences in the context into the dataset to make MD better defined and more appropriate for training deep neural network models.

Different from the MOH and TroFi datasets, the LCC dataset does not limit the source words to verbs. Another difference is that the labels in the LCC dataset are metaphoricity scores. This makes the LCC dataset more difficult to solve. Our model predicts 3 while the label is 2 for the word “form” in Sentence 8, for example. Possibly our model detects the metaphorical use of “copper” in the same sentence and decides to assign a higher metaphoricity score to the entire sentence. The prediction error of our model in Sentence 9 is in a similar case. Our model predicts a high score due to the synergy of the metaphorical words “example” and “invader”. The annotations in the LCC dataset are sometimes controversial as well. The word “reduce” in Sentence 10 perfectly matches the literal meaning “to cut down on”, but is annotated as 1 (weakly metaphorical) in LCC, for example.

On the other hand, since higher attention weights are put on the evidence for classification in Transformer-based models, we examine the attention maps on the last self-attention layer generated by our model under the sentence-level classification setting to interpret the performance of our model. Under the sentence-level classification setting, the predictions are made from the hidden states of the CLS token. So we evaluate the attention scores of the CLS token on all the other words in each sentence. The rule well applies to the word-level classification and sequential labeling settings, only with different tokens on which to base the predictions. To avoid duplication, we only display the attention heatmaps generated under the sentence-level setting in this paper. Figure 2 displays the attention heatmaps on MOH and TroFi examples to reflect the influence of metaphorical polarity on the attention scores and Figure 3 contains the heatmaps on LCC examples to show the effect of metaphorical intensities. In Figure 2a, the subject “the good player”, the verb “times” and the object “his swing” are all heavily attended, indicating the literal usage of the word “time” paired with the words “player” and “swing”. Quite on the contrary, the verb “visited” in Figure 2b is very lightly attended compared to “he” and “illness” in the same sentence, which is a signal of the metaphorical use of the word “visited” in its context. The same pattern applies to the examples in TroFi (Figure 2d and 2c) and LCC (Figure 3a, 3b, 3c and 3d) that the more heavily

our model attends on an aspect word, the lower chance it is used metaphorically in the context. It is worth noting that when the sentences grow longer, the amount of potential aspect words also increases. The use of these aspect words can be literal or metaphorical at the same time, which benefits classifying the metaphoricity of the sentence as a whole. In Figure 2a, for instance, the verb “hit” is used literally with the noun “ball” as well. But there are also cases where the multiple aspect words in one sentence hold different metaphoricities, e.g. the words “swallow” and “sink” in Sentence 2 of Table 4. These examples contribute to many prediction errors made by our sentence-level classification model but are generally not a problem for the aspect-based classification and sequential labeling models. As we stated before, examining the metaphoricity of given aspect words only simplifies MD. Given the powerful neural models in the NLP field, we do not need this type of simplification anymore. As our next step, we will keep working on labeling MD datasets at the sentence-level or on the aspect level with multiple aspect words per sentence. We will also introduce social media data to MD for richer metaphorical expressions and varied topics.

## 8 Conclusion and Future Work

Though difficult, MD has been an important task in the NLP community. In this paper, we refined the definitions of MD by defining a new task formulation. We also designed and evaluated a BERT-based model on three MD benchmark datasets. Our model largely outperformed the previous state-of-the-art methods. Through analysis of the prediction errors made by our model, we found that a large number of prediction errors can be attributed to the simplicity of the datasets and the annotation qualities. To validate this, we re-annotated the MOH dataset and manually verified the quality of our new annotations. We saw in the experiments that our model achieves very high accuracy on existing MD benchmark datasets, meaning that they are becoming overly simple for deep neural networks. Our future work will focus on collecting and annotating a new MD dataset with more complex texts. Regarding the prosperity of social media, we also plan to address the metaphor detection problem on informal text. We hope our work will attract more interest to MD and we call for future contributions to solve the problem.

## References

- Beata Beigman Klebanov, Chee Wee Leong, E. Dario Gutierrez, Ekaterina Shutova, and Michael Flor. 2016. [Semantic classifications for detection of verb metaphors](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 101–106, Berlin, Germany. Association for Computational Linguistics.
- Julia Birke and Anoop Sarkar. 2006. [A clustering approach for nearly unsupervised recognition of non-literal language](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy. Association for Computational Linguistics.
- Julia Birke and Anoop Sarkar. 2007. [Active learning for the identification of nonliteral language](#). In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 21–28, Rochester, New York. Association for Computational Linguistics.
- David B Bracewell, Marc T Tomlinson, Michael Mohler, and Bryan Rink. 2014. A tiered approach to the recognition of metaphor. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 403–414. Springer.
- Luana Bulat, Stephen Clark, and Ekaterina Shutova. 2017. [Modelling metaphor with attribute-based semantics](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 523–528, Valencia, Spain. Association for Computational Linguistics.
- Max Coltheart. 1981. The mrc psycholinguistic database. *The Quarterly Journal of Experimental Psychology Section A*, 33(4):497–505.
- Verna Dankers, Marek Rei, Martha Lewis, and Ekaterina Shutova. 2019. [Modelling the interplay of metaphor and emotion through multitask learning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2218–2229, Hong Kong, China. 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 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Erik-Lân Do Dinh and Iryna Gurevych. 2016. [Token-level metaphor detection using neural networks](#). In *Proceedings of the Fourth Workshop on Metaphor in NLP*, pages 28–33, San Diego, California. Association for Computational Linguistics.
- Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. [Neural metaphor detection in context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 607–613, Brussels, Belgium. Association for Computational Linguistics.
- Matt Gedigian, John Bryant, Sridhar Narayanan, and Branimir Cicic. 2006. [Catching metaphors](#). In *Proceedings of the Third Workshop on Scalable Natural Language Understanding*, pages 41–48, New York City, New York. Association for Computational Linguistics.
- Dirk Hovy, Shashank Srivastava, Sujay Kumar Jauhar, Mrinmaya Sachan, Kartik Goyal, Huiying Li, Whitney Sanders, and Eduard Hovy. 2013. [Identifying metaphorical word use with tree kernels](#). In *Proceedings of the First Workshop on Metaphor in NLP*, pages 52–57, Atlanta, Georgia. Association for Computational Linguistics.
- Hyeju Jang, Yohan Jo, Qinlan Shen, Michael Miller, Seungwhan Moon, and Carolyn Rosé. 2016. [Metaphor detection with topic transition, emotion and cognition in context](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 216–225, Berlin, Germany. Association for Computational Linguistics.
- G Lakoff and M Johnson. 1980. Metaphors we live by.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- George A Miller. 1998. *WordNet: An electronic lexical database*. MIT press.
- Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. [Metaphor as a medium for emotion: An empirical study](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 23–33, Berlin, Germany. Association for Computational Linguistics.
- Michael Mohler, Mary Brunson, Bryan Rink, and Marc Tomlinson. 2016. [Introducing the LCC metaphor datasets](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 4221–4227, Portorož, Slovenia. European Language Resources Association (ELRA).
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.

- Marek Rei, Luana Bulat, Douwe Kiela, and Ekaterina Shutova. 2017. [Grasping the finer point: A supervised similarity network for metaphor detection](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1537–1546, Copenhagen, Denmark. Association for Computational Linguistics.
- Ekaterina Shutova, Douwe Kiela, and Jean Maillard. 2016. [Black holes and white rabbits: Metaphor identification with visual features](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 160–170, San Diego, California. Association for Computational Linguistics.
- Ekaterina Shutova and Lin Sun. 2013. [Unsupervised metaphor identification using hierarchical graph factorization clustering](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 978–988, Atlanta, Georgia. Association for Computational Linguistics.
- Ekaterina Shutova, Lin Sun, and Anna Korhonen. 2010. [Metaphor identification using verb and noun clustering](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1002–1010, Beijing, China. Coling 2010 Organizing Committee.
- Gerard Steen. 2010. *A method for linguistic metaphor identification: From MIP to MIPVU*, volume 14. John Benjamins Publishing.
- Tomek Strzalkowski, George Aaron Broadwell, Sarah Taylor, Laurie Feldman, Samira Shaikh, Ting Liu, Boris Yamrom, Kit Cho, Umit Boz, Ignacio Cases, and Kyle Elliot. 2013. [Robust extraction of metaphor from novel data](#). In *Proceedings of the First Workshop on Metaphor in NLP*, pages 67–76, Atlanta, Georgia. Association for Computational Linguistics.
- Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54.
- Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman, Eric Nyberg, and Chris Dyer. 2014. [Metaphor detection with cross-lingual model transfer](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 248–258, Baltimore, Maryland. Association for Computational Linguistics.
- Yulia Tsvetkov, Elena Mukomel, and Anatole Gershman. 2013. [Cross-lingual metaphor detection using common semantic features](#). In *Proceedings of the First Workshop on Metaphor in NLP*, pages 45–51, Atlanta, Georgia. Association for Computational Linguistics.
- Peter Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. [Literal and metaphorical sense identification through concrete and abstract context](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 680–690, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Michael Wilson. 1988. Mrc psycholinguistic database: Machine-usable dictionary, version 2.00. *Behavior research methods, instruments, & computers*, 20(1):6–10.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.