MetaPro Online:
A Computational Metaphor Processing Online System

Rui Mao[*♠], **Xiao Li**[♠], **Kai He**[♠], **Mengshi Ge**[♣] and **Erik Cambria**[♣]
[♣]Nanyang Technological University, Singapore
[♠]Ruimao Tech, Shenzhen, China
{mao.r, li.x, he.kai}@ruimao.tech; {mengshi.ge, cambria}@ntu.edu.sg

Abstract

Metaphoric expressions are a special linguistic phenomenon, frequently appearing in everyday language. Metaphors do not take their literal meanings in contexts, which may cause obstacles for language learners to understand them. Metaphoric expressions also reflect the cognition of humans via concept mappings, attracting great attention from cognitive science and psychology communities. Thus, we aim to develop a computational metaphor processing online system, termed MetaPro Online[1], that allows users without a coding background, e.g., language learners and linguists, to easily query metaphoricity labels, metaphor paraphrases, and concept mappings for non-domain-specific text. The outputs of MetaPro can be directly used by language learners and natural language processing downstream tasks because MetaPro is an end-to-end system.

1 Introduction

Metaphors are frequently used in our everyday language. Relevant statistics (Cameron, 2003; Martin, 2006; Steen et al., 2010) found that about one-third of sentences in corpora contain metaphors. It is possible to use approximately six metaphors during a minute of speaking (Geary, 2011).

Metaphors are a non-negligible linguistic phenomenon for linguistic learners and researchers. On the other hand, Lakoff and Johnson (1980) argued that metaphors are not just a linguistic phenomenon, because metaphors also reflect human cognition. Human perceptions of LOVE are not complete without metaphors, e.g., love is a *journey, magic, attraction, madness, nurturance*[2], etc.

There are two important research domains in computational metaphor processing: linguistic and conceptual. The former area focuses on the surface realization of metaphors, e.g., metaphor identification and interpretation.

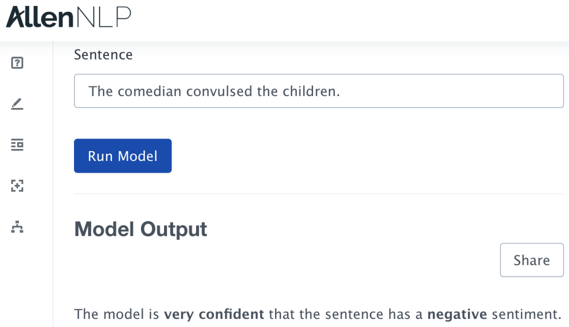

(a) Fine-tuning-based sentiment analysis.

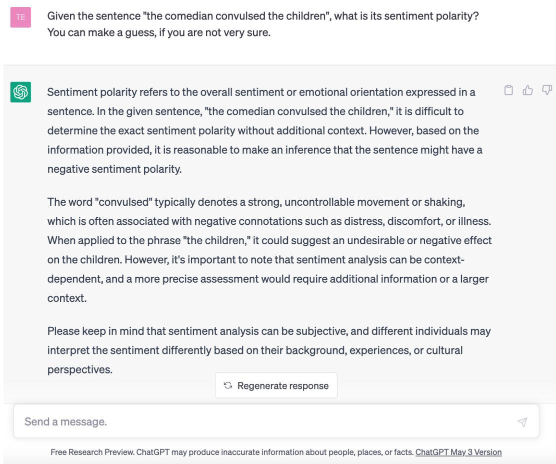

(b) Generative AI-based sentiment analysis.

(c) Machine translation.

Figure 1: Errors in current AI applications, caused by a metaphoric expression.

The motivation is that metaphors do not take their literal meanings in contexts, which may cause difficulties for language learners and machines to

* The work was done when he worked at Ruimao Tech.
[1]Website: https://metapro.ruimao.tech/
[2]Italics denotes metaphors.

Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics Volume 3: System Demonstrations, pages 127–135
July 10-12, 2023 ©2023 Association for Computational Linguistics

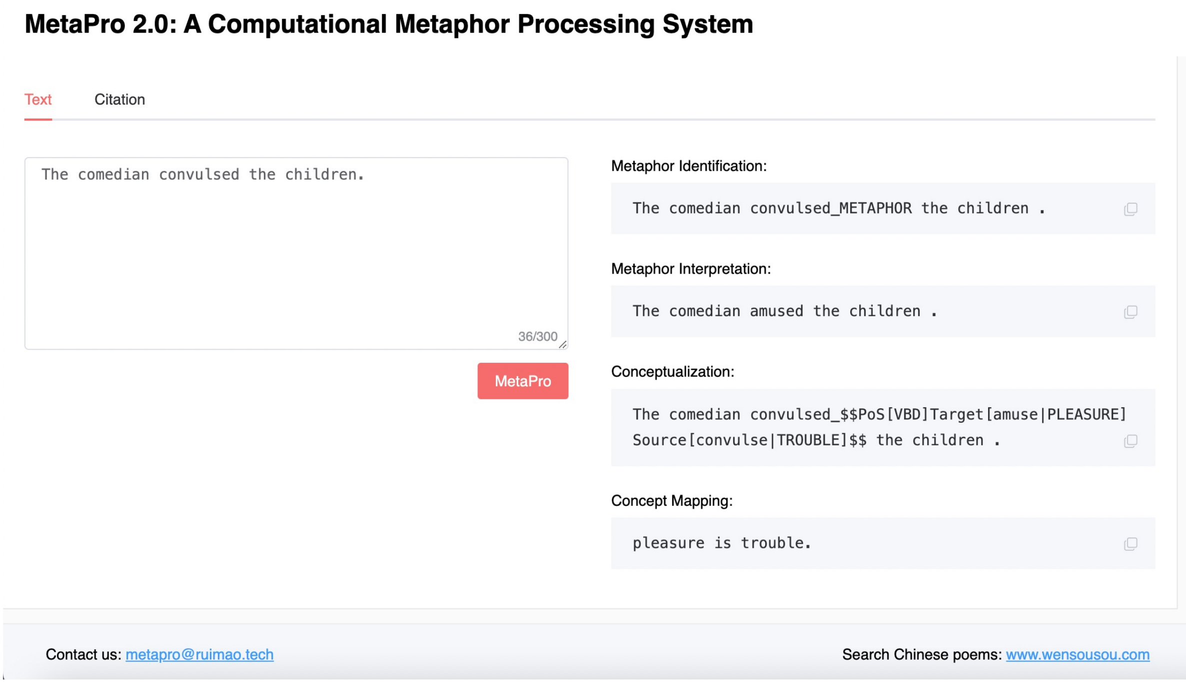

Figure 2: MetaPro Online.

understand the real meanings of metaphors (Cambria et al., 2017). As seen in Figure 1 (a), given the metaphoric expression, “the comedian *convulsed* the children”, a sentiment classifier yields an incorrect negative label for the sentence. In Figure 1 (b), the generative AI seems to explain why such an error occurs – “The word ‘convulsed’ typically denotes a strong, uncontrollable movement or shaking, which is often associated with negative connotations”. AI notes that “convulsed” triggers the negative prediction for the sentence based on its literal meaning. Figure 1 (c) demonstrates that literally translating a metaphor into another language may cause interpretation difficulties because of cultural differences. For example, the given sentence is translated as “the comedian caused the children convulsions” in Chinese, which carries a very different meaning from the real English meaning. Thus, a linguistic metaphor processing system (Mao et al., 2022) aims to automatically identify metaphors and paraphrase them into their literal counterparts to mitigate the interpretation barrier for users.

Conceptual metaphor processing focuses on identifying and mapping source and target concept domains of metaphors. The motivation is to understand the cognition of humans from their metaphoric expressions. For example, given a metaphoric expression, “she *attacked* his argument”, Lakoff and Johnson (1980) argued that conceptually the ARGUMENT concept (target domain) is mapped to a WAR concept (source domain). One uses WAR to metaphorically describe ARGUMENT, reflecting strategies, attack, and defense behaviors associated with ARGUMENT. Lakoff and Johnson (1980) also believed that these concepts naturally frame our thinking and behaviors when arguing. Researchers can study cognition from metaphor concept mappings (see Section 6). Thus, a conceptual metaphor processing system aims to automatically generate source and target concept agents for a metaphor to represent the concept mappings.

Many previous works focused on metaphor identification tasks (Mao et al., 2019; Feng and Ma, 2022), metaphor interpretation tasks (Shutova, 2010; Mao et al., 2018), and concept mappings (Rosen, 2018; Ge et al., 2022). However, these methods cannot process metaphors from end to end. We proposed an end-to-end metaphor identification and interpretation method, the first version of MetaPro (Mao et al., 2022), where metaphoric single words can be paraphrased into their literal counterparts, and metaphoric multi-word expressions can be explained with clauses. However, a programming package cannot be easily used by non-programmers, e.g., language learners and linguistic researchers. On the other hand, automatic concept mappings have not been achieved on the

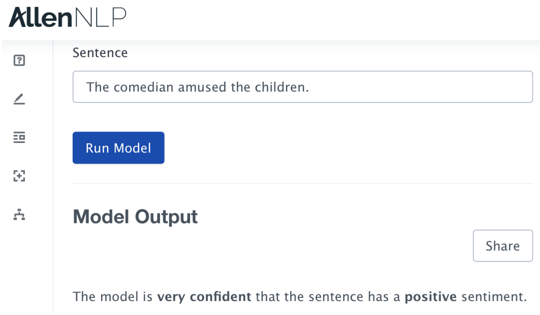

(a) Fine-tuning-based sentiment analysis.

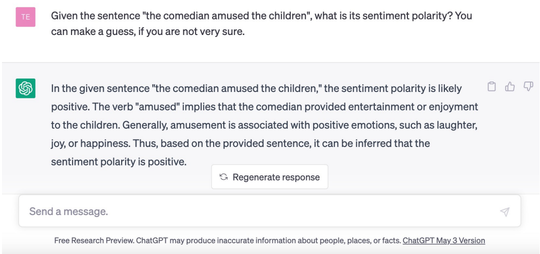

(b) Generative AI-based sentiment analysis.

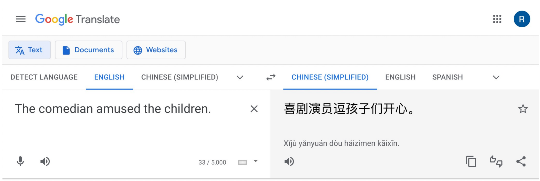

(c) Machine translation.

Figure 3: Correct outputs after MetaPro pre-processing.

sentence level, although we proposed a concept mapping method on the word-pair level in the work of Ge et al. (2022). Thus, we aim to develop a computational metaphor processing online system, MetaPro Online, allowing users without a coding background to easily query metaphoricity labels, metaphor interpretations, and concept mappings for non-domain-specific text.

As seen in Figure 2, given a query sentence, "the comedian *convulsed* the children", MetaPro can identify "*convulsed*" as a metaphor. Then, the sentence is paraphrased as "the comedian amused the children" for interpretation. The source concept is TROUBLE. The target concept is PLEASURE. Then, the concept mapping is represented as PLEASURE IS TROUBLE. We can also observe that the sentiment analysis classifier, the generative AI, and the machine translation system can all yield satisfying outputs after using the paraphrased sentence in Figure 3. The concept mapping "PLEASURE IS TROUBLE", generated by MetaPro, implies that extreme happiness can be accompanied by trouble, e.g., causing the children convulsions. Such

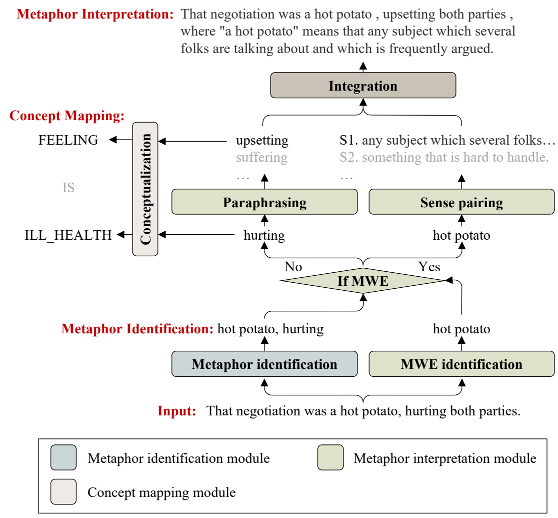

Figure 4: The framework of MetaPro. MWE denotes multi-word expressions.

implicit meanings can be used for analyzing human cognition because metaphors are stimulated by subjective experiences (Grady, 1997).

2 Related Work

According to the survey of Ge et al. (2023), there have been metaphor identification (Mao et al., 2019; Feng and Ma, 2022), interpretation (Shutova, 2010; Mao et al., 2018), and concept mapping (Rosen, 2018; Ge et al., 2022) research works in the computational linguistics community, while readily used end-to-end systems for supporting downstream applications are rare. Martin (1990) presented a metaphor interpretation system named Metaphor Interpretation, Denotation, and Acquisition System (MIDAS). The system was developed to answer metaphorical queries about Unix. Narayanan (1997) proposed a Knowledge-based Action Representations for Metaphor and Aspect (KARMA) system. The system is based on hand-coded executing schemes for motion-verb concept reasoning in an economic domain. Barnden and Lee (2002) proposed a first-order logic rule-based system for concept mapping analysis in the mental state description domain. However, the limited manual rules in the aforementioned works cannot be used for analyzing text from a broader domain. Mao et al. (2022) proposed a non-domain-specific end-to-end system for metaphor identification and interpretation. They evaluated the system on a sentiment analysis task. However, it cannot generate concept mappings and does not provide a graphical user interface for non-programmers.

3 MetaPro

The user interface of MetaPro Online is succinct. There is just an input window and a processing button (see Figure 2). The back-end consists of three modules: metaphor identification, metaphor interpretation, and concept mappings. The overall framework can be viewed in Figure 4.

First, the metaphor identification module detects metaphors on the token level, given an input sentence, e.g., "that negotiation was a *hot potato*, *hurting* both parties". Next, the metaphor interpretation module paraphrases the identified metaphoric single-word expressions (e.g., "*hurting*" means "upsetting"). It explains metaphoric multi-word expressions with clauses (e.g., "where 'a hot potato' means that any subject which several folks are talking about and which is frequently argued"). Finally, the concept mapping module abstracts the source ("ILL_HEALTH") and target ("FEELING") concepts of a metaphor via the lemma of the original metaphoric word ("hurt") and the lemma of its paraphrase ("upset"), respectively. The mapping is represented as "a target concept is a source concept" (e.g., FEELING IS ILL_HEALTH). The current version of MetaPro Online cannot paraphrase and conceptualize metaphoric multi-word expressions. Thus, we explain them via clauses and omit their concept mapping generations.

We introduce each module below at the concept level to help readers understand the mechanisms of MetaPro Online. Since MetaPro is an ensemble system combining three research outcomes, we omit the algorithmic details here. One may refer to our previous works, Mao and Li (2021); Mao et al. (2022); Ge et al. (2022), to understand the technical details of metaphor identification, metaphor interpretation, and concept mapping modules, respectively.

3.1 Metaphor identification

The used algorithm of our metaphor identification module was first proposed in the work of Mao and Li (2021). We embed this algorithm into MetaPro Online because it achieved state-of-the-art performance on token-level metaphor identification tasks.

Metaphor identification is a binary classification task, classifying a word as metaphoric or literal. Our metaphor identification module uses a multi-task learning framework, simultaneously learning sequential metaphor identification and Part-of-Speech (PoS) tagging tasks. The motivation is that previous works (Wu et al., 2018; Su et al., 2020) found PoS tags are effective features for learning metaphor identification. Thus, we introduced the learning of PoS tagging as an auxiliary task to fuse the features learned from PoS tagging.

To boost the multi-task learning framework, we also proposed a novel soft-parameter sharing mechanism, Gated Bridging Mechanism (GBM) in the work of Mao and Li (2021). The intuition is that GBM allows useful information from a neighbor tower to pass through the gate and fuse with hidden states learned in the private tower, while the gates filter out the useless information. We demonstrated the effectiveness of the proposed multi-task learning model on both metaphor identification and aspect-based sentiment analysis tasks.

3.2 Metaphor interpretation

The algorithm for the metaphor interpretation module of MetaPro was first proposed in the work of Mao et al. (2022). We embed this algorithm into MetaPro because it can paraphrase metaphoric single-word expressions and explain metaphoric multi-word expressions. The paraphrases and explanations help users understand the intended meanings of metaphors and can be processed by other downstream natural language processing systems because the outputs of the interpretation module are also natural language.

Given an identified metaphor, if the metaphor is a single-word expression, RoBERTa-large (Liu et al., 2019) and WordNet (Fellbaum, 1998) are used for paraphrasing the metaphor. We masked out a metaphor token and used the masked word prediction of RoBERTa to predict the probability distribution of candidate words. The candidate words are sourced from the WordNet hypernyms and synonyms of the lemma of the metaphor, providing paraphrase constraints for a metaphor. The hypothesis is that a metaphor can be paraphrased into one of its hypernyms and synonyms (Mao et al., 2018). The word forms of the candidate words are aligned with the original metaphor. In the work of Mao et al. (2022), we developed a word form alignment dictionary, e.g., `{..., 'upset': {'VBG': 'upsetting', 'VBD': 'upset', ...}, ...}` by parsing a Wikipedia dump. Thus, a lemma can map to any form, given a Penn Treebank PoS label (Marcus et al., 1993).

Mao et al. (2022) also developed a dictionary- and rule-based algorithm to identify metaphoric multi-word expressions. They did not use a neural network model, because the identified metaphoric multi-word expressions are finally mapped to their dictionary explanations for metaphor interpretation. The dictionary- and rule-based algorithm can directly identify and map them to their dictionary explanations without using another model for mapping. In order to improve the generalization ability of this algorithm, Mao et al. (2022) developed two feature sets, namely a lemma feature set that consists of the lemmas of multi-word expressions and a triplet feature set that consists of dependency triplet features. Both feature sets can map features with the corresponding multi-word expressions. An input sentence is pre-processed with lemmatization and dependency parsing first. If features from any feature sets are the subsets of the pre-processed sentence, the corresponding multi-word expressions are detected. If there is an overlap between an identified multi-word expression and metaphoric tokens given by the metaphor identification module, the multi-word expression is explained via a clause by using its selective dictionary explanation. The final metaphor interpretation output integrates the paraphrases of single-word metaphors and metaphoric multi-word expressions. The word form of a paraphrase was aligned with the original metaphor during the candidate word preparation stage.

3.3 Concept mapping

The used algorithm of our concept mapping module was first proposed in the work of Ge et al. (2022). We embed this algorithm into MetaPro, because it can abstract concepts for words. The effectiveness of the abstracted concepts has been proved in metaphor identification and human evaluation.

The conceptual metaphor theory was proposed by Lakoff and Johnson (1980). They empirically summarized some examples of concept mappings to explain their theory. However, there is no theoretical research to explain how to conceptualize metaphors and generate mappings. For example, it was not clearly defined, if a target concept should be abstracted from the paraphrase or the context word of a metaphor. By viewing concept mapping examples given by Lakoff and Johnson (1980); Lakoff (1994), evidence from both sides can be observed. Thus, we choose to generate target concepts from the paraphrases of metaphors. The source concepts are given by the original metaphor.

Another challenge is there is no theoretical research about what the abstractness level of a concept agent should be to represent a word, to our best knowledge. Thus, we follow the hypothesis of Ge et al. (2022) that an appropriate concept agent should represent the main senses of a word.

We developed a conceptualization algorithm based on WordNet and a statistical knee point algorithm (Satopaa et al., 2011) in the work of Ge et al. (2022). First, a word, e.g., a metaphor or a paraphrase, was aligned to its nominal form via WordNet and a gestalt pattern matching algorithm (Ratcliff and Metzener, 1988). Next, given a noun, we retrieved all hypernym paths from the noun node to the root node. Different paths represent different senses of the noun. The hypernyms on different levels of a path denote concepts with different abstractness levels because WordNet is a conceptually structured knowledge base. Next, the hypernyms on each path are rated by a linear score function. A higher score denotes that the hypernym is more abstract. The overall score of a hypernym is given by the sum of the hypernym scores on all its distributed paths. Next, we computed the knee point with the overall scores of all hypernyms. A hypernym is selected as a concept agent if it covers the same number of senses as the knee point hypernym and it is more concrete (a lower score) than the knee point hypernym. Otherwise, the knee point hypernym is selected as the concept agent.

Based on the above method, we can compute a source concept with a metaphor, and a target concept with a paraphrase. Then, the concept mapping is given by "a target concept is a source concept".

3.4 Training data and lexical resources

We use VU Amsterdam Metaphor Corpus (Steen et al., 2010) as the data source to train the metaphor identification module. The training set combines the training and validation sets prepared by Leong et al. (2018), including nominal, verbal, adjective, and adverb metaphors from conversations, fiction, academic text, and news. For the statistics of the dataset, please view the work of Mao et al. (2022).

Metaphor paraphrases are based on WordNet and masked word predictions. Thus, no training set is required for learning metaphor paraphrases. The metaphoric multi-word expression interpretation is a dictionary- and rule-based algorithm. The feature and explanation dictionaries contain idiomatic multi-word expressions that were sourced from The Idioms[3] and the collection of Agrawal et al. (2018). We have defined 3,560 lemma pairing features and 3,470 dependency triplet pairing features for 3,050 idiomatic multi-word expressions in the work of Mao et al. (2022). On average, each multi-word expression has 2.7 explanations.

The concept mapping module is based on a statistical learning algorithm, using WordNet as the only lexical resource. Thus, we do not use any training set to learn concept mappings.

4 Supporting Downstream Tasks

MetaPro was used as a text pre-processing technique, where the metaphor interpretation outputs were fed into sentiment analysis classifiers instead of the original inputs. We observe that MetaPro improved the performance of Vader (Hutto and Gilbert, 2014), AllenNLP sentiment analysis[4], and Azure sentiment analysis[5] on a financial news headline sentiment analysis task (Cortis et al., 2017) by 1.5%, 2.2%, and 4.7% accuracy, respectively, compared with the results given by the original news headlines (Mao et al., 2022). On the other hand, we also observe that using MetaPro-generated concept mappings as features could bring a classifier extra gains in accuracy (+1.1% and +1.9%) for a depression detection task (Han et al., 2022). The benchmarking dataset was from the work of Shen et al. (2017). More importantly, the concept mapping features help the classifier explain the common concept mappings between depressed and non-depressed groups. Besides the accuracy gains, this is particularly helpful for cognitive science because MetaPro provides an automatic solution for analyzing concept mapping patterns via metaphors at scale.

Detailed performance benchmarking of different technical components of MetaPro with state-of-the-art baselines on diverse evaluation tasks and datasets can be viewed from the works of Mao and Li (2021); Mao et al. (2022); Ge et al. (2022).

5 Evaluation

Besides the performance improvements and explainability enhancements in downstream tasks, we also qualitatively evaluate the practicality of MetaPro according to the criteria proposed by Shutova (2015) in the following sections. She proposed to evaluate a computational metaphor processing system from two aspects, namely the levels of analysis and applicability.

5.1 Levels of analysis

Linguistic metaphor. MetaPro has the capacity to analyze various forms of linguistic metaphors, encompassing both conventional and novel metaphors, including single- and multi-word expressions. It possesses the capability to handle a diverse range of linguistic metaphors without restriction on specific syntactic constructions, thus offering a comprehensive approach to metaphor processing.

Conceptual metaphor. MetaPro can abstract source and target concepts from original single-word metaphors and paraphrases, respectively. However, the current version is incapable of conceptualizing metaphoric multi-word expressions and metaphoric sentences.

Extended metaphor. The current version of MetaPro cannot process extended metaphors on the document level, due to the training set of the metaphor identification module does not contain extended metaphors. We cannot find a helpful dataset to study extended metaphors to our best knowledge.

Metaphorical inference. The current version of MetaPro cannot process metaphorical inference on the document level, because we cannot find a helpful dataset to study metaphorical inference to our best knowledge.

5.2 Applicability

Task coverage. MetaPro can identify metaphors and interpret them from both linguistic and conceptual perspectives.

Easily to integrate. The outputs of metaphor interpretation and concept mapping modules are natural language. Thus, MetaPro can be integrated with other natural language processing systems.

Unrestricted text. MetaPro can process unrestricted real-world natural language text. The current version cannot directly process emojis and spelling errors commonly appearing on social media. The maximum input length is 512 tokens after Byte-Pair Encoding (Radford et al., 2019). For efficient computing, we set up the maximum input length as 300 characters for MetaPro Online at the current version.

Be open-domain. MetaPro is not domain-specific. The training data for metaphor identification were sourced from VU Amsterdam Metaphor Corpus,

[3] https://www.theidioms.com/
[4] https://allennlp.org/
[5] https://azure.microsoft.com

including diverse topics, different genres, and concept domains. Our metaphor interpretation is based on a pre-trained language model, trained with open-domain text without fine-tuning. The used knowledge bases, e.g., WordNet and multi-word expression dictionaries, offer general semantic knowledge of English words.

Task and knowledge dependency. Current metaphor interpretation and concept mapping modules of MetaPro depend on WordNet and multi-word expression processing dictionaries. This is because we cannot find a large annotated dataset to train neural network models to achieve the functions from end to end in a supervised fashion. We use WordNet for supporting the task of paraphrasing metaphors because simply using a medium size pre-trained language model can hardly yield satisfying results in the context of unsupervised learning.

Word class and syntax diversity. The inputs of MetaPro are sentences with diverse syntactic constructions. The current version of MetaPro targets to identify, interpret and conceptualize open-class metaphors. It also explains metaphoric multi-word expressions that contain other PoS. We focus on open-class metaphors because they contain richer semantic information than closed-class ones. This is more helpful for downstream tasks.

6 Use Case

Besides the examples in Figures 1 and 3, we reported the use cases of MetaPro in sentiment analysis in the work of Mao et al. (2022). For example, given "Rio Tinto CEO Sam Walsh rejects fears over China growth, demand", the three examined classifiers yielded incorrect "negative" predictions. This is probably because "fears" and "reject" likely appear in negative contexts. However, after MetaPro paraphrasing the original input as "Rio Tinto CEO Sam Walsh eliminates concerns over China growth, demand", the classifiers can yield correct "positive" predictions.

Han et al. (2022) reported the use cases in depression detection, where concept mappings are additional features besides the original text. They believe that metaphor concept mappings reflect the inner world of people because they were not explicitly presented in the text. For example, they found that LEVEL IS IMPORTANCE is a representative concept mapping for depressed people. This may result in more stress, if a person frequently maps an objective measure in the LEVEL concept to a subjective feeling concept, e.g., IMPORTANCE.

7 Conclusion

We proposed MetaPro online in this work, which is a computational metaphor processing online system. Compared with previous works, MetaPro can identify metaphors, paraphrase them into their literal counterparts, and generate concept mappings from end to end. The system can process unrestricted and non-domain-specific English text. The user interface is very friendly to non-programmers. Thus, it can help language learners to understand the real meanings of English metaphors. We also demonstrated the performance improvements of using MetaPro on downstream AI applications, e.g., using MetaPro to automatically obtain concept mappings from social media posts to study cognitive patterns exhibited by individuals diagnosed with depression. The above use cases show that MetaPro has huge application potential in diverse domains.

However, the current version cannot paraphrase and conceptualize metaphoric multi-word expressions, which is important for sentic computing (Cambria et al., 2022). It cannot process non-English text, extended metaphors, and metaphorical inference as well. We will fill this gap in future work and strive to enhance the precision and information processing capacities of MetaPro by developing more advanced algorithms, thereby providing enhanced support for linguistic and cognitive science research endeavors.

Ethics and Broader Impact Statement

This article follows the ACL Code of Ethics. We comply with the licenses of all used datasets. Although there was no sensitive data used for training our models or developing our knowledge bases in MetaPro, we encourage all downstream applications can honor the ethical code for conducting linguistic and cognitive research. The broader impacts include but are not limited to using the tool to study the cognitive patterns of a certain group of people or a person, and using this tool to falsify original text. According to Mao et al. (2023), there are certain biases in pre-trained language models. We cannot guarantee that MetaPro can yield unbiased outputs, because it depends on a pre-trained language model. Besides, MetaPro is not a perfect system. The errors generated by MetaPro may also introduce biases for downstream applications.

References

Ruchit Agrawal, Vighnesh Chenthil Kumar, Vigneshwaran Muralidharan, and Dipti Misra Sharma. 2018. No more beating about the bush: A step towards idiom handling for Indian language NLP. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC)*, pages 319–324.

John A Barnden and Mark G Lee. 2002. An artificial intelligence approach to metaphor understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.

Erik Cambria, Qian Liu, Sergio Decherchi, Frank Xing, and Kenneth Kwok. 2022. SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3829–3839.

Erik Cambria, Soujanya Poria, Alexander Gelbukh, and Mike Thelwall. 2017. Sentiment analysis is a big suitcase. *IEEE Intelligent Systems*, 32(6):74–80.

Lynne Cameron. 2003. *Metaphor in Educational Discourse*. A&C Black.

Keith Cortis, André Freitas, Tobias Daudert, Manuela Huerlimann, Manel Zarrouk, Siegfried Handschuh, and Brian Davis. 2017. SemEval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval)*, pages 519–535.

Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.

Huawen Feng and Qianli Ma. 2022. It's better to teach fishing than giving a fish: An auto-augmented structure-aware generative model for metaphor detection. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 656–667.

Mengshi Ge, Rui Mao, and Erik Cambria. 2022. Explainable metaphor identification inspired by conceptual metaphor theory. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10681–10689.

Mengshi Ge, Rui Mao, and Erik Cambria. 2023. A survey on computational metaphor processing techniques: From identification, interpretation, generation to application. *Artificial Intelligence Review*.

James Geary. 2011. *I Is an Other: The Secret Life of Metaphor and How It Shapes the Way We See the World*. Harper Collins.

Joseph Edward Grady. 1997. *Foundations of Meaning: Primary Metaphors and Primary Scenes*. University of California, Berkeley.

Sooji Han, Rui Mao, and Erik Cambria. 2022. Hierarchical attention network for explainable depression detection on Twitter aided by metaphor concept mappings. In *Proceedings of the 29th International Conference on Computational Linguistics (COLING)*, pages 94–104.

Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8, pages 216–225.

George Lakoff. 1994. *Master Metaphor List*. University of California.

George Lakoff and Mark Johnson. 1980. *Metaphors We Live by*. University of Chicago press.

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 (FigLang)*, pages 56–66.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv e-prints*, pages arXiv–1907.

Rui Mao and Xiao Li. 2021. Bridging towers of multi-task learning with a gating mechanism for aspect-based sentiment analysis and sequential metaphor identification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13534–13542.

Rui Mao, Xiao Li, Mengshi Ge, and Erik Cambria. 2022. MetaPro: A computational metaphor processing model for text pre-processing. *Information Fusion*, 86-87:30–43.

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 (ACL)*, volume 1, pages 1222–1231.

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 (ACL)*, pages 3888–3898.

Rui Mao, Qian Liu, Kai He, Wei Li, and Erik Cambria. 2023. The biases of pre-trained language models: An empirical study on prompt-based sentiment analysis and emotion detection. *IEEE Transactions on Affective Computing*.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The penn treebank. *Computational Linguistics*, 19(2):313–330.

James H 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. *Trends in Linguistics Studies and Monographs*, 171:214.

Srini Narayanan. 1997. Knowledge-based action representations for metaphor and aspect (KARMA). *Computer Science Division, University of California at Berkeley dissertation*.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI*.

John W Ratcliff and David E Metzener. 1988. Pattern matching: The gestalt approach. *Dr Dobbs Journal*, 13(7):46.

Zachary Rosen. 2018. Computationally constructed concepts: A machine learning approach to metaphor interpretation using usage-based construction grammatical cues. In *Proceedings of the Workshop on Figurative Language Processing (FigLang)*, pages 102–109.

Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. 2011. Finding a "kneedle" in a haystack: Detecting knee points in system behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops (ICDCS-Workshops)*, pages 166–171. IEEE.

Guangyao Shen, Jia Jia, Liqiang Nie, Fuli Feng, Cunjun Zhang, Tianrui Hu, Tat-Seng Chua, and Wenwu Zhu. 2017. Depression detection via harvesting social media: A multimodal dictionary learning solution. In *Joint Conference on Artificial Intelligence (IJCAI)*, pages 3838–3844.

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 (NAACL-HLT)*, pages 1029–1037.

Ekaterina Shutova. 2015. Design and evaluation of metaphor processing systems. *Computational Linguistics*, 41(4):579–623.

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.

Chuandong Su, Fumiyo Fukumoto, Xiaoxi Huang, Jiyi Li, Rongbo Wang, and Zhiqun Chen. 2020. Deepmet: A reading comprehension paradigm for token-level metaphor detection. In *Proceedings of the Second Workshop on Figurative Language Processing (FigLang)*, pages 30–39.

Chuhan Wu, Fangzhao Wu, Yubo Chen, Sixing Wu, Zhigang Yuan, and Yongfeng Huang. 2018. Neural metaphor detecting with cnn-lstm model. In *Proceedings of the Workshop on Figurative Language Processing (FigLang)*, pages 110–114.