

What *Drives* the Use of Metaphorical Language? Negative Insights from Abstractness, Affect, Discourse Coherence and Contextualized Word Representations

Prisca Piccirilli and Sabine Schulte im Walde

Institute for Natural Language Processing, University of Stuttgart, Germany
{prisca.piccirilli, schulte}@ims.uni-stuttgart.de

Abstract

Given a specific discourse, which discourse properties trigger the use of metaphorical language, rather than using literal alternatives? For example, what drives people to say *grasp the meaning* rather than *understand the meaning* within a specific context? Many NLP approaches to metaphorical language rely on cognitive and (psycho-)linguistic insights and have successfully defined models of discourse coherence, abstractness and affect. In this work, we build five simple models relying on established cognitive and linguistic properties – frequency, abstractness, affect, discourse coherence and contextualized word representations – to predict the use of a metaphorical vs. synonymous literal expression in context. By comparing the models' outputs to human judgments, our study indicates that our selected properties are not sufficient to systematically explain metaphorical vs. literal language choices.

1 Introduction

Metaphors are "not just nice", but represent a "necessary" element of everyday thought and communication (Ortony, 1975; Lakoff and Johnson, 1980; van den Broek, 1981; Schäffner, 2004, i.a.), and are ubiquitous in natural language text corpora (Gedigian et al., 2006; Shutova and Teufel, 2010; Steen et al., 2010, i.a.). From the perspective of natural language processing (NLP), automatic approaches to metaphor processing are therefore important for any task that requires natural language understanding, and NLP has been concerned with the detection (Köper and Schulte im Walde, 2016; Alnafesah et al., 2020; Ehren et al., 2020; Dankers et al., 2020, i.a.), the interpretation (Shutova, 2010; Bizzoni and Lappin, 2018; Mao et al., 2018, i.a.) and, the generation (Stowe et al., 2021; Zhou et al., 2021) of metaphors.¹

¹See Tong et al. (2021) for a systematic, comprehensive review and discussion of the most recent metaphor processing systems and datasets.

As to our knowledge, however, no study so far has raised the question of WHY a metaphorical expression is used within a specific discourse, rather than an equally plausible literal alternative. For example, consider the discourse in table 1, where both the metaphorical expression *grasp the meaning* and its synonymous literal alternative *understand the meaning* seem equally acceptable (Piccirilli and Schulte im Walde, 2021, 2022). What are the factors priming for one use or the other? Are there cues within the discourse which influence the selection of one usage over the other? To which extent can computational approaches based on discourse properties model human behavior regarding the choice between synonymous metaphorical vs. literal language usage?

According to psycholinguistics and computational linguistics research, the processing of words is a function of their **frequency of occurrence** in the language (van Jaarsveld and Rattink, 1988; Wittmann et al., 2017, i.a.). Is the choice between a metaphorical and a literal expression therefore just an effect of frequency? In contrast, conceptual metaphor theory establishes metaphorical language as a figurative device for transferring knowledge from a concrete domain to a more abstract domain (Lakoff and Johnson, 1980), and the hypothesis that metaphorical usages correlate with the **abstractness of the context** has been supported in numerous NLP studies on automatic metaphor identification (Turney et al., 2011; Tsvetkov et al., 2013; Köper and Schulte im Walde, 2016; Alnafesah et al., 2020; Hall Maudslay et al., 2020). **Affect** has also been explored with regard to metaphoricity. Not only has metaphorical language been found to carry a stronger emotional load than literal language (Blanchette and Dunbar, 2001; Crawford, 2009), but metaphorical words and sentences are also judged to be "more emotionally engaging" than their synonymous literal paraphrases (Citron and Goldberg, 2014; Mohammad et al., 2016), and

This wasn't just a play on words, rather it was a demand that they should 'maintain a consistency between their words and their actions'. But I agree, that still does not absolve them from the need to speak truth to power. In our times when people spend so much time with TV and the internet, do they have the interest and time to read poetry? Many people believe that it is difficult to read poetry. Can everyone [*grasp* / *understand the meaning*] of a good poem, or is a skill necessary?

Table 1: Example of a discourse from the dataset introduced by Piccirilli and Schulte im Walde (2021). Both the literal expression *understand meaning* and its metaphorical counterpart *grasp meaning* seem equally acceptable in the discourse.

informing NLP models with **emotion** features has proven useful for metaphor detection (Gargett and Barnden, 2015; Köper and Schulte im Walde, 2018; Dankers et al., 2019). When analyzing metaphorical discourse features (Glucksberg, 1989; Steen, 2004) and the interactions of **discourse coherence** with the contextual salience of metaphorical vs. literal expressions (Inhoff et al., 1984; Gibbs, 1989; Giora, 1997; Giora and Fein, 1999; Gibbs, 2002; Kövecses, 2009), findings are directly connected with the theory of discourse cohesion, i.e., the principle that discourse should be a "group of collocated, structured, and coherent sentences" (Halliday and Hasan, 1976; Jurafsky and Martin, 2019). Models relying on the coherence of lexical semantic discourse structures have therefore been successful in identifying metaphors (Sporleder and Li, 2009; Bogdanova, 2010; Mesgar and Strube, 2016; Dankers et al., 2020). From a yet different perspective, Transformer-based pretrained language models (T-PLMs) (Vaswani et al., 2017), pre-trained on the language modeling task, are able to predict masked items in context and produce **contextual embeddings** accounting for both left and right contexts, and resulting in word representations that are dynamically informed by the surrounding words.

The rich previous interdisciplinary research on figurative language seems to agree on tight interactions between metaphorical language detection and properties of the respective discourses, i.e., **cognitive aspects** (abstractness and affect), **discourse coherence** and **contextual properties**. Do these discourse properties indeed trigger the use of a metaphorical vs. its synonymous counterpart? We address this question by exploring five simple discourse-based models inspired by the above-mentioned properties, namely frequency, abstractness, affect, discourse coherence and contextualized word representations. We approach the question within two studies. First, we explore discourse features of metaphorical usages occurring in natural language by applying the cognitive and linguistically-inspired models to existing discourses from the English corpus *ukWaC* (Baroni

et al., 2009). Then, we zoom into the models' predictions and evaluate them against human preferences (i.e., annotations) for metaphorical vs. literal language within the same discourses. By modeling the prediction for synonymous metaphorical vs. literal expressions motivated by the above discourse perspectives, we gain insight into human behavior for metaphorical vs. literal language choices.

2 Dataset

Research on figurative language has produced impressive resources on the metaphoricity of lexical items. However, these resources have limitations regarding the specific task we are addressing in this work: (i) the context in which words are metaphorically used is not large enough, providing human judgments only on the *word-level* (Steen et al., 2010) or on the *sentence-level* (Stefanowitsch, 2008; Shutova and Teufel, 2010; Mohammad et al., 2016), (ii) the target words or expressions are *ambiguous*, i.e., they may have both a metaphorical and a literal sense (Tsvetkov et al., 2013; Mohler et al., 2016), (iii) they are *extended* metaphors (Gibbs, 2006; Martin, 2008), (iv) the paraphrases to metaphorically-used words are automatically generated and no manual annotations were provided to evaluate the outputs (Bollegala and Shutova, 2013; Bizzoni and Lappin, 2018, i.a.).

We therefore use our recently released dataset specifically designed to investigate the choice of metaphorical vs. literal expressions in context (Piccirilli and Schulte im Walde, 2021, 2022). It contains a total of 1,000 discourses of five to six sentences (98 words on average), in which the final sentence of each discourse contains either a metaphorical expression or its literal alternative from a pair of synonymous subject-verb (SV) or verb-object (VO) expressions. Table 1 presents an example for the VO pair *grasp/understand meaning*. The overall 50 pairs of English expressions were selected from Shutova (2010) and Mohammad et al. (2016), and for each of the pairs, we extracted 20 discourses from the *ukWaC* (Baroni et al., 2009), 10 of which containing the metaphorical usage (e.g., *grasp*) and

10 of which containing the literal paraphrase (e.g., *understand*). To gain insight into human preferences between the selected pairs of metaphorical vs. literal expressions, we also collected crowdsourced human judgments for these 1,000 discourses, asking annotators to choose which expression they favored given the preceding discourse. This dataset is therefore optimal for the task at hand, as it provides (i) synonymous metaphorical vs. literal expressions, (ii) at the *discourse-level*, (iii) manually annotated. In the present work, we make use of the 1,000 discourses containing the original expressions in the *ukWaC* and the annotators’ choices for metaphorical vs. literal usages within a subset of 287 discourses, where 70% or more annotators agreed on the preference (metaphorical or literal).

3 Models and Experimental Setup

3.1 Prediction models

We approach the task of predicting the use of metaphorical vs. literal expressions as a *prompting task*: given an input prompt², the models predict whether the missing span should be the metaphorical or the literal expression. We apply five models relying on discourse properties, where each model approaches the task from a different perspective, to give us insight on which discourse features influence the metaphorical vs. literal selection.

Frequency approach When given the choice between a metaphorical expression and its synonymous counterpart, do we tend to favor the most frequent usage? **Baseline (Freq.)**: Our baseline relies on the occurrences of the SV and OV tuples in the original *ukWaC* corpus: the model receives the prefix prompt as input, and always outputs the most frequent expression of the pair.

Cognitively-inspired approaches The cognitive interaction between abstractness/affect and metaphorical language raises the question: to which extent (i) a more abstract discourse and (ii) a more emotionally-loaded discourse favor a metaphorical usage? **Abstractness (Abstr.)**: We measure the abstractness of a discourse preceding the target expression within four settings, based on the norms from Brysbaert et al. (2014). We assign abstractness scores to *all* words (Abstr.all), only *nouns* (Abstr.n), only *verbs* (Abstr.v) or only *adjectives* (Abstr.adj). We then obtain an overall

²The input prompt is a *prefix prompt* for all models except for BERT, whose input is a *cloze prompt*.

rating of abstractness for each discourse by computing the median of the respective lexical items’ abstractness scores. For each setting, we use as threshold the abstractness median of the respective part-of-speech class, and the model predicts the metaphorical expression if the overall abstractness score of the discourse is below that threshold (i.e., more abstract/less concrete), and the literal counterpart if above (i.e., less abstract/more concrete). **Emotionality (Emo.)**: We build a model predicting the target expression based on the emotionality of the preceding discourse, which we represent using the English emotion lexicon from Buechel et al. (2020). Each lexical item from the preceding discourse is assigned an emotionality score, and the median represents the overall emotionality score of the discourse. Appendix A.1 provides details on the emotionality score.

Discourse coherence approach Is the choice of an expression from a synonymous pair driven by the semantic relatedness between the components of that expression and the lexical items in the surrounding context? We adapt the **Lexical Coherence Graph (LCG)** introduced by Mesgar and Strube (2016) to measure the semantic relatedness between the words in the preceding context and the target expression contained in the final discourse sentence: we compute the cosine scores for the two components in the SV/VO expressions (i.e., the verb and its argument) and each word in the preceding discourse, relying on contextualized BERT embeddings (Devlin et al., 2019). The output is a graph connecting the preceding context to both the metaphorical expression and its synonymous alternative; edge weights are represented by the average of the respecting cosine values. The expression – metaphorical or literal – with the maximum weight (i.e., the largest average cosine score) is selected. Appendix A.2 provides details on the LCG score.

Contextualized discourse properties How do contextualized word representations prime for the use of a metaphorical vs. literal preference? Our question triggers a *cloze-task* style (Taylor, 1953) and applies the **Pre-trained Language Model BERT** (Devlin et al., 2019) in a zero-shot manner. We give a *cloze prompt* as input to the model, as in table 1: we mask the target expression, and the model selects the most probable answer amongst the two candidates (metaphorical or literal). We experiment with both BERT_{base} and BERT_{large}.

3.2 Experimental Setup

We compare our models’ predictions against two gold standards. We first evaluate the models’ outputs against the 1,000 discourses collected from *ukWaC*, containing the balanced *originally*-used metaphorical or literal expressions (**Orig. Data**) to investigate which discourse aspects (abstractness, emotionality, coherence, word representations) might have primed the metaphorical vs. literal selection. We expect the model predictions to provide insight into the features that influenced the *original* speakers’ preferences. We then zoom into the *annotated* version of the same data (**Anno. Data**), for which participants were asked which expression, metaphorical or literal, they favored in the given 287 discourses (cf. Section 2). We analyze whether (dis-)agreements between humans are also reflected in the models’ predictions.

4 Results and Analyses

We first analyze the models’ predictions with regard to metaphorical vs. literal usages in the *originally*-extracted data (**Orig. Data**), to address the question: do cognitive and linguistically-inspired models reflect metaphorical language usage encountered in natural language corpora? The first column of figure 1 presents the accuracy score of each model with regard to Orig. Data. All models reach around 50% of accuracy; however, they behave very differently regarding their individual predictions.³ Figure 1 presents the percentages of overlapping output decisions between our five models, revealing interesting insights: we observe 77% overlapping predictions between the PLMs with the frequency baseline, suggesting that the majority of PLM predictions is frequency-driven, which is not the case for most abstractness settings and emotionality. Emotionality itself correlates with abstractness in all settings but the one where we consider only adjectives, which is surprising as adjectives tend to carry a lot of emotions (Mohammad and Turney, 2010; Bostan and Klinger, 2019). The overlap between LCG and BERT reaches 62%, whereas its correlations with abstractness and emotionality are rather low, suggesting that decisions based on abstractness and emotionality are different to those based on word representations and semantic relatedness. At first sight, our different perspectives seem rather complementary, which is yet to be verified in future studies.

³Appendix B.1 presents further prediction differences.




Figure 1: Accuracy scores of each model with regard to the original data (1st column in red) and percentages of overlapping output decisions between models.

Overall, none of our models seems to reliably predict what is observed in natural language data. We thus cannot derive what might have triggered the original speakers to favor a metaphorical or literal expression over the other.

We then compare our models’ predictions to human behavior: are the proportions of metaphorical vs. literal predictions from our models similar to human perceptions (**Anno. Data**)? Figure 2 presents the proportions of the metaphorical expressions predicted by four of the models in relation to the proportions of metaphorical usages favored by the human judges.⁴ If humans and models were making similar decisions, all data points, each representing the proportion of metaphorical uses for a pair of expressions, would be on the regression line. We make different observations depending on the model in question. As far as frequency is concerned, the model predicts either the metaphorical or the literal expression (100% or 0% respectively, in the graph). Many literal expressions that are favored by participants (low x-axis %) seem to correlate with their higher occurrences in natural language corpora, e.g., *make* vs. *throw remark* (VO). However, humans also favor many metaphorical expressions which are less frequently used, e.g., *stir* vs. *cause excitement* (VO) and inversely, a few metaphorical expressions which are more frequent than their literal counterparts are not necessarily favored by people, e.g., *poison* vs. *corrupt mind* (VO). Thus, frequency does not seem to be a systematic factor for metaphorical/literal choices. Regard-

⁴Appendix B.2 shows the graphs for all models.




Figure 2: Proportions of predicted metaphorical expressions for each model, with regard to the proportions of these metaphorical expressions favored by annotators. Appendix B.2 shows the graphs for all models.

ing abstractness, many expressions fall perfectly on the regression line, e.g., *twist vs. misinterpret word* (VO), *story grab vs. intrigue* (SV), suggesting at first some interactions between the expression preferences and the abstractness of the respective discourses. However, the numerous metaphorical expressions that are predicted by the model (80%+) but not by humans, e.g., *taste vs. experience freedom* (VO), *factor shape vs. determine* (SV) do not confirm that hypothesis. Concerning the LCG, the picture is less clear. Many literal expressions are favored by both the model and humans, e.g., *color vs. affect judgement* (VO), but many metaphorical usages are preferred by humans when the model predicts the literal ones, e.g., *breathe vs. instill life* (VO). Therefore, metaphorical vs. literal coherence does not seem either to be a determining factor for metaphorical vs. literal preference, respectively, which does not support the context-salience hypothesis, where one would expect a metaphorical expression to be favored following a metaphorical discourse, and ditto for a literal expression/literal discourse (Kövecses, 2009). Finally, the overall low metaphorical predictions by the PLM are expressions for which humans provide high metaphorical proportions, such as *abuse vs. drink alcohol* (VO). This suggests a potential lack of metaphorical language representation, in line with our observation concerning abstractness and emotionality.

5 Discussion and Future Work

Previous research has provided evidence for interactions between metaphorical language and discourse properties, namely frequency, abstractness, affect, discourse coherence and contextualized word representations, on the *word-* and *sentence-*

level. In this work, we took a step further: we looked at the task of predicting the use of a metaphorical vs. synonymous literal expression and built models based on the above-mentioned features on the *discourse-level*. Our findings show that these discourse properties do not seem to be indicative of metaphorical usage.

We propose several directions for future work. First of all, we considered discourses of around five sentences, but the decision on the context window might have an impact on the findings. Further work exploring the optimal size of preceding context would be interesting. Another promising direction might analyze PLMs’ attention mechanisms (Tenney et al., 2019; Clark et al., 2019) on the presented task, and also explore the extent to which modifying the attention of such models, i.e., fine-tuning (Peters et al., 2019; Zhao and Bethard, 2020), improves their performance to mimic human preferences for metaphorical vs. literal usages. Finally, we have looked at five features; needless to say that exploring further discourse properties is necessary, such as co-reference, complexity, aptness, creativity, prototypicality, and the influence of genre and specific domains (e.g., religious/scientific texts).

6 Conclusion

We suggested five simple models to investigate WHY humans choose to use a metaphorical expression in a specific discourse. Regardless of the perspectives, our work demonstrates that a range of previously suggested salient discourse properties do not seem to influence preferences on the choice between synonymous metaphorical vs. literal expressions. Our findings thus ask for a more nuanced approach to metaphorical language choices in NLP.

Acknowledgements

We thank the anonymous reviewers for their useful and quality feedback. This research was supported by the DFG Research Grant SCHU 2580/4-1 *Multimodal Dimensions and Computational Applications of Abstractness*.

References

- Ghadi Alnafesah, Harish Tayyar Madabushi, and Mark Lee. 2020. [Augmenting Neural Metaphor Detection with Concreteness](#). In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 204–210, Online. Association for Computational Linguistics.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. [The WaCky Wide Web: A Collection of Very Large Linguistically Processed Web-Crawled Corpora](#). *Language Resources and Evaluation*, 43(3):209–226.
- Yuri Bizzoni and Shalom Lappin. 2018. [Predicting Human Metaphor Paraphrase Judgments with Deep Neural Networks](#). In *Proceedings of the Workshop on Figurative Language Processing*, pages 45–55, New Orleans, Louisiana. Association for Computational Linguistics.
- Isabelle Blanchette and Kevin Dunbar. 2001. [Analogy Use in Naturalistic Settings: The Influence of Audience, Emotion, and Goals](#). *Memory & Cognition*, 29(5):730–735.
- Daria Bogdanova. 2010. [A Framework for Figurative Language Detection Based on Sense Differentiation](#). In *Proceedings of 2010 Conference the Association for Computational Linguistics: Student Research Workshop*, pages 67–72, Uppsala, Sweden. Association for Computational Linguistics.
- Danushka Bollegala and Ekaterina Shutova. 2013. [Metaphor Interpretation Using Paraphrases Extracted from the Web](#). *PLOS ONE*, 8(9):1–10.
- Laura Ana Maria Bostan and Roman Klinger. 2019. [Exploring Fine-Tuned Embeddings that Model Intensifiers for Emotion Analysis](#). In *Proceedings of the Tenth Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 25–34, Minneapolis, USA. Association for Computational Linguistics.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. [Concreteness Ratings for 40 Thousand Generally Known English Word Lemmas](#). *Behavior Research Methods*, 64(3):904–911.
- Sven Buechel and Udo Hahn. 2018a. [Emotion Representation Mapping for Automatic Lexicon Construction \(Mostly\) Performs on Human Level](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2892–2904, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Sven Buechel and Udo Hahn. 2018b. [Word Emotion Induction for Multiple Languages as a Deep Multi-Task Learning Problem](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1907–1918, New Orleans, Louisiana. Association for Computational Linguistics.
- Sven Buechel, Susanna Rücker, and Udo Hahn. 2020. [Learning and Evaluating Emotion Lexicons for 91 Languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1202–1217, Online. Association for Computational Linguistics.
- Francesca M.M. Citron and Adele E. Goldberg. 2014. [Metaphorical Sentences Are More Emotionally Engaging Than Their Literal Counterparts](#). *Journal of cognitive neuroscience*, 26(11):2585–2595.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What Does BERT Look at? An Analysis of BERT’s Attention](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- L. Elizabeth Crawford. 2009. [Conceptual Metaphors of Affect](#). *Emotion review*, 1(2):129–139.
- Verna Dankers, Karan Malhotra, Gaurav Kudva, Volodymyr Medentsiy, and Ekaterina Shutova. 2020. [Being Neighbourly: Neural Metaphor Identification in Discourse](#). In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 227–234, Online. Association for Computational Linguistics.
- Verna Dankers, Marek Rei, Martha Lewis, and Ekaterina Shutova. 2019. [Modelling the Interplay of Metaphor and Emotion through Multitask Learning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2218–2229, Hong Kong, China. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rafael Ehren, Timm Lichte, Laura Kallmeyer, and Jakob Waszczuk. 2020. [Supervised Disambiguation of German Verbal Idioms with a BiLSTM Architecture](#). In *Proceedings of the Second Workshop on*

- Figurative Language Processing*, pages 211–220, Online. Association for Computational Linguistics.
- Paul Ekman. 1992. [An Argument for Basic Emotions](#). *Cognition and Emotion*, 6(3-4):169–200.
- Andrew Gargett and John Barnden. 2015. [Modeling the Interaction Between Sensory and Affective Meanings for Detecting Metaphor](#). In *Proceedings of the Third Workshop on Metaphor in NLP*, pages 21–30, Denver, Colorado. Association for Computational Linguistics.
- Matt Gedigian, John Bryant, Srini Narayanan, and Branimir Ciric. 2006. [Catching Metaphors](#). In *Proceedings of the Third Workshop on Scalable Natural Language Understanding*, pages 41–48, New York City, New York. Association for Computational Linguistics.
- Raymond W. Gibbs. 1989. [Understanding and Literal Meaning](#). *Cognitive Science*, 13(2):243–251.
- Raymond W. Gibbs. 2002. [A New Look at Literal Meaning in Understanding What Is Said and Implicated](#). *Journal of Pragmatics*, 34(4):457–486.
- Raymond W. Gibbs. 2006. [Metaphor Interpretation as Embodied Simulation](#). *Mind & Language*, 21(3):434–458.
- Rachel Giora. 1997. [Understanding Figurative and Literal Language: The Graded Salience Hypothesis](#). *Cognitive Linguistics*, 8(3):183–206.
- Rachel Giora and Ofer Fein. 1999. [On Understanding Familiar and Less-Familiar Figurative Language](#). *Journal of Pragmatics*, 31(12):1601–1618.
- Sam Glucksberg. 1989. [Metaphors in Conversation: How Are They Understood? Why Are They Used?](#) *Metaphor and Symbolic Activity*, 4(3):125–143.
- Rowan Hall Maudslay, Tiago Pimentel, Ryan Cotterell, and Simone Teufel. 2020. [Metaphor Detection using Context and Concreteness](#). In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 221–226, Online. Association for Computational Linguistics.
- Michael A.K. Halliday and Ruqaiya Hasan. 1976. *Cohesion in English*. Longman.
- Albrecht Werner Inhoff, Susan D. Lima, and Patrick J. Carroll. 1984. [Contextual Effects on Metaphor Comprehension in Reading](#). *Memory and Cognition*, 12(6):558–567.
- Daniel Jurafsky and James H. Martin. 2019. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, 3rd edition. Prentice Hall, USA.
- Maximilian Köper and Sabine Schulte im Walde. 2016. [Distinguishing Literal and Non-Literal Usage of German Particle Verbs](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 353–362, San Diego, CA, USA. Association for Computational Linguistics.
- Maximilian Köper and Sabine Schulte im Walde. 2018. [Analogies in Complex Verb Meaning Shifts: the Effect of Affect in Semantic Similarity Models](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 150–156, New Orleans, Louisiana. Association for Computational Linguistics.
- Zoltán Kövecses. 2009. [The Effect of Context on the Use of Metaphor in Discourse](#). *Ibérica: Revista de la Asociación Europea de Lenguas para Fines Específicos.*, 17(17):11–23.
- George Lakoff and Mark Johnson. 1980. *Metaphors We Live By*. University of Chicago Press, Chicago.
- Rui Mao, Chenghua Lin, and Frank Guerin. 2018. [Word Embedding and WordNet Based Metaphor Identification and Interpretation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1222–1231, Melbourne, Australia. Association for Computational Linguistics.
- James H. Martin. 2008. [A Corpus-based Analysis of Context Effects on Metaphor Comprehension](#). In A. Stefanowitsch and S. Th. Gries, editors, *Corpus-Based Approaches to Metaphor and Metonymy*, pages 214–236. De Gruyter Mouton.
- Mohsen Mesgar and Michael Strube. 2016. [Lexical Coherence Graph Modeling Using Word Embeddings](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1414–1423, San Diego, California. Association for Computational Linguistics.
- Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. [Metaphor as a Medium for Emotion: An Empirical Study](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 23–33, Berlin, Germany. Association for Computational Linguistics.
- Saif Mohammad and Peter Turney. 2010. [Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon](#). In *Proceedings of the 2010 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA. Association for Computational Linguistics.

- Michael Mohler, Mary Brunson, Bryan Rink, and Marc Tomlinson. 2016. [Introducing the LCC Metaphor Datasets](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation*, pages 4221–4227, Portorož, Slovenia. European Language Resources Association.
- Andrew Ortony. 1975. [Why Metaphors Are Necessary and Not Just Nice](#). *Educational Theory*, 25(1):45–53.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Prisca Piccirilli and Sabine Schulte im Walde. 2021. [Synonymous Pairs of Metaphorical and Literal Expressions in Context: An Empirical Study and Dataset to tackle or to address the question](#). In *Proceedings of the Workshop DiscAnn*, Tübingen, Germany.
- Prisca Piccirilli and Sabine Schulte im Walde. 2022. [Features of Perceived Metaphoricity on the Discourse Level: Abstractness and Emotionality](#). In *Proceedings of the 13th Conference on Language Resources and Evaluation*, Marseille, France. European Language Resources Association. To appear.
- Christina Schäffner. 2004. [Metaphor and Translation: Some Implications of a Cognitive Approach](#). *Journal of Pragmatics*, 36(7):1253–1269.
- Ekaterina Shutova. 2010. [Automatic Metaphor Interpretation as a Paraphrasing Task](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1029–1037, Los Angeles, California. Association for Computational Linguistics.
- Ekaterina Shutova and Simone Teufel. 2010. [Metaphor Corpus Annotated for Source - Target Domain Mappings](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, Valletta, Malta. European Language Resources Association.
- Caroline Sporleder and Linlin Li. 2009. [Unsupervised Recognition of Literal and Non-Literal Use of Idiomatic Expressions](#). In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, pages 754–762, Athens, Greece. Association for Computational Linguistics.
- Gerard J. Steen. 2004. [Can Discourse Properties of Metaphor Affect Metaphor Recognition?](#) *Journal of Pragmatics*, 36(7):1295–1313.
- Gerard J. Steen, Aletta G. Dorst, J. Berenike Herrmann, Anna A. Kaal, Tina Krennmayr, and Tryntje Pasma. 2010. [A Method for Linguistic Metaphor Identification. From MIP to MIPVU](#). John Benjamins.
- Anatol Stefanowitsch. 2008. [Words and Their Metaphors: A Corpus-Based Approach](#), pages 63–105. De Gruyter Mouton.
- Kevin Stowe, Nils Beck, and Iryna Gurevych. 2021. [Exploring Metaphoric Paraphrase Generation](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 323–336, Online. Association for Computational Linguistics.
- Wilson L. Taylor. 1953. [“Cloze Procedure”: A New Tool for Measuring Readability](#). *Journalism quarterly*, 30(4):415–433.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT Rediscovered the Classical NLP Pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Xiaoyu Tong, Ekaterina Shutova, and Martha Lewis. 2021. [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*, pages 4673–4686, Online. Association for Computational Linguistics.
- Yulia Tsvetkov, Elena Mukomel, and Anatole Gershman. 2013. [Cross-Lingual Metaphor Detection Using Common Semantic Features](#). In *Proceedings of the First Workshop on Metaphor in NLP*, pages 45–51, Atlanta, Georgia. Association for Computational Linguistics.
- Peter Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. [Literal and Metaphorical Sense Identification through Concrete and Abstract Context](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 680–690, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Raymond van den Broek. 1981. [The Limits of Translatability Exemplified by Metaphor Translation](#). *Poetics Today*, 1(4):73–87.
- Henk J. van Jaarsveld and Gilbert E. Rattink. 1988. [Frequency Effects in the Processing of Lexicalized and Novel Nominal Compounds](#). *Journal of Psycholinguistic Research*, 17:447–473.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). In *ADvances in Neural Information Processing Systems*, pages 5998–6008.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. [Norms of Valence, Arousal, and Dominance for 13,915 English Lemmas](#). *Behavior Research Methods*, 45(4):1191–1207.

Moritz Wittmann, Maximilian Köper, and Sabine Schulte im Walde. 2017. [Exploring Soft-Clustering for German \(Particle\) Verbs across Frequency Ranges](#). In *Proceedings of the 12th International Conference on Computational Semantics*, Montpellier, France.

Yiyun Zhao and Steven Bethard. 2020. [How does BERT’s attention change when you fine-tune? An analysis methodology and a case study in negation scope](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4729–4747, Online. Association for Computational Linguistics.

Jianing Zhou, Hongyu Gong, and Suma Bhat. 2021. [PIE: A Parallel Idiomatic Expression Corpus for Idiomatic Sentence Generation and Paraphrasing](#). In *Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)*, pages 33–48, Online. Association for Computational Linguistics.

A Models

A.1 Emotionality

Buechel et al. (2020) presented a new methodology to automatically generate lexicons for 91 languages comprising eight emotional variables: *Valence, Arousal, Dominance* (VAD) as well as the five basic emotions *Joy, Anger, Surprise, Fear, Surprise* (BE5) (Ekman, 1992). As a source dataset, they used the English emotion lexicon from Warriner et al. (2013), comprising about 14K entries in VAD format collected via crowdsourcing. They applied the BE5 ratings from Buechel and Hahn (2018a) to convert the VAD ratings. Via their monolingual state-of-the-art multi-task feed-forward network (Buechel and Hahn, 2018b), they projected ratings on these eight variables, resulting in an English lexicon containing 2M word type entries with very high correlation with human judgments (around 90% for each variable).

We are interested in the emotionality of lexical items, i.e., the emotional load that a term conveys, rather than the actual emotion a term refers to. We therefore use the BE5 ratings of the English lexicon for our study. Out of the five scores that a term receives – one for each emotion, we assume that its highest score is reflective of the "emotional load" of that term, i.e., how much emotion it conveys. For example, the lexical item "truth" obtained the scores 2.24, 1.46, 1.4, 1.49, 1.46 for Joy, Anger, Sadness, Fear and Surprise, respectively. In our experiments, the term "truth" is therefore attributed the score of 2.24 as its emotional load.

A.2 Lexical Coherence Graph

Following Mesgar and Strube (2016), we measure the semantic relatedness between words represented by their word embeddings, computing the cosine score between the two words of the expression (SV or VO) with each word of the preceding discourse. Consider v_a, v_b, v_c , the word vectors for word a in the preceding discourse A , word b in the metaphorical expression B_m contained in the last sentence B and word c in the literal expression C_l contained in the last sentence C , respectively. The cosine scores $\cos(v_a, v_b)$ and $\cos(v_a, v_c)$ between the two word vectors is a measure of semantic connectivity of the two words. The range of $\cos(v_a, v_b)$ and $\cos(v_a, v_c)$ is between $[-1, +1]$, showing how well the two words are semantically correlated. Figure 4 shows how relatedness is measured, and figure 5 shows the output of the graph.

B Results

B.1 Evaluation: Predictions vs. Original Data

Figure 3 presents the percentages of metaphorical expressions predicted by each model. As mentioned in section 4, the accuracy scores of all models reach around 50%, but they actually behave very differently with regard to performances for metaphorical vs. literal predictions. Remember that the originally-collected discourses are perfectly balanced (cf. Section 2), where half of the discourses contains metaphorical expressions and the other half contains literal expressions from synonymous pairs. As expected, the frequency baseline therefore reaches 50% of accuracy. We note however that the literal usage of the pairs is the most frequent in the *ukWaC* corpus (37/50 pairs), which leads the model to predict the literal counterpart in 74% of the cases. The abstractness and emotionality models mostly select for the metaphorical usages (up to 100% for Emo.), as the discourses are considered more abstract and more emotionally-loaded based on the respective norms. This suggests that original speakers did not perceive the degree of abstractness and emotionality of the discourse as triggers to favor one usage over the other. This aligns with the findings in Piccirilli and Schulte im Walde (2022), who analyzed the relationship between the abstractness and emotionality of the preceding discourses with human preferences for metaphorical vs. literal expressions. The LCG model favors the literal expressions as well (62%), while both BERT_{base} and BERT_{large} predominantly predict the literal expressions (81.80% and 80.10%, respectively).




Figure 3: Percentages of metaphorical expressions predicted by each model.




Figure 4: Discourse A with three words $\{w_1, w_2, w_3\}$ and sentences B and C with four words, where $\{w_5, w_6\}$ are the two words composing the metaphorical expression, and $\{w_9, w_{10}\}$ are composing the literal paraphrase. Depending on the expression input (SV or VO), the respective subject or object is identical in $\{w_5, w_6\}$ and $\{w_9, w_{10}\}$, as only the verb is used either as a metaphorical or a literal variant. The semantic relatedness between each word in $\{w_5, w_6\}$ and in $\{w_9, w_{10}\}$ is computed with each word in A .




Figure 5: The word relation with the maximum weight (here B_m as indicated by the plain line) represents a stronger connection with the preceding discourse A .

B.2 Evaluation: Predictions vs. Annotated Data

Figure 6 presents an overview of the proportion of expressions that are predicted metaphorically by the models with regard to the preferences of the same expressions by human annotators.

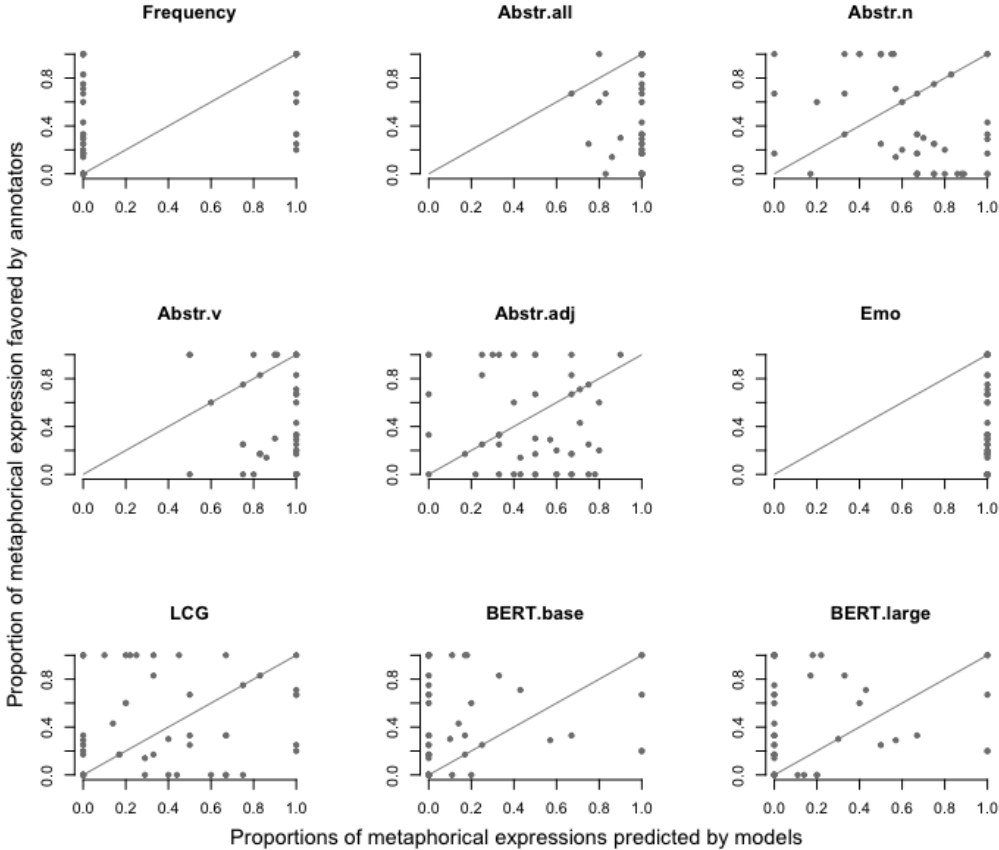


Figure 6: Proportions of the metaphorical expressions predicted by the models with regard to the proportions of these usages to be favored by the participants.