

SCIENCE IS EXPLORATION: Computational Frontiers for Conceptual Metaphor Theory

Rebecca M. M. Hicke^{1,2,*}, Ross Deans Kristensen-McLachlan^{2,3}

¹Department of Computer Science, Cornell University, USA

²Center for Humanities Computing, Aarhus University, Denmark

³Department of Linguistics, Cognitive Science, and Semiotics, Aarhus University, Denmark

Abstract

Metaphors are everywhere. They appear extensively across all domains of natural language, from the most sophisticated poetry to seemingly dry academic prose. A significant body of research in the cognitive science of language argues for the existence of *conceptual metaphors*, the systematic structuring of one domain of experience in the language of another. Conceptual metaphors are not simply rhetorical flourishes but are crucial evidence of the role of analogical reasoning in human cognition. In this paper, we ask whether Large Language Models (LLMs) can accurately identify and explain the presence of such conceptual metaphors in natural language data. Using a novel prompting technique based on metaphor annotation guidelines, we demonstrate that LLMs are a promising tool for large-scale computational research on conceptual metaphors. Further, we show that LLMs are able to apply procedural guidelines designed for human annotators, displaying a surprising depth of linguistic knowledge.

Keywords

conceptual metaphor theory, large language models, pragglejazz

"Metaphors are much more tenacious than facts."

- Paul de Man [19]

1. Introduction

Metaphor is commonly understood as the description of one concept in the vocabulary of another, typically for poetic, rhetorical, or otherwise literary effect. However, some metaphors are so systematic and conventionalized that we barely recognize them as such. Consider, for example, the commonplace metaphor LIFE IS A JOURNEY. While this is certainly a metaphor in its own right, it can also be instantiated in ways which do not directly refer to either life or journeys: "We have *come a long way*"; "I'm *going through a rough patch*"; "She's at a *crossroads*." A metaphor like LIFE IS A JOURNEY is more than simple figurative language. Instead, it allows us to make sense of the process of a normal human life through the rich language of journeys and voyages. These kinds of systematic mappings are known as *conceptual metaphors* [16].

CHR 2024: Computational Humanities Research Conference, December 4–6, 2024, Aarhus, Denmark

*Corresponding author.

✉ rmh327@cornell.edu (R. M. M. Hicke); rdkm@cc.au.dk (R. D. Kristensen-McLachlan)

🌐 <https://rmatouschekh.github.io> (R. M. M. Hicke)

🆔 0009-0006-2074-8376 (R. M. M. Hicke); 0000-0001-8714-1911 (R. D. Kristensen-McLachlan)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



Conceptual metaphors are so deeply ingrained in our everyday language that the most frequent usage of a word or phrase may be metaphorical. This can make it difficult to recognize the presence of metaphorical language, even for linguists and other specialists [14]. Moreover, processing and understanding conceptual metaphors necessarily involves encyclopedic general knowledge of the world, as well as shared cultural norms and stereotypes between speaker/hearer or writer/reader. Identifying conceptual metaphors has therefore consistently proven to be a challenge for both natural language processing (NLP) and artificial intelligence (AI). Given that contemporary transformer-based models ‘learn’ by consuming huge amounts of natural language and build knowledge by looking at how words are commonly used, it seems that they should similarly struggle to recognize these metaphors.

Nevertheless, being able to consistently identify the presence of conceptual metaphors in discourse would be greatly beneficial for fields such as computational literary studies and cultural analytics more broadly. In this paper, we therefore set out to test empirically whether LLMs can be used productively in the context of conceptual metaphor theory (CMT). If the answer is no, it may be evidence that developing computational approaches to conceptual metaphor remains an intractable problem. If, on the other hand, LLMs *can* consistently identify the presence of conceptual metaphors, the question then arises of exactly how and where LLMs encode metaphoricality.

2. Related Work

Cognitive Linguistic Foundations

The study of CMT was inaugurated in the 1980s in Lakoff and Johnson [16]. Central to the argument in [16] is the fundamental role that metaphor plays in structuring everyday discourse. Rather than simply being a literary or rhetorical flourish, Lakoff and Johnson argue that conceptual metaphors structure our common understanding of entities, processes, or entire semantic domains. CMT played a crucial role in the development of non-generativist cognitive linguistics [7, 17, 18, 27, 28]. An important development came with the introduction of *conceptual blending* (otherwise known as *conceptual integration*), in which the kinds of analogical reasoning underlying conceptual metaphors are argued to be more fundamental to human language and cognition [11, 12, 8]. While there are dissenting voices [13], recent decades have seen a growing body of empirical research which claims to provide experimental evidence in favour of CMT [4, 6, 29].

Metaphor and Artificial Intelligence

The extensive role and function of metaphor in natural language has not escaped researchers in NLP. Indeed, a great deal of early work in AI concerned exactly how a computer could make sense of metaphorical language [5, 23, 24] (see [2] for a comprehensive overview). Much of this work relied on so-called symbolic AI, but struggled because GOF AI (good, old-fashioned AI) could not provide the analogous reasoning skills and extensive general knowledge required to process metaphorical language. In the context of conceptual blending specifically, there have been tentative attempts to model the underlying processes involved computationally [10].

1. Read the entire text-discourse to establish a general understanding of the meaning.
2. Determine the lexical units in the text-discourse.
3.
 - a. For each lexical unit in the text, establish its meaning in context, that is, how it applies to an entity, relation, or attribute in the situation evoked by the text (contextual meaning). Take into account what comes before and after the lexical unit.
 - b. For each lexical unit, determine if it has a more basic contemporary meaning in other contexts than the one in the given context. For our purposes, basic meanings tend to be
 - i. More concrete; what they evoke is easier to imagine, see, hear, feel, smell, and taste.
 - ii. Related to bodily action.
 - iii. More precise (as opposed to vague).
 - iv. Historically older.
 Basic meanings are not necessarily the most frequent meanings of the lexical unit.
 - c. If the lexical unit has a more basic current-contemporary meaning in other contexts than the given context, decide whether the contextual meaning contrasts with the basic meaning but can be understood in comparison with it.
4. If yes, mark the lexical unit as metaphorical.

Figure 1: The metaphor identification procedure (MIP) introduced by the Pragglejaz Group in [14].

However, the encyclopedic world knowledge required in CMT is challenging to encode since the domain is essentially unbounded. Moreover, a core issue at the heart of computational metaphor analysis has historically been the lack of a shared definition of what is being studied, along with the scarcity of robust data and evaluation strategies [26].

Recent Developments

However, recent years have again seen a growing interest in quantitative and computational approaches to metaphor, primarily in more applied fields such as corpus linguistics and stylistics [9, 15, 25]. With the advent of neural language models, research in NLP has increasingly attempted to discover if (and how) word embeddings capture aspects of metaphoricity [20, 30, 21, 22]. More recently, work has been done on explicitly probing pre-trained language models to inspect how metaphorical knowledge is encoded [1] and creating shared task datasets [30]. In the context of LLMs specifically, GPT-3 has been shown to perform reasonably well at predicting source domains for English metaphors, although it performs significantly worse for Spanish [31]. These are promising developments and point to the increasing awareness of the practical and theoretical challenges underlying the nature of figurative language in NLP. To date, though, we are unaware of existing research which seeks to operationalize existing metaphor identification procedures for LLMs.

For each **word** in the **following** text, determine if it has a ***more*** basic contemporary meaning in other contexts than the one in the given context. For our purposes, basic meanings tend to be:

- more concrete; what they evoke is easier to imagine, see, here, feel, smell, and taste
- related to bodily action
- more precise (as opposed to vague)
- historically older

Basic meanings are not necessarily the most frequent meanings of the **word**. Remember to only report YES for words whose use in this context is not their most basic meaning.

Examples:

"I demolished his argument."

- I: NO (pronoun, basic meaning is the same)
- demolished: YES (verb, more basic meaning refers to destroying a building)
- his: NO (possessive pronoun, basic meaning is the same)
- argument: NO (noun, basic meaning is the same)

"They write about their family."

- They: NO (pronoun, basic meaning is the same)
- write: NO (verb, basic meaning is the same)
- about: NO (preposition, basic meaning is the same)
- their: NO (possessive pronoun, basic meaning is the same)
- family: NO (noun, basic meaning is the same)

Figure 2: The modified prompt used in all experiments. All significant (non-formatting) changes from Step 3b of the MIP procedure are highlighted in red. The sentence to be analyzed is appended in quotations to the bottom of the prompt when querying models.

3. Methods

We are interested in studying whether LLMs are capable of leveraging the contextual and cultural knowledge necessary to perform difficult metaphor identification tasks, such as recognizing conceptual metaphors. To do this, we operationalize the Metaphor Identification Procedure (MIP), a set of annotation guidelines introduced by the Pragglejaz Group in 2007 [14]. The steps of MIP are detailed in Figure 1. MIP and its variants are among the most commonly used procedures for metaphor annotation in work in corpus linguistics and stylistics; we therefore take it as a reasonable proxy for the metaphor identification process in humans.

Examining MIP, the first few steps have clear parallels in transformer-based LLMs. Step 1 is performed by the attention mechanism. Creating the representation of each word, in traditional attention mechanisms, requires ‘looking’ at every other word in the text, thus allowing the model to “establish a general understanding of [its] meaning.” Step 2, determining the lexical units in a text, is directly paralleled by tokenization. Finally, Step 3a is achieved via contextual embeddings, whose explicit purpose is to represent the specific meaning of a word in context. However, Step 3b of MIP has no clear counterpart in transformer-based models. This is because the ‘basic meaning’ of a lexical unit, as defined by MIP, is not obviously stored or represented within these models.

We want to know if state-of-the-art LLMs are capable of replicating Step 3b of the MIP

procedure despite these barriers, so we transform it into a prompt for instruction-tuned models (Figure 2). In creating the prompt, we leave the text of Step 3b almost entirely intact. “Lexical units” are replaced by “words” in order to encourage consistent outputs. Emphasis is added in several locations to only annotating words which have a *more* basic meaning than their usage in the input text. Finally, two examples are appended to the prompt. These examples are meant to facilitate in-context learning and provide a structure for the models to mimic in their outputs. They instruct the model to provide a parenthetical expression next to each word consisting of the word’s part of speech and its more basic meaning, if one exists. These expressions were introduced by the gpt-4o model itself during initial exploration and were then included in the prompt because they force the model to provide explanations for its annotations. This allows for further evaluation of the outputs and ultimately means they play a similar role to Chain of Thought prompts, which have been shown to improve model performance across a variety of tasks [32].

We use this prompt to query three models from OpenAI’s GPT family, chosen because of their size, popularity, and high performance on a wide range of tasks. Specifically, we look at gpt-3.5-turbo,¹ gpt-4-turbo,² and gpt-4o.³ We access each model using OpenAI’s Chat Completion API. A system prompt (“*You are a helpful assistant. You have extensive linguistic knowledge.*”) is provided each time the model is queried. Then, the main prompt (Figure 2) is passed as a user message, with the current text of interest appended in quotations at the end. All parameters except nuclear sampling (top_p) are left at their default values during prompting for all models. We set top_p to 0.1, which means that the models only select from the tokens which make up the top 10% of the probability mass when choosing each next token. We adjust the top_p value because we are not prompting for diverse or creative text, but instead for reliable and accurate outputs, which we hypothesize will be more likely produced from high-probability tokens. The cost of all experiments using the OpenAI API was ~\$100.

4. Data

We evaluate each model on two datasets. First, we wish to study the models’ basic ability to use MIP to correctly identify metaphorical words. To this end, we use the Trope Finder (TroFi) dataset [3], which consists of sentences drawn from Wall Street Journal issues published between 1987 and 1989. Each sentence contains one of a list of fifty words whose usage is annotated as literal or non-literal. We only consider sentences from the dataset that have been labeled by humans and take their annotations as ground truth, excluding all sentences labeled only by the TroFi clustering system. This leaves us with 3,736 sentences. Each model is then prompted as described above for each sentence and evaluated only on the words annotated in TroFi. A response is marked as correct if the model says “YES” there is a more basic meaning for the word of interest and the label from TroFi is “literal” and vice versa. Occasionally, the models provide labels for a phrase in the input sentence instead of individual words; the label provided for the entire phrase is then applied to the word of interest if necessary. If a model

¹gpt-3.5-turbo-0125

²gpt-4-turbo-2024-04-09

³gpt-4o-2024-05-13

Table 1

The precision and recall values for the identification of literal and metaphorical word usage in the TroFi dataset by all models.

	Metaphorical		Literal	
	Precision	Recall	Precision	Recall
3.5-turbo	58.30	97.90	66.42	5.59
4-turbo	74.80	86.90	77.41	60.53
4o	73.40	93.66	83.69	54.24

does not provide a label for a word, we act as if the incorrect label was given.

However, the TroFi dataset does not permit for a very fine-grained analysis of the models' capabilities. In order to facilitate such an analysis, we create a new dataset consisting of example sentences pulled from Lakoff and Johnson's original work on CMT [16]. We collect all full sentence examples from the text that we believe can be understood as metaphorical without context, leaving us with 544 sentences, which we refer to as the *MWLB dataset*. The dataset is made publicly available, along with all model outputs and human annotations.⁴

Each model is again prompted to provide word-level annotations for every sentence. Then, we perform detailed qualitative analysis on each model's output for a subset of 100 sentences. Specifically, we evaluate each output on five binary categories:

- **L&J Metaphor(s) – Identified:** Whether the model has correctly identified *all* metaphors highlighted by Lakoff and Johnson. For metaphorical phrases, the model's response is considered correct if it has identified at least one key word as having a more basic meaning.
- **L&J Metaphor(s) – Correct Basic Meanings:** Whether the model has provided a correct, or plausible, more basic meaning for *all* metaphorical words identified in the category above. A label of 1 is applied only if all metaphors from Lakoff and Johnson have been correctly identified.
- **Additional Annotations:** Whether the model has labeled any words not highlighted by Lakoff and Johnson as having a more basic meaning.
- **Additional Annotations – Metaphorical:** Whether *all* the additional words labeled by the model are plausibly metaphorical or may have a more basic meaning.
- **Additional Annotations – Correct Basic Meanings:** Whether the model has provided a correct, or plausible, basic meaning for *all* additional labeled words. It is possible for models to give the correct basic meaning for a word even if it is not plausibly metaphorical.

All annotations are performed by one of the authors, using MIP as a guideline. Confusing or difficult annotations are then discussed between both authors till a consensus is reached. These annotations are necessarily subjective, and we do not claim that they represent a ground truth. Instead, since these evaluations are informed by the same procedure used to create the

⁴<https://github.com/rmatouschekh/science-is-exploration>




Figure 3: Confusion matrices for results on the TroFi dataset for each model. The rows represent the true values for each sample, and the columns represent the model labels.

outputs, we claim that they allow us to study the *plausibility* of the models’ performance, if not strictly their *accuracy*. In future work, further evaluation of model outputs with multiple trained annotators may allow for a more concrete analysis of model performance.

5. Results and Discussion

On the TroFi dataset, all models demonstrate an ability to distinguish between literal and non-literal word usage with generally better than chance performance (Table 1). Thus, they seem to be able to apply the MIP procedure for identifying metaphors and determine computationally whether a word has a more basic meaning. This is surprising since, as discussed above, basic word meanings are not necessarily the most frequent. Nothing about the language modeling approach to pre-training models guarantees that a model would be able to differentiate between the most frequent and most basic meaning of a word.

We also find, however, that the models are prone to over-labeling words as metaphorical (Figure 3). Whereas all three models achieve high recall for metaphor identification, they frequently label words used literally as having a more basic meaning, leading to low recall for the ‘literal’ category and low precision for the ‘metaphorical’ category. 3.5-turbo particularly struggles with accurately identifying literal word usage. It is worth noting that, although there are clear cases where the models have failed to properly apply MIP, there may also be a disconnect between annotations in the TroFi corpus and what is considered metaphorical by the MIP procedure, as MIP was not used for annotation in TroFi.

For the annotated subset of the MWLB dataset, all the models found all the Lakoff and Johnson metaphors in over 60% of sentences and further provided a correct basic meaning for the metaphorical words in over 50% of sentences (Figure 4). 4o performed the best on this task, achieving a remarkable performance given the subtlety and complexity of many of the metaphors and the cultural knowledge required to identify both their metaphoricity and their basic meanings. 4o’s high performance may be due in part to the fact that it was most likely to label words as metaphorical overall, having annotated additional words in the largest percent-




Figure 4: The percentage of samples from Lakoff and Johnson where all metaphors and basic meanings are correctly identified.

age of sentences (Figure 5). In only 48% of sentences did these additional annotations identify only plausibly metaphorical words. Nonetheless, the inclusion of correct basic meanings for such a large proportion of the Lakoff and Johnson metaphors suggests that MIP was overall being correctly applied and true ‘knowledge’ was being demonstrated, not just random chance.

3.5-turbo struggled the most at replicating the desired output structure, occasionally excluding words or truncating sentences, inconsistently formatting lists, and sometimes dropping the parenthetical explanations. This affected its accuracy, as the model sometimes failed to provide annotations for words central to the metaphors of interest. Additionally, 4-turbo was least likely to annotate a word as having a more basic meaning (Figure 5). This meant it identified fewer Lakoff and Johnson metaphors, but a greater proportion of the additional metaphors it annotated were plausible. The basic meanings provided by 4-turbo were usually of high quality, and tended to be more accurate and specific than those produced by 3.5-turbo. Despite the benefits, however, this model’s ‘caution’ led it to perform worse than 4o overall.

Each model struggled to correctly identify when smaller function words were used metaphorically, particularly prepositions like ‘in.’ This made the models worse at identifying so-called container metaphors, such as LIFE IS A CONTAINER and ACTIVITIES ARE CONTAINERS. For example, in the sentence “That’s *in* the *center* of my *field* of vision.”, labeled by Lakoff and Johnson as VISUAL FIELDS ARE CONTAINERS, the word ‘in’ was overlooked as metaphorical by all three models.

The models were also challenged by metaphors in which an entity is being treated as a different kind of object, like instances of personification, PLACE FOR INSTITUTION, or PRODUCER FOR PRODUCT metaphors. However, the GPT-4 models had a much greater ability to detect and provide accurate basic meanings for these. For example, in “Let’s not let Thailand become another *Vietnam*.”, both 4-turbo and 4o correctly identified ‘Vietnam’ as a metaphor. 4o explained that the more basic meaning of the word “refers to the country, whereas in this context it refers to a situation similar to the Vietnam War” and 4-turbo provides a similar response. Likewise, in “I hate to read *Heidegger*.”, 4o recognized that ‘Heidegger’ is being used metaphorically and stated that the “more basic meaning refers to the person Martin Heidegger, a German philosopher, rather than his works.” Identifying and explaining both of these metaphors requires a nuanced understanding of both the semantics of the sentence and the cultural context surrounding them.




Figure 5: The percentage of samples from Lakoff and Johnson in which additional metaphors have been labeled. Smaller bars are included which represent the percentage of additional words for which the correct basic meaning is provided (upper bar) and which are plausible metaphors (lower bar).

The models cannot always perform this analysis (all three miss that ‘the Alamo’ is metaphorical in “Remember *the Alamo!*”), but it is impressive that they are ever able to do so.

In addition, forcing the models to annotate word-by-word makes it challenging for them to identify metaphors comprising multi-word units. For example, it is difficult to determine which words should be marked as having more basic meanings in the sentence “*Get the most out of life.*” Analyzing texts word-by-word can also make identifying and providing the basic meanings of metaphorical compound words more complex. For example, both ‘underage’ and ‘brainchild’ cause problems for 3.5-turbo; it does not recognize that the *under* of underage is metaphorical and it says that the more basic meaning of brainchild is “a child conceived in the mind.” In contrast, 4-turbo says that the “more basic meanings of ‘brain’ and ‘child’ are more concrete and related to physical objects or beings” and notes that the “more basic meaning of ‘under’ is physically beneath something” for underage. 4o recognizes both words as metaphorical, but provides worse basic meanings.

While the models are clearly fallible, they nevertheless demonstrate an impressive ability to synthesize semantic, syntactic, and cultural information. They are frequently able to recognize when a word is being used metaphorically and often provide a correct basic meaning for the word. This ability sometimes holds even for nuanced and complex metaphors.

6. Conclusion

We find that large, generative LMs are capable of applying the classic metaphor annotation procedure, MIP. In doing so, they demonstrate an ability to discern the “basic meaning” of words and thus a depth of linguistic understanding that is not obviously gleaned from language modeling pre-training. Notably, they also demonstrate that LLMs are able to execute linguistic procedures designed for human annotators. This capacity means that generative LLMs are a promising tool for large-scale computational research on conceptual metaphors, which has previously been largely infeasible. In addition, it suggests that further research on how

metaphoricity is learned by models may provide insight into their ability to acquire complex linguistic knowledge that often relies on people’s embodied experiences. These findings also suggest several avenues for future research, including studies into where information about words’ basic meanings or metaphoricity is stored by models, an exploration of open source models ability to annotate for conceptual metaphor, and research on operationalizing other linguistic procedures for execution by LLMs.

Acknowledgments

We would like to thank Lavinia Cerioni for her help in brainstorming this project and our colleagues at the Center for Humanities Computing for their support and advice. Additionally, part of the computation done for this project was performed on the UCloud interactive HPC system, which is managed by the eScience Center at the University of Southern Denmark.

References

- [1] E. Aghazadeh, M. Fayyaz, and Y. Yaghoobzadeh. “Metaphors in Pre-Trained Language Models: Probing and Generalization Across Datasets and Languages”. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, 2022, pp. 2037–2050. DOI: 10.18653/v1/2022.acl-long.144.
- [2] J. A. Barnden. “Metaphor and Artificial Intelligence: Why They Matter to Each Other”. In: *The Cambridge Handbook of Metaphor and Thought*. Cambridge University Press, 2008, pp. 311–338. DOI: 10.1017/cbo9780511816802.020.
- [3] J. Birke and A. Sarkar. “A Clustering Approach for Nearly Unsupervised Recognition of Nonliteral Language”. In: *11th Conference of the European Chapter of the Association for Computational Linguistics*. Trento, Italy: Association for Computational Linguistics, 2006, pp. 329–336.
- [4] L. Boroditsky. “Metaphoric Structuring: Understanding Time Through Spatial Metaphors”. In: *Cognition* 75.1 (2000), pp. 1–28. DOI: 10.1016/s0010-0277(99)00073-6.
- [5] J. G. Carbonell. “Metaphor: An Inescapable Phenomenon in Natural Language Comprehension”. In: *Strategies for Natural Language Processing*. Pittsburgh, PA: Psychology Press, 1982, pp. 415–434. DOI: 10.4324/9781315802671.
- [6] D. Casasanto and L. Boroditsky. “Time in the Mind: Using Space to Think About Time”. In: *Cognition* 106.2 (2008), pp. 579–593. DOI: 10.1016/j.cognition.2007.03.004.
- [7] W. Croft and D. A. Cruse. *Cognitive Linguistics*. Cambridge, UK: Cambridge University Press, 2004. DOI: 10.1017/cbo9780511803864.
- [8] B. Dancygier and E. Sweetser. *Mental Spaces in Grammar: Conditional Constructions*. Cambridge Studies in Linguistics. Cambridge, UK: Cambridge University Press, 2005. DOI: 10.1017/cbo9780511486760.

- [9] A. Deignan and E. Semino. “Corpus Techniques for Metaphor Analysis”. In: *Metaphor Analysis*. London, UK, 2010, pp. 161–179.
- [10] M. Eppe, E. Maclean, R. Confalonieri, O. Kutz, M. Schorlemmer, E. Plaza, and K.-U. Kühnberger. “A Computational Framework for Conceptual Blending”. In: *Artificial Intelligence* 256 (2018), pp. 105–129. DOI: 10.1016/j.artint.2017.11.005.
- [11] G. Fauconnier. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, UK: Cambridge University Press, 1994. DOI: 10.1017/cbo9780511624582.
- [12] G. Fauconnier and M. Turner. *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities*. New York City, NY: Basic Books, 2002.
- [13] R. W. Gibbs. “Why Do Some People Dislike Conceptual Metaphor Theory?” In: *Cognitive Semiotics* 5.1 (2009), pp. 14–36. DOI: 10.1515/cogsem.2013.5.12.14.
- [14] P. Group. “MIP: A Method for Identifying Metaphorically Used Words in Discourse”. In: *Metaphor and Symbol* 22.1 (2007), pp. 1–39. DOI: 10.1207/s15327868ms2201_1.
- [15] T. Krennmayr. “What Corpus Linguistics Can Tell Us About Metaphor Use In Newspaper Texts”. In: *Journalism Studies* 16.4 (2015), pp. 530–546. DOI: 10.1080/1461670x.2014.937155.
- [16] G. Lakoff and M. Johnson. *Metaphors We Live By*. Chicago: The University of Chicago Press, 1980. DOI: 10.7208/chicago/9780226470993.001.0001.
- [17] R. W. Langacker. *Foundations of Cognitive Grammar: Volume I: Theoretical Prerequisites*. Stanford, CA: Stanford University Press, 1987.
- [18] R. W. Langacker. *Foundations of Cognitive Grammar: Volume II: Descriptive Application*. Stanford, CA: Stanford University Press, 1991. DOI: 10.1515/9780804764469.
- [19] P. de Man. “Semiotics and Rhetoric”. In: *Diacritics* 3.3 (1973), pp. 27–33. DOI: 10.2307/464524.
- [20] R. Mao, C. Lin, and F. Guerin. “Word Embedding and WordNet Based Metaphor Identification and Interpretation”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, 2018, pp. 1222–1231. DOI: 10.18653/v1/p18-1113.
- [21] P. V. Panicheva, I. D. Mamaev, and T. A. Litvinova. “Towards Automatic Conceptual Metaphor Detection for Psychological Tasks”. In: *Information Processing & Management* 60.2 (2023), pp. 1–17. DOI: 10.1016/j.ipm.2022.103191.
- [22] M. Ptiček and J. Dobša. “Methods of Annotating and Identifying Metaphors in the Field of Natural Language Processing”. In: *Future Internet* 15.6 (2023), p. 201. DOI: 10.3390/fi15060201.
- [23] S. W. Russell. “Computer Understanding of Metaphorically Used Verbs”. In: *American Journal of Computational Linguistics* 3.3 (1976), pp. 1–74.
- [24] S. W. Russell. “Information and Experience in Metaphor: A Perspective From Computer Analysis”. In: *Metaphor and Symbolic Activity* 1.4 (1986), pp. 227–270. DOI: 10.1207/s15327868ms0104_1.

- [25] E. Semino, Z. Demjén, A. Hardie, S. Payne, and P. Rayson. *Metaphor, Cancer and the End of Life: A Corpus-Based Study*. Vol. 19. Routledge Advances in Corpus Linguistics. Oxford, UK: Routledge, 2018. DOI: 10.4324/9781315629834.
- [26] E. Shutova. “Design and Evaluation of Metaphor Processing Systems”. In: *Computational Linguistics* 41.4 (2015), pp. 579–623. DOI: 10.1162/coli_a_00233.
- [27] L. Talmy. *Toward a Cognitive Semantics, Volume 1: Concept Structuring Systems*. Cambridge, MA: The MIT Press, 2000. DOI: 10.7551/mitpress/6847.001.0001.
- [28] L. Talmy. *Toward a Cognitive Semantics, Volume 2: Typology and Process in Concept Structuring*. Cambridge, MA: The MIT Press, 2000. DOI: 10.7551/mitpress/6848.001.0001.
- [29] P. H. Thibodeau, R. K. Hendricks, and L. Boroditsky. “How Linguistic Metaphor Scaffolds Reasoning”. In: *Trends in Cognitive Sciences* 21.11 (2017), pp. 852–863. DOI: 10.1016/j.tics.2017.07.001.
- [30] X. Tong, E. Shutova, and M. Lewis. “Recent Advances in Neural Metaphor Processing: A Linguistic, Cognitive and Social Perspective”. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, 2021, pp. 4673–4686. DOI: 10.18653/v1/2021.naacl-main.372.
- [31] L. Wachowiak and D. Gromann. “Does GPT-3 Grasp Metaphors? Identifying Metaphor Mappings with Generative Language Models”. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada: Association for Computational Linguistics, 2023, pp. 1018–1032. DOI: 10.18653/v1/2023.acl-long.58.
- [32] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou. “Chain-of-Thought Prompting Elicits Reasoning in Large Language Models”. In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. New Orleans, LA: Curran Associates Inc., 2024, pp. 1–14. DOI: 10.48550/arXiv.2201.11903.