



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/177027/>

Version: Submitted Version

Article:

Mao, R., Lin, C. and Guerin, F. (Submitted: 2021) Interpreting verbal metaphors by paraphrasing. arXiv. (Submitted)

© 2021 The Author(s). For reuse permissions, please contact the Author(s).

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Interpreting Verbal Metaphors by Paraphrasing

Rui Mao¹, Chenghua Lin², and Frank Guerin³

¹Department of Computing Science, University of Aberdeen, AB24 3UE, UK

¹r03rm16@abdn.ac.uk

²Department of Computer Science, University of Sheffield, S1 4DP, UK

²c.lin@sheffield.ac.uk

³Department of Computer Science, University of Surrey, GU2 7XH, UK

³f.guerin@surrey.ac.uk

Abstract

Metaphorical expressions are difficult linguistic phenomena, challenging diverse Natural Language Processing tasks. Previous works showed that paraphrasing a metaphor as its literal counterpart can help machines better process metaphors on downstream tasks. In this paper, we interpret metaphors with BERT and WordNet hypernyms and synonyms in an unsupervised manner, showing that our method significantly outperforms the state-of-the-art baseline. We also demonstrate that our method can help a machine translation system improve its accuracy in translating English metaphors to 8 target languages.

1 Introduction

Metaphor is defined as using one or several words to illustrate a meaning different to the basic meaning of the words (Steen et al., 2010). Due to the difficulty of inferring underlying meanings of metaphors, metaphoric expressions challenge diverse Natural Language Processing (NLP) tasks, e.g., sentiment analysis (Ghosh et al., 2015) and machine translation (MT) (Mao et al., 2018). Currently, BERT has achieved large improvements on diverse downstream tasks (Devlin et al., 2019). However, it has not been examined on metaphor interpretation. In this paper, we interpret verbal metaphors with BERT and WordNet (Fellbaum, 1998). We focus on verbal metaphors, because verbs are the most common metaphoric expressions among all PoS categories (Steen et al., 2010), forming the most popular tasks in automatic metaphor identification and interpretation (Shutova, 2015; Leong et al., 2018).

We consider metaphor interpretation as a paraphrasing task, predicting literal counterparts for metaphors. Unlike fine-tuning or feature based BERT applications, we use BERT as a language modelling (LM) method to predict a semantically

similar word to the real meaning of a metaphor, which has the highest probability of appearing in a given context as the literal paraphrase of the metaphor (see § 3 for the reason). The probability is given by BERT based LM prediction and the semantically similar word is constrained by WordNet hypernyms and synonyms. Therefore, our approach is fully unsupervised which does not require labelling a large dataset for model training.

By running an automatic evaluation on a publicly available dataset (MOH, Mohammad et al., 2016), our model achieves 8% accuracy gains on verbal metaphor paraphrasing, compared with the state-of-the-art baseline (Mao et al., 2018). On the human evaluated paraphrasing task, the average gain of our model is 13% on two benchmark datasets, i.e., MOH and VUA (Steen et al., 2010). We also examine the metaphor paraphrases on MT tasks based on Google translator¹ and 8 diverse target languages. Evaluation results show that our method can improve the accuracy of translating English metaphors by an average gain of 20.9%, outperforming the baseline by 12.9%.

The contribution of this work can be summarised as follows: (1) we introduce a simple yet effective method by using BERT as a LM model for metaphor interpretation, yielding significant improvements against the state-of-the-art baseline; (2) with our model, a MT system can significantly improve its accuracy of translating English metaphors into 8 diverse target languages.

2 Related work

Due to the lack of large annotated metaphor interpretation corpora, previous works addressed metaphor interpretation mainly by unsupervised modelling of co-occurrences of words and their contexts with different knowledge bases and lexical

¹<https://translate.google.com/>

resources (Shutova, 2010; Shutova et al., 2012; Bollegala and Shutova, 2013). These works focused on interpreting verbal metaphors from word-pairs with specific syntactic structures, e.g., verb-subject and verb-direct object, which is inconvenient for real-world applications. Another trend in metaphor interpretation focused on a specific domain, e.g., question answering about Unix (Martin, 1990), event descriptions in economics (Narayanan, 1997) and mental states descriptions (Barnden and Lee, 2002) with hand-coded knowledge and logic rules. Mao et al. (2018) extended the word-pair models, interpreting metaphors from full sentences in open domains. They modelled the co-occurrences of words and their contexts with word2vec CBOW (Mikolov et al., 2013) input and output vectors to generate appropriate paraphrases of metaphors. Compared with context-independent word2vec embeddings, BERT as a context-dependent word embedding method, has shown remarkable performance on diverse NLP tasks (Devlin et al., 2019). However, to the best of our knowledge, BERT has not been applied in metaphor interpretation.

3 Methodology

Inspired by Mao et al. (2018), we introduce a hypernym and synonym constrained missing word prediction method for unsupervised metaphor interpretation. The difference is our missing word prediction is based on pre-trained BERT (BERT-LARGE-CASED²) that has 24 Transformer layers (Vaswani et al., 2017), whereas Mao et al. (2018) used word2vec input and output vectors.

BERT is a pre-trained Language Modelling method. The training target is to predict randomly selected masked WordPieces (Wu et al., 2016) and the next sentence of a current processing sentence. Since the masked word prediction is bidirectional, combining its surrounding context information, a pre-trained BERT model on open corpora can be naturally applied as a missing word prediction model. In our task, we predict the probability of a missing word (m_i) with its context (w) by using pre-trained BERT

$$p(m_i|w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_n) \\ = \text{BERT}([\text{CLS}], w_1, \dots, [\text{MASK}]_i, \dots, w_n, [\text{SEP}]),$$

where the position of m_i is replaced with $[\text{MASK}]_i$; $[\text{CLS}]$ and $[\text{SEP}]$ are special tokens, representing

²<https://github.com/google-research/bert>

the start and end of an input sequence with a length of n . For a sentence that has multiple metaphors, we mask one metaphoric word each time.

Next, we predict an appropriate paraphrase for the metaphor. In order to connect a metaphor (the target word) with its possible literal counterparts, we introduce the semantic constraints of WordNet hypernyms and synonyms for candidate literal counterpart development. It is likely that one of the hypernyms and synonyms has a similar meaning to the underlying meaning of a metaphor, thus, they help the BERT model filter out irrelevant predictions. The hypernyms $\{h_{i,j}\}$ and synonyms $\{s_{i,k}\}$ of a metaphoric verb i and their different verb forms (f , inflections) are considered as our candidate word set ($\mathcal{C}_i = \{h_{i,j}, h_{i,j}^f, s_{i,k}, s_{i,k}^f\}$).

The best fit word (b_i) in \mathcal{C}_i is given by a candidate word with the highest probability in its context

$$b_i = \underset{m_i \in \mathcal{C}_i}{\operatorname{argmax}} p(m_i|w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_n).$$

b_i is considered as literal, because according to relevant statistics from Cameron (2003); Martin (2006); Steen et al. (2010) and Shutova (2016), literals are more common in typical corpora. Thus, a literal has higher probability appearing in a context than a metaphor with the similar meaning. Noticeably, the limitation of paraphrasing a metaphor with a single literal word is that the paraphrase may lose some nuance in the original metaphor. However, such a paraphrase method can help a machine to better translate the real meanings of metaphors into human comprehensible languages (see § 6).

4 Dataset

MOH. MOH was formed by Mohammad et al. (2016), sourcing from WordNet example sentences. We select 315 sentences containing 315 metaphoric verbs whose metaphoricity was agreed by at least 70% annotators, forming our automatic evaluation test set (MOH₃₁₅). The average sentence length is 8.8. We conduct human evaluation with 50 randomly selected sentences (MOH₅₀) from MOH₃₁₅. **VUA₅₀.** We also randomly select 50 metaphoric verbs and their associated sentences from VU Amsterdam Metaphor Corpus (Steen et al., 2010) for human evaluation, where the sentences are originated from British National Corpus (Burnard, 2000). VUA₅₀ may contain multiple metaphoric verbs in a sentence. Thus, 36 different sentences from all four different genres e.g., academic text, conversation,

news and fiction are selected. The average sentence length of VUA₅₀ is 27.7.

5 Baselines

SIM-CBOW_{I+O} (Mao et al., 2018) modelled word co-occurrences with CBOW based word2vec (Mikolov et al., 2013) input and output vectors. Their paraphrases were also constrained by WordNet hypernyms and synonyms.

BERT-LARGE-CASED-WWM (Cui et al., 2019) masked whole words during the BERT pre-training procedure, rather than WordPieces. We test this method as an alternative of BERT-LARGE-CASED for the best fit word prediction.

6 Result

The evaluation is conducted as three phases. First, we automatically evaluate model performance on MOH₃₁₅. Since MOH₃₁₅ originated from WordNet example sentences, a paraphrasing prediction is correct, if the paraphrase belongs to the WordNet sense class which the test sentence was in.

As seen under the first column in Table 1, the accuracy of BERT-LARGE-CASED on the automatic evaluation are 49%, outperforming SIM-CBOW_{I+O} by 8%. It shows that using BERT can better paraphrase a target word in a sentence, based on constrained missing word prediction. This is likely because Transformer based BERT can better model the long-term dependency and word orders in sentences, while a shallow neural network based CBOW only represents word co-occurrences within a window of contexts. In CBOW, the position of a target word was not specified in a sentence, while BERT used positional embedding for each word during its pre-training procedure. Compared with BERT-LARGE-CASED-WWM, there is a gain of 3% in BERT-LARGE-CASED. Masking whole words during pre-training does not outperform the original BERT that masked WordPieces, probably because the learned stem information from WordPiece is more useful for predicting a missing word than the whole-word-masking approach. Thus, the original BERT can better retrieve a missing word.

In the second phase, we test model performance on MOH₅₀ and VUA₅₀ with human evaluation. As seen in Table 1, the accuracy of automatic evaluation on MOH₅₀ given by SIM-CBOW_{I+O}, BERT-LARGE-CASED-WWM and BERT-LARGE-CASED is similar to their performance on MOH₃₁₅. For human evaluation, we in-

cluded 3 native English speakers with master degrees. A participant would be asked if a lemmatized paraphrase of a target word is acceptable in the given context with a questionnaire. A paraphrase is acceptable, if it captures the original meaning of a target word in the given context, where the grammatical tense of the lemmatized paraphrase is not a part of consideration. An example of the questionnaire can be viewed in Figure 1a. The final decision is agreed by at least 2 annotators.

In the human evaluation (Table 1), BERT-LARGE-CASED surpasses SIM-CBOW_{I+O} by 14%, 12% and 13% on MOH₅₀, VUA₅₀ and their combination datasets (Cohen’s $\kappa = 0.61$), respectively, yielding an average gain of 13%. The average gain of BERT-LARGE-CASED over BERT-LARGE-CASED-WWM is 4%. The accuracy of the three models on MOH₅₀ with human evaluation is higher than their automatic evaluation, because candidate words from other WordNet sense classes are possibly acceptable from practical aspects. E.g., “read” is an acceptable paraphrase of “scan” in “She scanned the newspaper headlines while waiting for the taxi” for human, although it is incorrect for automatic evaluation. BERT-LARGE-CASED performs similarly on MOH₅₀ and VUA₅₀, although the average length of VUA sentences (27.7) is longer than that of MOH (8.8).

In the third phase, for evaluating the quality of paraphrases on a downstream task, we examine the MOH₅₀ paraphrases on English MT with 8 target languages, namely *German, Russian, Greek, Italian, Dutch, Chinese, Thai* and *Japanese*. We invited 3 participants per target language, to annotate the quality of a translation of each target word in an original sentence, a SIM-CBOW_{I+O} paraphrased sentence, and a BERT-LARGE-CASED paraphrased sentence. BERT-LARGE-CASED-WWM is excluded in this phase, because many predictions of it are the same as the predictions given by BERT-LARGE-CASED by our observation (e.g., B and C in Figure 1a), whereas BERT-LARGE-CASED surpasses BERT-LARGE-CASED-WWM on the paraphrase evaluations. The translations are given by Google translator. The participants are native in the target languages, and had learnt English in school, and had each spent a minimum of one year living and studying in England as a postgraduate. Participants were provided with senses of the original metaphors in WordNet, so that they can fully understand the true meanings of the metaphors. The

<p>1. Am I supposed to swallow that story?</p> <p>For the target word “swallow” (the bold):</p> <p style="text-align: right;">Yes/No</p> <p>A. “talk” is an acceptable paraphrase; <input type="checkbox"/></p> <p>B. “believe” is an acceptable paraphrase; <input type="checkbox"/></p> <p>C. “believe” is an acceptable paraphrase. <input type="checkbox"/></p>	<p>1. Am I supposed to swallow that story?</p> <p>swallow: believe or accept without questioning or challenge. Acceptable/Unacceptable</p> <p>A. Я должен проглотить эту историю? <input type="checkbox"/></p> <p>B. Я должен рассказать эту историю? <input type="checkbox"/></p> <p>C. Я должен верить этой истории? <input type="checkbox"/></p>
(a) A is paraphrased by SIM-CBOW _{I+O} ; B is by BERT-LARGE-CASED; C is by BERT-LARGE-CASED-WWM.	(b) A is translated by original sentence; B is by SIM-CBOW _{I+O} paraphrasing; C is by BERT-LARGE-CASED paraphrasing.

Figure 1: Example questions for human evaluation. (a) Paraphrasing task. (b) Machine Translation task.

Model	Automatic Evaluation		Human Evaluation		
	MOH ₃₁₅	MOH ₅₀	MOH ₅₀	VUA ₅₀	MOH ₅₀ +VUA ₅₀
SIM-CBOW _{I+O}	0.41	0.42	0.68	0.70	0.69
BERT-LARGE-CASED-WWM	0.46	0.46	0.76	0.80	0.78
BERT-LARGE-CASED	0.49	0.50	0.82	0.82	0.82

Table 1: Metaphor interpretation evaluation on paraphrasing tasks, measured by accuracy.

	German	Russian	Greek	Italian	Dutch	Chinese	Thai	Japanese	Avg
Cohen’s kappa (κ)	0.59	0.67	0.60	0.63	0.67	0.69	0.71	0.55	-
Original	0.48	0.32	0.32	0.50	0.50	0.38	0.30	0.28	0.385
SIM-CBOW _{I+O}	0.38	0.40	0.36	0.56	0.46	0.60	0.52	0.44	0.465
BERT-LARGE-CASED	0.60	0.42	0.46	0.72	0.66	0.68	0.70	0.50	0.594

Table 2: Metaphor interpretation evaluation on Machine Translation tasks, measured by accuracy.

quality of a translation is measured by a binary choice, e.g., acceptable or unacceptable. An acceptable translation is qualified, if a translation expresses the true meaning of the target word and it matches its context (cohesion) in their mother language, where errors in context words and tenses are ignored. The final decision was agreed by at least two annotators (κ ranges from 0.55 to 0.71 in Table 2). An example can be viewed in Figure 1b.

As seen in Table 2, the average accuracy of original metaphor translations is 38.5%. By paraphrasing the metaphors with BERT-LARGE-CASED, the average accuracy is improved to 59.4%, achieving a gain of 20.9%, outperforming SIM-CBOW_{I+O} by 12.9%. The average gain of BERT-LARGE-CASED on the target Asian languages (30.7%) against the original metaphor translations are higher than that of European languages (15.0%). This is probably because the larger language and cultural differences mean that English metaphors which are directly translated are unlikely to be understandable by Asians. After paraphrasing, the incomprehensible phrases become comprehensible. E.g., “swallow” in “Am I supposed to swallow that story?”

and “tugged” in “She tugged for years to make a decent living.” are badly translated by Google across all the target languages. However, the translations based on BERT-LARGE-CASED paraphrased target words, e.g. “Am I supposed to believe that story?” and “She struggled for years to make a decent living.” are acceptable in each target language. We also noticed that there are some translation errors introduced due to incorrect paraphrasing. However, the accuracy gain achieved by paraphrasing metaphors substantially surpasses the small impact from the error, yielding a significant overall improvement (20.9% in accuracy) in the MT task.

7 Conclusion

We propose an unsupervised model for metaphor interpretation based on BERT and WordNet, yielding large gains against the baseline on metaphor paraphrasing tasks. Our evaluation demonstrates that paraphrasing English metaphors into their literal counterparts can help a MT system improve the accuracy of translating the metaphors into more comprehensible target languages. In future work, we will test our model on other downstream tasks.

References

- John A Barnden and Mark G Lee. 2002. An artificial intelligence approach to metaphor understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.
- D Bollegala and E Shutova. 2013. Metaphor interpretation using paraphrases extracted from the web. *PLoS ONE*, 8(9).
- Lou Burnard. 2000. Reference guide for the British national corpus (world edition).
- Lynne Cameron. 2003. *Metaphor in Educational Discourse*. A&C Black.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for Chinese BERT. *arXiv preprint arXiv:1906.08101*.
- 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*, pages 4171–4186.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Aniruddha Ghosh, Guofu Li, Tony Veale, Paolo Rosso, Ekaterina Shutova, John Barnden, and Antonio Reyes. 2015. Semeval-2015 task 11: Sentiment analysis of figurative language in twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 470–478.
- Chee Wee Ben Leong, Beata Beigman Klebanov, and Ekaterina Shutova. 2018. A report on the 2018 VUA metaphor detection shared task. In *Proceedings of the Workshop on Figurative Language Processing*, pages 56–66.
- Rui Mao, Chenghua Lin, and Frank Guerin. 2018. Word embedding and WordNet based metaphor identification and interpretation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1222–1231.
- James Martin. 1990. *A Computational Model of Metaphor Interpretation*. Academic Press Professional.
- James H Martin. 2006. A corpus-based analysis of context effects on metaphor comprehension. Technical Report CU-CS-738-94, Boulder: University of Colorado: Computer Science Department.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- 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.
- Srini Narayanan. 1997. Knowledge-based action representations for metaphor and aspect (KARMA). *Computer Science Division, University of California at Berkeley dissertation*.
- Ekaterina Shutova. 2010. Automatic metaphor interpretation as a paraphrasing task. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1029–1037. Association for Computational Linguistics.
- Ekaterina Shutova. 2015. Design and evaluation of metaphor processing systems. *Computational Linguistics*, 41(4):579–623.
- Ekaterina Shutova. 2016. Design and evaluation of metaphor processing systems. *Computational Linguistics*.
- Ekaterina Shutova, Tim Van de Cruys, and Anna Korhonen. 2012. Unsupervised metaphor paraphrasing using a vector space model. In *24th International Conference on Computational Linguistics*, page 1121.
- Gerard J Steen, Aletta G Dorst, J Berenike Herrmann, Anna Kaal, Tina Krennmayr, and Trijntje Pasma. 2010. *A Method for Linguistic Metaphor Identification: From MIP to MIPVU*, volume 14. John Benjamins Publishing.
- 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 *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.