Metaphor processing in the medical domain via Linked Data and Language Models*

Lucía Pitarch[0000−0002−6734−8808]

Univeristy of Zaragoza, Zaragoza, Spain
lpitarch@unizar.es

Abstract. This thesis proposes a hybrid approach that benefits from Natural Language Processing and Semantic Web technologies for computational metaphor processing. Metaphors are linguistic devices that enable us to perceive and express a concept in terms of another similar one. Designing systems that allow their explicit identification and interpretation can highly facilitate communication in sensitive and obscure contexts such as the medical one. This proposal seeks the identification, understanding, generation, and manipulation of metaphors while providing novel datasets and baselines to exploit Languages Models and Linked Data in the context of figurative knowledge. The developed methodologies will be validated by their application into a specific communication tool between cancer patients and healthcare professionals.

Keywords: Metaphors · LLOD · Ontologies · PTLMs · Medical Domain.

1 Introduction

Metaphors are not only poetic resources to embellish communication, but very common linguistic devices that enable us to talk about one thing in terms of another in every kind of communication [19]. Metaphors establish correspondences between a target domain[1], which is the one implicitly trying to be expressed, and one or more source domains, which are explicitly represented on the text. By doing so, some characteristics of the target domain are highlighted while others are shadowed. In short, metaphors focus and shape how we perceive the world [19]. By establishing associations between source and target domains, metaphors make communications more economical and efficient. They can be used to fill in lexical gaps, motivate semantic change, describe personal experience and ground it into some common knowledge shared between speakers, or influence decision making by priming the speaker with a different idea [35].

* Supported by the Spanish project PID2020-113903RB-I00 (AEI/FEDER, UE), by DGA/FEDER, and by the EU research and innovation program HORIZON Europe 2021 through the "4D PICTURE" project under grant agreement 101057332

[1] Domains in this linguistic field are understood differently than in the Semantic Web community, in this work we consider domains as the background knowledge needed to understand a concept, it is to some extent similar to a semantic field.

Having such big impact in communication, it is no surprise how common they are in the medical environment. As studied by Semino et al. [32], patients used journey and violence metaphors around 1.5 times per 1000 words in their discourse. Complementary, Casarett [3] pointed our that oncologists use at least one metaphor in every conversation with advanced cancer patients. Benefits of metaphor usage in the medical domain are twofold. For patients, they have been proven as useful communicative devices by aiding them make sense of difficult and abstract experiences [17] and when used as coping mechanisms [16]. For clinicians, they can help them explain complex medical terms [28], reinforce or reorient potentially unhealthy misconceptions [10], help patients gain foresight on their condition [27], aid the initiation of difficult topics [9] or make patients aware about risks [20]. Yet, as shown by Landau et al., using the wrong metaphor can backfire [20]; thus, studying what metaphors are beneficial for individual patients under particular conditions (e.g., illness, treatment phase) becomes crucial.

While previous research on metaphors in the medical domain has been mainly conducted using manual effort, it is very time and resource-consuming, and hardly adaptable to new patients or to changes in the patient's condition. Thus, complementing manual effort with Computational Metaphor Processing (CMP) seems like the logical way to go. The aim of this thesis is to generate a CMP tool that facilitates the understanding and generation of metaphors in the medical domain. The tool must be transparent, highly interpretable, adaptable to individual needs, scalable, and dynamic.

The underlying hypothesis behind this thesis is that a mixed approach that uses both Language Models and Linked Data will provide the needed flexibility, adaptability, interpretability, transparency, and dynamic processing of metaphors. The development of such hybrid technology will lead to a powerful communication tool between clinicians and patients. Moreover, by pursuing this goal, the thesis will contribute to the technological field by gaining insights into the abilities of Natural Language Processing (NLP) and current Ontological models to encode and represent figurative knowledge, and in the humanistic domain by providing large structured data which can further be used to understand the cognitive processes that guide metaphoric expression or the socio-cultural and medical preferences for different metaphors in the medical domain.

2 State of the Art

2.1 Computational Metaphor Processing

Tasks: In previous systematic reviews[14, 30] CMP has been split into 4 different tasks: a) metaphor identification, which consists in finding the words in a text being used metaphorically; b) metaphor interpretation, which groups metaphor paraphrasing, metaphor best fitting definition selection, and metaphoric implication selection; c) metaphor mapping, consisting in explicitly showing the correspondences between source and target domains, and d) metaphor generation.

Methods: CMP has gained attention in recent years [14], since the seminal work of Shutova [33]. Shutova described how in 2015 approaches towards metaphor processing were shifting from hand-coded knowledge to statistical modeling.

Then, after the arousal of transformer architectures [8], a renewed attention towards metaphors can be seen in the literature, and the usage of Pre-Trained Language Models (PTLMs) has become the state of the art for metaphor identification [1] and interpretation [34]. Distributional methodologies (e.g., with the usage of Language Models) rely on a discourse analysis approach, and while being promising to identify the metaphorical expressions used in a text or being able to superficially interpret them by paraphrasing them, they lack the ability to run a finer graded analysis and conceptualization of metaphors.

A more symbolic approach (e.g., through Linked Data and Knowledge Graphs) would contribute in structuring the different elements involved in metaphors and the relations between them. Structured representation of knowledge would provide models which enable inference discovery to uncover individual preferences for particular metaphors, the patterns used to create metaphors, and an understanding of how the different lexical entries in the text metaphorically relate to each other. Overall, using Semantic Web technologies would facilitate the prediction of the effect, usage, and preference for a particular metaphor. Efforts towards symbolic representations of metaphor have already been initiated [13, 12, 15], yet, they are too generic and should be adapted to cover the medical and individual-centered approach of this thesis.

Hybrid approaches mixing symbolic and distributional models, while promising, are very scarce and new. Song et al., [34] represent similes (a similar device to metaphor) as triples and extract them using LMs, but they only encode three entity types (source domain, target domain, and attribute) which is not enough to represent a metaphor.

Datasets: Two different kinds of data should be used in this thesis to meet the hybrid approach necessities. On the one hand, Structured data covering metaphoric knowledge [13, 12, 15], and, on the other hand, available datasets to train and test Language Models with metaphors, a summary of which can be found in Ge's et al., survey [14]. Among these, most focus on computational metaphor identification [18, 2, 25, 31] and leave mostly uncovered other tasks such as metaphor interpretation, mapping, and generation. Available datasets for computational metaphor processing are too general, they do not differentiate between different figurative (e.g., between analogies, metaphors, idioms, or metonymy), they usually cover a small number of examples, and they do not control socio-cultural or individual differences though being key aspects in metaphor comprehension, and, finally, and as a core limitation that motivates one of the goals of this thesis, they do not cover the medical domain.

Baselines: Among the different tasks in CMP, metaphor identification has obtained the most attention, thus, most baselines available are for this task [14], and results are already very promising with baselines around 80% of accuracy. Yet, this is only for English and little has been done for other languages. Moreover, among the different things which could be controlled for bias in the prediction of metaphoric expressions in the texts (e.g., different underlying metaphors, domains, or metaphor types), just part of speech is controlled. Regarding the rest of the tasks, research is not only scarce but also there is a lot of heterogeneity in the approaches taken. Baselines and evaluation criteria are defined particularly per paper, making comparability among them very hard.

2.2 Metaphors in the Medical Domain

Studies regarding metaphors in the medical domain have mostly been conducted by manually harvesting physicians' and patients' discourse to identify, classify and study the implications of using a limited group of metaphors in the medical domain [10, 28, 32]. Most studies have focused either on studying the usage of figurative expressions by non-neurotypical patients to gain insights into cognition or, on analyzing the discourse surrounding oncological processes. In this latter setting violence and journey metaphors have been the most discussed, leaving the rest unattended. It is relevant to extend this kind of research to other metaphors and gain a bigger adaptation ability to each patient [10, 32], their sociocultural background [11], treatment phase, and illness [26]. At the moment, large databases and computational tools to automate this process, to the best of our knowledge are not available, thus, conveying a key goal of this thesis. The only communication tool through metaphors currently available is the Metaphor Menu [32], an inventory of narratives exploring different metaphors that can be used to talk about cancer. A limitation of such a tool is that it covers only a small and closed number of metaphors, it is only openly available in English and it has not been thoroughly researched on how to be adapted to individual patients' necessities. An additional limitation of the Metaphor Menu is that it only covers the underlying metaphor in a text but it does not provide an structured inventory of the metaphoric mappings that can be selected to express a particular metaphor or to reframe an existing one through particular lexical expressions (e.g., if a nurse wants to express treatment in terms of the journey metaphor, a tool containing such correspondences, would relate the patient to the traveler, and the treatment to the path). Extension of tools such as the Metaphor Menu in this direction becomes relevant as researchers such as Landau et al., [20] have pointed out extending the same metaphor (through different lexical entries) throughout the text to talk about both the risks and prevention possibilities is more effective than changing metaphor across discourse.

3 Problem Statement and Contributions

Metaphoric expressions identification in texts, or what it has usually been called metaphor identification, has been the most research task in CMP. However, unless some further effort is conducted afterwords, the identified words, relations between them, and what they are implicitly trying to express, cannot be interpreted. Thus, metaphoric expressions identification needs to be complemented by the joint development of resources and research in related tasks such as metaphor interpretation, generation and mapping. The first contribution from this thesis is then, shifting towards a more holistic exploration of metaphors. The proposed pipeline towards fulfilling this kind of approach is represented in Figure 1 and described bellow.

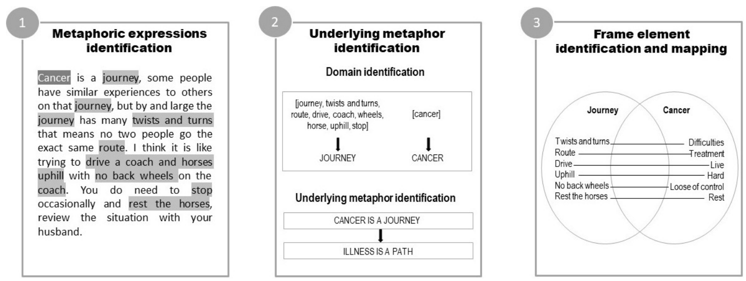

Fig. 1. Metaphor processing road map

1. **Metaphoric expressions identification**, traditionally coined *metaphor identification*. This task consists of identifying the words used metaphorically in a text.
2. **Underlying metaphor identification**, this next task derives from a rather cognitive approach to metaphor processing in comparison to the main discourse analysis approach[2]. It consists of, given a list of the lexical entries used metaphorically in the text, identifying the source and target domains of the metaphor.
3. **Frame elements' identification and mapping** consists of given the lexical entries used metaphorically in the discourse, the underlying metaphors that relate them, and their source and target domains, go a step further and uncover the semantic roles they play inside that the expressed domain in the text and the correspondent entities in the target domain.

The described pipeline makes use, on the one hand, of Knowledge Graphs to represent the metaphors and elements in them (e.g., domains, frame elements, lexical entries or mappings), and, on the other hand, of LMs for the identification of new metaphoric expressions in the text, prediction of relations between them, and further population of the ontology.

Right now, the main challenges that need to be addressed to make the execution of the pipeline possible are:

[2] Cognitive approach to metaphor complements discourse approach and provides cues to following the patterns of creation and understanding of metaphors as well as structured and abstracted examples which facilitate ontological representation of metaphors.

1. Finding the best way to encode not only metaphors in the ontology, but also the context and corpus characteristics they appear in. We hypothesize this could be done by adapting current available ontologies and linking them to resources as Ontolex-Frac [4] for corpus annotation.
2. Understanding how figurative knowledge is encoded in LMs and how to exploit them to further populate the metaphor Knowledge Graph.
3. There is a need of datasets that cover metaphors at a finer-graded level. For instance previous metaphoric expression identification datasets do not differenciate between different kinds of figurative knowledge, this is problematic as different kinds of figurative language should be represented differently in the graph. Moreover there is a lack of datasets for metaphor interpretation and mapping.
4. Benchmarks and Baselines for CMP should also be created to compare the obtained results.

This pipeline should be conducted separately for discourse produced by different subgroups of people, controlling socio-cultural variables, type of illness or affliction, and treatment phase factors. Once discourse has been processed, the outcome should consist of different subgraphs per population type. These graphs should be compared to try and grasp if any different patterns arise from them.

By pursuing these outcomes, this thesis will contribute to the technological field by providing further understanding and development of NLP and Ontological models to encode metaphorical knowledge, algorithms, and architectures mostly based on Language Models that enhance them, databases that can be further exploited, and, baselines for the newly created tasks. Moreover, the outcoming structured datasets will provide data that can be used to explore cognitive and sociocultural patterns behind metaphor processing.

4 Research Methodology and Approach

The first step of this thesis should cover the available **ontological model's adaptation**. A great ontology to start could be Framester [12] and derived resources from it such as ImageSchemaNet [15]. After its adaptation to our necessities, it should be linked to cognitive resources such as Small World of Words [7], factual knowledge bases such as Wikidata [37] and lexicographic resources such as Ontolex-Frac [4] and Ontolex-Core [24] modules.

The next step should cover the **dataset creation**. This dataset should consist on texts produced on the medical environment, annotated with metadatate controlling for the patient and clinician particular circumstances. A starting point, which could also serve to validate the ontological models creating in the first step, is the transformation of the Metaphor Menu [32] in structured data. Further data should be collected and parsed from questionnaires and internet forums, similarly to [32].

The third step consists of the **exploitation of Language Models** to perform the tasks described in section 3 while further populating the ontology.

At this point technical questions such as the following ones will be addressed, pushing the SotA in CMP forward: What figurative knowledge is encoded in Language Models? How? What kind of knowledge injection boosts Language Models performance in each of the CMP tasks described in section 3.

Our first approach covers only Pre-Trained Language Models, as they provide higher control and interpretability of the outputs. Working with Large Language Models and comparing them to our contributions, is currently out of the scope of this thesis and would remain as future work, that could be pursued after the delivery of the thesis.

Finally, the applicability of our outcomes and their effect in the medical domain will be validated in a real-world scenario by interviews in collaboration with the Horizon Europe 4DPicture project.

5 Evaluation Plan

The proposed **ontological model** for representing metaphors should follow the best practices described in the Linguistic Linked Data community [5].

Metaphoric expressions identification will be tested on general domain metaphor identification datasets and benchmarks as the ones described in [14].

Identification of the underlying metaphors behind texts and **metaphor mapping** can be evaluated by using as gold dataset the data extracted from ontological resources such as Framester [12], or datasets with metaphoric analogies encoding their domains such as the one from Czinczoll et al., [6]. Superficially it can be compared to similar works such as the one conducted by Song et al., [34]; yet, the work with analogies rather than metaphors. Given the lack of resources to compare the latter-named tasks with other works, we would consider at this point the creation of a benchmark for CMP.

The resulting communication tool will be validated in a **real-world scenario** by interviews with healthcare professionals, linguists, and patients. This will be possible through collaboration with the Horizon Europe 4DPicture project.

6 Preliminary Results

Our first experiments dealt with the metaphor identification task. In a first attempt to uncover how much linguistic knowledge is encoded in the available pre-trained language models to classify words between figurative and literal expressions. Even when following a very simple approach we obtained promising and competitive results when compared with the SotA approaches which support the model with external linguistic features and theories [1, 21, 38, 39].

Our approach is in line with research such as [36, 29] that goes towards the idea that PTLMs already encode a large amount of linguistic knowledge, and thus can be directly exploited, or with minimal fine-tuning to perform a wide range of language-related tasks.

For this first approach, we use minimal prompting to fine-tune and exploit RoBERTa [23] model. This prompting procedure derives from the idea that

PTLMs have been trained with a masked language model objective where given a correctly verbalize sentence they need to predict a masked word in it. Thus, providing the model with the instructions in a well-verbalized way, more similar to how the model was trained, is supposed to boost the model's performance, this has been coined as prompting [22]. In our first experiments, we fine-tune a PTLM with a sequence classification layer on the top. We train the model with the input prompt "`[SEP] sentence with target word [SEP] target word`" and the regarding label, 1, if the target word is a metaphor in the sentence, otherwise 0. The following Table 1 summarizes the obtained results, which even with such a simple methodology are results are not far behind the current SotA. In the future and by taking advantage of the metaphor modeling as structured data, we aim to inject the PTLMs with additional lexical features from BabelNet or Ontolex following the minimal prompting technique.

Table 1. F1 Score report for metaphor identification datasets

	Babieno, 22	Lin, 21	Yang, 21	Wan, 21	Ours
VUA-20	**72.5**	—	—	—	68.9
VUA-V	68.8	75.6	**80.7**	75.0	71.1
MOH-X	80.8	**84.7**	—	—	76.9
TroFi	61.7	74.5	—	**89.3**	73.1

Further analysis should make use of visualization techniques to check how the model is learning through different layers and epochs.

7 Conclusions

Explicit metaphor representation should aid communication in sensitive contexts such as the medical one. While previous tools such as the Metaphor Menu required huge manual efforts to produce them, they are still very case oriented, small, not flexible and with limitations. This thesis proposes a mixed exploitation of Natural Language Processing and Semantic Web technologies to further enrich them, while providing new extense databases and guidelines that will boost figurative knowledge computational processing.

Acknowledgements I would like to thank my supervisors Dr.Jorge Gracia del Río and Dr.Jorge Bernad Lusilla for their support and feedback, and Javier Fernández for his valuable insight and questions.

References

1. Babieno, M., Takeshita, M., Radisavljevic, D., Rzepka, R., Araki, K.: Miss roberta wilde: Metaphor identification using masked language model with wiktionary lexical definitions. Applied Sciences **12**(4) (2022)

2. Birke, J., Sarkar, A.: A clustering approach for nearly unsupervised recognition of nonliteral language. In: Conference of the European Chapter of the Association for Computational Linguistics (2006)
3. Casarett, D., Pickard, A., Fishman, J.M., Alexander, S.C., Arnold, R.M., Pollak, K.I., Tulsky, J.A.: Can metaphors and analogies improve communication with seriously ill patients? Journal of Palliative Medicine **13**(3), 255–260 (2010)
4. Chiarcos, C., Ionov, M., de Does, J., Depuydt, K., Khan, F., Stolk, S., Declerck, T., McCrae, J.P.: Modelling frequency and attestations for ontolex-lemon. In: Proceedings of the 2020 Globalex Workshop on Linked Lexicography. pp. 1–9 (2020)
5. Cimiano, P., Chiarcos, C., McCrae, J.P., Gracia, J., Cimiano, P., Chiarcos, C., McCrae, J.P., Gracia, J.: Linguistic linked open data cloud. Springer (2020)
6. Czinczoll, T., Yannakoudakis, H., Mishra, P., Shutova, E.: Scientific and creative analogies in pretrained language models. arXiv preprint arXiv:2211.15268 (2022)
7. De Deyne, S., Navarro, D.J., Perfors, A., Brysbaert, M., Storms, G.: The "small world of words" english word association norms for over 12,000 cue words. Behavior research methods **51**, 987–1006 (2019)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 NACL. ACL, Minneapolis, Minnesota (Jun 2019)
9. Fadul, N., Elsayem, A.F., Palmer, J.L., Fabbro, E.G.D., Swint, K., Li, Z., Poulter, V.A., Bruera, E.: Supportive versus palliative care: What's in a name? Cancer **115** (2009)
10. Fatehi, A., Table, B., Peck, S., Mackert, M., Ring, D.: Medical metaphors: Increasing clarity but at what cost? The Archives of Bone and Joint Surgery **10**(8), 721–728 (2022)
11. Fernandez, J.R., Richmond, J., Nápoles, A.M., Kruglanski, A.W., Forde, A.T.: Everyday discrimination and cancer metaphor preferences: The mediating effects of needs for personal significance and cognitive closure. SSM - Population Health **17** (2021)
12. Gangemi, A., Alam, M., Asprino, L., Presutti, V., Recupero, D.R.: Framester: A wide coverage linguistic linked data hub. In: International Conference Knowledge Engineering and Knowledge Management (2016)
13. Gangemi, A., Alam, M., Presutti, V.: Linked metaphors. In: International Workshop on the Semantic Web (2018)
14. Ge, M., Mao, R., Cambria, E.: A survey on computational metaphor processing techniques: From identification, interpretation, generation to application (2022)
15. Giorgis, S.D., Gangemi, A., Gromann, D.: Imageschemanet: A framester graph for embodied commonsense knowledge. Semantic Web (2022)
16. Gustafsson, A.W., Hommerberg, C., Sandgren, A.: Coping by metaphors: the versatile function of metaphors in blogs about living with advanced cancer. Medical Humanities **46**, 267 – 277 (2019)
17. Harrington, K.J.: The use of metaphor in discourse about cancer: a review of the literature. Clinical journal of oncology nursing **16 4**, 408–12 (2012)
18. Krennmayr, T., Steen, G.: VU Amsterdam Metaphor Corpus, pp. 1053–1071. Springer Netherlands, Dordrecht (2017)
19. Lakoff, G., Johnson, M.: Metaphors We Live By. The University of Chicago Press (1980)
20. Landau, M.J., Cameron, L.D., Arndt, J., Hamilton, W., Swanson, T.J., Bultmann, M.N.: Beneath the surface: Abstract construal mindset increases receptivity to metaphors in health communications. Social cognition **37 3**, 314–340 (2019)

21. Lin, Z., Ma, Q., Yan, J., Chen, J.: Cate: A contrastive pre-trained model for metaphor detection with semi-supervised learning. In: Proceedings of the 2021 EMNLP. pp. 3888–3898 (2021)
22. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G.: Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. ACM Computing Surveys **55**, 1–35 (9)
23. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
24. McCrae, J.P., Bosque-Gil, J., Gracia, J., Buitelaar, P., Cimiano, P.: The ontolex-lemon model: development and applications. In: Proceedings of eLex 2017 conference. pp. 19–21 (2017)
25. Mohammad, S.M., Shutova, E., Turney, P.D.: Metaphor as a medium for emotion: An empirical study. In: International Workshop on Semantic Evaluation (2016)
26. Munday, I., Newton-John, T.R.O., Kneebone, I.I.: Clinician experience of metaphor in chronic pain communication. Scandinavian Journal of Pain **23**, 88 – 96 (2022)
27. Penson, R.T., Schapira, L., Daniels, K.J., Chabner, B.A., Lynch, T.J.: Cancer as metaphor. The oncologist **9 6**, 708–16 (2004)
28. Pinheiro, A.P.M., Pocock, R.H., Dixon, M.D., Shaib, W.L., Ramalingam, S.S., Pentz, R.D.: Using metaphors to explain molecular testing to cancer patients. The oncologist **22 4**, 445–449 (2017)
29. Pitarch, L., Dranca, L., Bernad, J., Gracia, J.: Lexico-semantic relation classification with multilingual finetuning. In: LLOD Approaches for language data research and management. pp. 86–88 (2022)
30. Rai, S., Chakraverty, S.: A survey on computational metaphor processing. ACM Computing Surveys (CSUR) **53**(2), 1–37 (2020)
31. Sanchez-Bayona, E., Agerri, R.: Leveraging a new spanish corpus for multilingual and crosslingual metaphor detection. arXiv preprint arXiv:2210.10358 (2022)
32. Semino, E., Demjén, Z., Hardie, A., Payne, S., Rayson, P.: Metaphor, cancer and the end of life: A corpus-based study. Routledge (2017)
33. Shutova, E.: Design and evaluation of metaphor processing systems. Computational Linguistics **41**(4), 579–623 (2015)
34. Song, W., Guo, J., Fu, R., Liu, T., Liu, L.: A knowledge graph embedding approach for metaphor processing. IEEE/ACM Transactions on Audio, Speech, and Language Processing **29**, 406–420 (2021)
35. Thibodeau, P.H., Matlock, T., Flusberg, S.J.: The role of metaphor in communication and thought. Language and Linguistics Compass **13**(5), e12327
36. Ushio, A., Camacho-Collados, J., Schockaert, S.: Distilling relation embeddings from pretrained language models. pp. 9044–9062. Association for Computational Linguistics (2021)
37. Vrandeić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. Commun. ACM **57**, 78–85 (2014)
38. Wan, H., Lin, J., Du, J., Shen, D., Zhang, M.: Enhancing metaphor detection by gloss-based interpretations. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021. pp. 1971–1981 (2021)
39. Yang, L., Zeng, J., Li, S., Shen, Z., Sun, Y., Lin, H.: Metaphor recognition and analysis via data augmentation. In: Natural Language Processing and Chinese Computing: 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part I 10. pp. 746–757. Springer (2021)