

HYU at SemEval-2022 Task 2: Effective Idiomaticity Detection with Consideration at Different Levels of Contextualization

Youngju Joung*

Dept. of Industrial Security
Dept. of Applied Statistics
Chung Ang University
ojoo.yj@gmail.com

Taeuk Kim**

Dept. of Computer Science
Dept. of Artificial Intelligence
Hanyang University
kimtaeuk@hanyang.ac.kr

Abstract

We propose a unified framework that enables us to consider various aspects of contextualization at different levels to better identify the idiomaticity of multi-word expressions. Through extensive experiments, we demonstrate that our approach based on the inter- and inner-sentence context of a target MWE is effective in improving the performance of related models. We also share our experience in detail on the task of SemEval-2022 Task 2 such that future work on the same task can be benefited from this.

1 Introduction

Multi-word expressions (MWEs) are a group of linguistic components containing two or more words with outstanding collocation (Baldwin and Kim, 2010; Constant et al., 2017). MWEs are valuable in that they contribute to enriching the expressiveness of a language, allowing diverse interpretations of their meaning according to the context in which they are located. That is, the semantics of an MWE can be originated from either (i) the direct composition of the literal definitions of its constituents (i.e., compositional meaning) or (ii) its conventional usage in the language (i.e., idiomatic meaning). For instance, given an expression called *wet blanket*, its compositional meaning is ‘a piece of cloth soaked in liquid’, whereas its idiomatic meaning is ‘a person who spoils the mood’ (see Table 1).

While MWEs function as an effective means of improving the abundance of a language, they are also one of the main obstacles that complicate natural language processing (NLP), from the perspective that an NLP model should be able to precisely identify their mode. In addition, the current trend where the goal of most NLP models is chiefly focused on capturing compositionality raises the question of how properly to deal with

Category	Meaning	Example
Compositional (Non-idiomatic)	A piece of cloth soaked in liquid.	And finally, the snow falls again, this time in a thick, <i>wet blanket</i> that encapsulates everything.
Idiomatic	A person who spoils the mood.	When Marie brings him down to earth, it’s not clear if she’s being a passive-aggressive <i>wet blanket</i> or if she might have a point.

Table 1: Comparison between the compositional and idiomatic meanings of the expression *wet blanket*.

idiomatic aspects of linguistic expressions (Garcia et al., 2021a,b; Zeng and Bhat, 2021).

An intuitive solution for mitigating the aforementioned problem is an introduction of a sophisticated method designed to estimate the idiomaticity of a given expression, which enables the separate processing of the expression according to its category. In a similar vein, we propose a series of techniques for better detecting the idiomaticity of a target MWE by actively exploiting its surrounding context in addition to considering the relationship between metaphors and the notion of idiomaticity.

Participating in SemEval-2022 Task 2, we focus on classifying two-noun compounds into idiomatic and non-idiomatic. The task provides two configurations. In the zero-shot setting, a model’s performance is evaluated on the set of sequences that include MWEs never appeared in the training phase. Meanwhile, in the one-shot setting, our model is exposed to a pair of instances per each MWE during training, one of which shows the idiomatic use of the MWE while the other is an example for the non-idiomatic case.¹ We present a unified framework that can be used in both kinds, paying slightly more attention to the one-shot setting. In extensive experiments, we show that most of our considerations lead to improvement in performance. We also present discourse on the task specification to promote a fair comparison between related models.

*Work done while Youngju was an intern at HYU.

**Corresponding author.

¹For more details on the task’s specification, refer to the task description paper (Tayyar Madabushi et al., 2022).

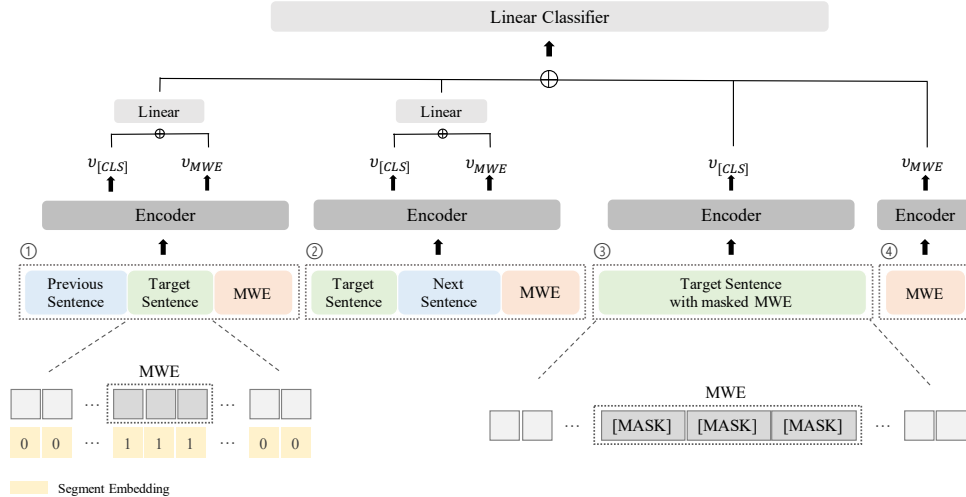


Figure 1: Proposed framework. Two features on the left (① and ②) are based on the surrounding context (Section 3.1), while the remaining two (③ and ④) are originated from considering the inner-sentence context (Section 3.2).

2 Related Work

Idiomat�city detection has been widely studied in the literature (Reddy et al., 2011; Liu and Hwa, 2019; Garcia et al., 2021a,b; Zeng and Bhat, 2021). Above all, Tayyar Madabushi et al. (2021) present a dataset that is the foundation of SemEval 2022 Task 2. This dataset consists of sentences that contain potential idiomatic MWEs with two surrounding sentences and annotations about the fine-grained set of meanings. The authors also evaluate a model’s ability to detect idiomatic usage depending on whether context and MWE are included. They report that reflecting the context in the way of simply concatenating surrounding sentences is not generally helpful, and that adding the corresponding MWE at the end of the input sequence improves performance. In the following sections, we re-examine their findings and present our own revision.

On the other hand, we investigate the viability of applying techniques for metaphor detection (Gao et al., 2018; Mao et al., 2019; Lin et al., 2021) into idiomat�city classification, inspired by the resemblance of the two tasks. A metaphor is a form of figurative expression used to implicitly compare two things seemingly unrelated on the surface at the attribute level (Baldwin and Kim, 2010). Not all metaphors have the property of idiomat�city, but some idioms rely on metaphorical composition.

In practice, Choi et al. (2021) introduce two metaphor identification theories (Metaphor Identification Procedure (MIP; Group (2007), Steen (2010)) and Selectional Preference Violation (SPV; Wilks (1975)) into their model to better capture

metaphors, which we expect also might be helpful for the procedure of identifying idiomatic expressions. The basic ideas of MIP and SPV are that a metaphor can be identified when we discover the difference between its literal and contextual meaning, and that it can also be detected when its semantics is distinguishable from that of its context. To realize the concepts, for MIP, Choi et al. (2021) employ a target word’s contextualized and isolated representations, while for SPV, they utilize the contextualized representations of the target word and the sentence including the word. We adopt some of their ideas and customize them for our purpose, i.e., modeling features for idiomat�city detection.

3 Proposed Method

As a participant of SemEval-2022 Task 2, we propose a framework powered on four features devised to facilitate the detection of idiomatic expressions. These features are computed by the same foundation model,² but distinguished from each other by what is inserted into the model as input to compute the features. A simple linear classifier is introduced on top of the concatenation of the four features to finally gauge the idiomat�city of an MWE in an input sequence. Figure 1 presents the overall picture of our method.

3.1 Features Based on Surrounding Context

We first focus on the fact reported by Tayyar Madabushi et al. (2021) that the *surrounding context*

²In this work, a ‘foundation model’ refers to a Transformer encoder pre-trained on large corpora, e.g., BERT and XLM-R.

(*a.k.a. inter-sentence context*), which we define as sentences located right before and after a target sentence, is uninformative in predicting idiomaticity. We hypothesize that this disappointing outcome is partly due to the way such context was exploited.

To be specific, given a sentence containing an MWE and its previous and following sentences, [Tayyar Madabushi et al. \(2021\)](#) propose simply putting all the three together in order. Despite its simplicity, this approach has an explicit drawback that a model should automatically learn how to distinguish the target sentence from its surrounding context. Moreover, when combined with context without caution, the input sequence becomes much (approximately $3\times$) longer than its original form, which might cause a negative effect on performance by merely intensifying the complexity of the problem rather than providing additional information.

To alleviate the aforementioned problems, we suggest a new approach of combining a target sentence with its context. Concretely, we first concatenate the target sentence with its ① previous and ② next sentences *respectively* (i.e., previous-target & target-next), and then inject each chunk into our encoder to derive $v_{[CLS]}$ and v_{MWE} . By doing so, we expect that the target sentence can be relatively more emphasized than its context, as the target sentence naturally appears twice while its context is exposed only once. Plus, by dividing the whole sequence into two parts, it is anticipated that the encoder struggles less to extract useful information from the input. Note that $v_{[CLS]}$ is the representation for the entire chunk, which is obtained by taking the [CLS] embedding from the last layer of the encoder, and that v_{MWE} is the average of the representations of the subwords that constitute the target MWE. Lastly, the final context-sensitive feature is computed by conducting a linear transformation of the concatenation of $v_{[CLS]}$ and v_{MWE} .

On the other hand, we propose two extra techniques in order to provide a clue on the location of MWEs. While constructing token-level representations for our encoder, we employ trainable segment embeddings that draw the line between tokens for the target MWE and others. Moreover, the target MWE is repeated at the end of each chunk, following [Tayyar Madabushi et al. \(2021\)](#).

3.2 Features Based on Inner-Sentence Context

Second, we consider adding features dedicated to more effectively leveraging the information em-

bedded in the target sentence, regarding the MWE and its neighboring words as separate objects. We import some ideas from prior work for metaphor detection ([Mao et al., 2019](#); [Choi et al., 2021](#)), exploiting the conceptual relationship between metaphors and idiomatic expressions.

Initially, we assume that similar to Metaphor Identification Procedure (MIP), whose core idea is that a metaphoric word’s semantics become distinct from its lexical meaning when it is contextualized, we consider an MWE as idiomatic when its static and contextualized embeddings are heterogeneous. While the contextualized representation of the target MWE is already available from the features provided in Section 3.1, we have not yet introduced the MWE’s static representations. To implement this, we again make use of the same encoder, however, only the MWE itself (removed from its context) is presented as input to the model. We call the output of this procedure the ④ *MWE-exclusive* representation, which becomes an ingredient for realizing the ‘idiomatic’ version of MIP. Note also that according to [Garcia et al. \(2021b\)](#), static models have been considered as competitive or even better to/than contextualized models for idiomaticity detection. Therefore, we aim to reinforce our framework by employing both the options.

Meanwhile, [Choi et al. \(2021\)](#) use Selectional Preference Violation (SPV) for metaphor detection, assuming that the semantics of a metaphoric word should be distinctive from that of its context.³ We basically adopt their idea, but revise its implementation, arguing that their implementation might be somewhat defective. In detail, [Choi et al. \(2021\)](#) compute the embeddings of a target and its context exactly as we do when computing $v_{[CLS]}$ and v_{MWE} in Section 3.1. However, it is highly probable that $v_{[CLS]}$ and v_{MWE} contain similar information as they are intertwined with each other by the attention mechanism, which is undesirable when estimating separate semantics of the target and context. We thus introduce the ③ *context-exclusive* representation of the target sentence by providing our encoder with a variant of the sentence where the target MWE is masked. When combined with the features from the previous section, we expect that the inner-sentence context independent from the target MWE at all can be useful for applying the concept of SPV into idiomaticity detection.

³This time, we limit the scope of the context as the sentence encompassing a target expression (i.e., inner-sentence context), following [Choi et al. \(2021\)](#).

Model / Lang.	English	Portuguese	Galician	Overall
<i>Zero-shot setting</i>				
Baseline (BERT)	70.70	68.03	50.65	65.40
Baseline (XLM-R)	72.29	65.68	46.16	63.21
Ours (submitted)	76.42	72.82	62.92	72.27
<i>One-shot setting</i>				
Baseline (BERT)	88.62	86.37	81.62	86.46
Baseline (XLM-R)	88.45	85.03	84.02	86.56
Ours (submitted)	91.59	84.57	82.87	87.50
Ours (post-eval)	92.29	88.05	87.10	89.96

Table 2: Main results on the test set. Numbers are from the best configuration (random seed) of each model.

4 Experiments

4.1 Experimental Setup

For all experiments, we present five instances per each model with the corresponding random seeds (42, 360, 2578, 5925, 9463). Each instance is trained for 10 epochs, and its best checkpoint which shows the top performance on the development set in terms of the macro F1-score is chosen for the inference of the test set. We use a max sequence length of 300, a learning rate of $3e-5$ for the AdamW (Loshchilov and Hutter, 2019) optimizer, and a batch size of 16 for the training set and 8 for the validation and test sets. The vectors of each representation ($v_{[CLS]}$ and v_{MWE}) have 768 dimensions respectively and the learning rate is scheduled to linearly decrease after the second epoch. We leverage XLM-R(-base) (Conneau et al., 2020) as our foundation model.

4.2 Main Results

We compare the results of our method (submitted) against those of the baseline offered by the task organizers (Tayyar Madabushi et al., 2022). Although the original baseline is powered on Multilingual BERT(-base), for a comparison, we also report the performance of the baseline equipped with XLM-R(-base). Evaluation is conducted on the test set, and each model’s performance is reported according to the language on which it is tested (English, Portuguese, and Galician) and the setting it is trained (zero- and one-shot).

From Table 2, we can see that both in the zero- and one-shot settings, our model largely outperforms baselines. Notice that in the zero-shot setting, our model outperforms the baseline powered on the same foundation model (XLM-R) by more than 16% in Galician. Considering that Galician was not included in the training data, this result con-




Figure 2: Ablation study.

firms that our model is more generalizable than the baselines from the perspective of input language.

4.3 Ablation Study

We perform an ablation study to confirm whether the elements of our framework are significant. Note that all the results used for comparison are the average of the scores of different instances with five random seeds. Overall, when tested on the validation set, the final version of our approach succeed in outperforming most of the variations, especially in the one-shot setting where our decisions for selecting the final model were made. We present more detailed analysis in the following.

First, we compare our method with the variation (A) which uses only target sentences without surrounding context and the variation (B) which reflects the context by concatenating three sentences. Tayyar Madabushi et al. (2021), where the authors employ the variation (B), previously reported that it is not helpful for idiomaticity detection to consider the surrounding context of a target MWE. However, as shown in Figure 2, we find that taking the context into account following our approach (i.e., separating the context into two chunks) is in fact advantageous in all experimental settings. Furthermore, we observe that the deviation of the scores of our method is much smaller and more stable than that of not considering context. This implies that if there exists a data instance not having much information available from its target sentence, the surrounding context of the target sentence can complements the lack of information.

Contrary to our expectation, it is shown that our method is not always better than the three variations (C), (D) and (E). The variation (C) removes segment embeddings, the variation (D) stops the repetition of MWEs at the tails of ① and ②, and in the variation (E) the target MWE is recovered (not

masked) in the computation of ‘*context-exclusive*’ representation (③). We leave a detailed examination regarding these as a follow-up study.

Lastly, it turns out that the variation (F) which removes the ‘*MWE-exclusive*’ representation (④), is more favored in the zero-shot setting. Unlike the one-shot setting, where a pair of positive and negative examples for a particular MWEs can be provided, the zero-shot setting requires the evaluation of MWEs not presented in the training set, which is a more harsher condition for idiomaticity detection models. Therefore, we conjecture that static representations for the MWEs unseen during training become a little bit noisy in the zero-shot setting, failing to function following our intention.

5 Discussion

5.1 Issue on Validation Set in One-shot Setting

In the one-shot setting, we expect that a pair of data instances (one for idiomatic and the other for non-idiomatic) per every MWE in the test set should be provided to the model we train. Likewise, if one wants to confirm that the validation set is rigorous enough to be a substitute for the test set in the procedure of selecting hyperparameters, every MWE in the validation set should have two corresponding instances in the training set. During the competition for SemEval 2022 Task 2, we have discovered that the necessary condition holds for the validation set in the practice phase, while **it does not hold in the evaluation phase**. In other words, the training set provided in the practice phase incorporated data instances that correspond to MWEs in the validation set. However, as the training set has been substituted with a new version, a problem has arisen where MWEs in the newly released training set do not match with those in the validation set. We conjecture that this discrepancy prevents one’s optimal actions in choosing the best models.

To prove our hypothesis, we test a variant whose performance on the validation set is not optimal, but has the potential of working well when evaluated on the test set. Specifically, we replicate our experiments, but do not choose the best instance based on validation performance. Instead, we simply choose the model instance trained until 9 epochs and compare it to baselines. As shown in Table 2, we find that the instance chosen based on the validation set (i.e., ours (submitted)) is worse than the randomly selected one (i.e., ours (post-eval)), implying that the inappropriateness of the

Form of MWEs	Validation	Test
<i>Zero-shot setting</i>		
Original form	76.34	71.72
Inflectional form	76.36	70.01
<i>One-shot setting</i>		
Original form	88.14	89.80
Inflectional form	89.33	88.95

Table 3: Performance gap with form changes in MWEs.

validation set in the evaluation phase might hinder correct comparisons between submitted models.

5.2 Impact of Form Changes in MWEs

When MWEs are repeated at the end of input sequences in ① and ② and embedded solely in ④ in our implementation, we copy them from target sentences so that we can preserve their inflectional form appearing in the sentences, rather than adopting their original form. To confirm the effectiveness of this approach, we conduct an experiment where we replace the MWEs with their original form.

From Table 3, we observe that unlike the case on the validation set where applying inflectional form is always helpful, it turns out that when evaluated on the test set, employing inflectional form is not beneficial for performance improvement, contrary to our expectation. The idea of having utilized the inflectional form of MWEs is from our conjecture that compositional and static representations of MWEs should be computed from the same form for a fair comparison between them. However, it seems that it is more effective to provide a model with a MWE’s original form in addition to its inflectional form such that the model can extract more information from the both sides. We leave thorough analysis on this phenomenon as future work.

6 Conclusion

In this work, we investigate the method of implementing better idiomaticity detection models by considering different levels of contextualization. We propose four features grounded on the surrounding and inner-sentence context of a target MWE, showing that these features are effective in improving performance. Moreover, we present a discussion on the issue related to the validation set in the one-shot setting and the impact of the form of MWEs. As future work, we plan to develop a method of designing the interaction between related features in a more sophisticated fashion, instead of simply concatenating them.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2020-0-01373, Artificial Intelligence Graduate School Program (Hanyang University)).

References

- Timothy Baldwin and Su Nam Kim. 2010. Multiword expressions. *Handbook of natural language processing*, 2:267–292.
- Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee, and Jongwuk Lee. 2021. [MelBERT: Metaphor detection via contextualized late interaction using metaphorical identification theories](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1763–1773, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Mathieu Constant, Gülşen Eryiğit, Johanna Monti, Lonneke Van Der Plas, Carlos Ramisch, Michael Rosner, and Amalia Todirascu. 2017. Multiword expression processing: A survey. *Computational Linguistics*, 43(4):837–892.
- 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.
- Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2021a. [Assessing the representations of idiomaticity in vector models with a noun compound dataset labeled at type and token levels](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2730–2741, Online. Association for Computational Linguistics.
- Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2021b. [Probing for idiomaticity in vector space models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3551–3564, Online. Association for Computational Linguistics.
- Pragglejaz Group. 2007. Mip: A method for identifying metaphorically used words in discourse. *Metaphor and symbol*, 22(1):1–39.
- Zhenxi Lin, Qianli Ma, Jianguye Yan, and Jieyu Chen. 2021. [CATE: A contrastive pre-trained model for metaphor detection with semi-supervised learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3888–3898, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Changsheng Liu and Rebecca Hwa. 2019. A generalized idiom usage recognition model based on semantic compatibility. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6738–6745.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#).
- Rui Mao, Chenghua Lin, and Frank Guerin. 2019. [End-to-end sequential metaphor identification inspired by linguistic theories](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3888–3898, Florence, Italy. Association for Computational Linguistics.
- Siva Reddy, Diana McCarthy, and Suresh Manandhar. 2011. [An empirical study on compositionality in compound nouns](#). In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 210–218, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- Gerard Steen. 2010. *A method for linguistic metaphor identification: From MIP to MIPVU*, volume 14. John Benjamins Publishing.
- Harish Tayyar Madabushi, Edward Gow-Smith, Marcos Garcia, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2022. SemEval-2022 Task 2: Multilingual Idiomaticity Detection and Sentence Embedding. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*. Association for Computational Linguistics.
- Harish Tayyar Madabushi, Edward Gow-Smith, Carolina Scarton, and Aline Villavicencio. 2021. [AStitchInLanguageModels: Dataset and methods for the exploration of idiomaticity in pre-trained language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3464–3477, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yorick Wilks. 1975. A preferential, pattern-seeking, semantics for natural language inference. *Artificial intelligence*, 6(1):53–74.

Ziheng Zeng and Suma Bhat. 2021. [Idiomatic expression identification using semantic compatibility](#). *Transactions of the Association for Computational Linguistics*, 9:1546–1562.