

Relational words have high metaphoric potential

Anja Jamrozik

Northwestern University
Department of Psychology - 2029 Sheridan Road
Evanston, IL 60208-2710, USA
a.jamrozik@u.northwestern.edu

Eyal Sagi

Northwestern University
Kellogg School - 2001 Sheridan Road
Evanston, IL 60208-2710, USA
eyal@u.northwestern.edu

Micah Goldwater

Northwestern University
Department of Psychology - 2029 Sheridan Road
Evanston, IL 60208-2710, USA
micahbg@gmail.com

Dedre Gentner

Northwestern University
Department of Psychology - 2029 Sheridan Road
Evanston, IL 60208-2710, USA
gentner@northwestern.edu

Abstract

What influences the likelihood that a word will be used metaphorically? We tested whether the likelihood of metaphorical use is related to the relationality of a word's meaning. Relational words name relations between entities. We predicted that relational words, such as verbs (e.g., *speak*) and relational nouns (e.g., *marriage*) would be more likely to be used metaphorically than words that name entities (e.g., *item*). In two experiments, we collected expert ratings of metaphoricity for uses of verbs, relational nouns, and entity nouns collected from a corpus search. As predicted, uses of relational words were rated as more metaphorical than uses of entity words. We discuss how these findings could inform NLP models of metaphor.

1 Introduction

Our goal is to assess the *metaphoric potential* of words and word classes—by which we mean the likelihood that the word (or word class) will be used metaphorically. By *metaphorical use*, we mean the use of a word to convey ideas that are not part of its basic or standard meaning. We note that metaphoric potential does not equate to metaphoric salience. Many common metaphorical uses are not particularly salient. These include non-spatial, abstract uses of prepositions (e.g., *in love*, *between assignments*) and metaphorical uses of verbs (e.g., *run for office*, *fall behind*).

One could question whether it is useful to identify the kind of metaphorical uses just mentioned. Shutova, Tuefel, & Korhonen (2012) point out that

it may not be relevant to NLP applications to identify highly conventionalized or “dead” metaphorical uses, ones for which a metaphorical sense has become dominant and earlier literal senses have become obsolete. An example is the verb *impress*, which was originally used in printing contexts and meant ‘to make a mark with pressure’ but now is typically used to mean ‘to produce admiration in someone’. While we agree that identifying such ‘dead’ metaphors may not be useful, we note that there are many conventional metaphors that also retain a healthy literal sense; and in these cases, identifying their metaphorical uses can be challenging. An example is the word *glow*, which can be used literally (The lamp *glows* dimly) as well as metaphorically (Her face *glows* with joy). In our research, we will consider both conventional and novel extensions of a word's meaning but will focus more on conventional metaphorical uses.

Many factors influence a word's metaphoric potential—including its conventionality as a prior metaphoric source, its familiarity, and whether it belongs to a conceptual system whose other members are often used metaphorically¹. We focus here on a relatively unexplored factor: namely, the relationality of the word's meaning.

Relational words are words that take more than one argument. These include verbs (KNOW(Sue, Ida)), prepositions (ON(fence, hill)), and relational

¹ For example, one might be able to say “Let me slide this to him” meaning “Let me communicate this to him in a smooth manner,” because the “conduit” metaphoric system (Reddy, 1979; Lakoff & Johnson, 1980) includes other instances of caused-physical-motion verbs used to convey communication of ideas (e.g., “Is this message getting across to you?”).

nouns² (FRIEND OF(John, George). Relational nouns (e.g., *guest, host, party*), which name relations or systems of relations, can be contrasted with entity nouns (e.g., *zebra, thing*), which name entities defined by their intrinsic properties (see Gentner & Kurtz, 2005; Goldwater Markman & Stilwell, 2011, and Markman & Stilwell, 2001).

Goldwater and Willits (2010) explored ways to distinguish relational from entity nouns based on their distributional patterns. All of the nouns from the Goldwater et al. (2011), Gentner & Kurtz (2005) and Gentner & Asmuth (2008) studies were normed for their relationality by naïve participants, who rated to what degree each word expressed relational or entity meanings. Goldwater and Willits analyzed the distributions of the top 50 highest rated relational nouns and top 25 entity nouns on a 10,000 word corpus from Wikipedia.com. The two kinds of nouns were found to have distinct distributional patterns. For relational nouns, the most frequent immediate following word is a preposition connecting the noun to another term (as in ‘proportion *of* X’ or ‘barrier *to* X’)³. In contrast, the most frequent immediate follower for entity nouns is *and*⁴. These distributional patterns can be used to predict noun type. Given two words, their distributional similarity can predict whether they are of the same noun type or different noun type with close to 90% accuracy. Although further study is needed of how well these results extend to a larger sample of nouns, we believe this is a promising direction.

Our hypothesis that metaphorical potential is related to relationality is supported by evidence that relational words are more *mutable* than entity words—that is, the meanings of relational words adjust more to fit their contexts than do the meanings of entity nouns (Gentner, 1981). Psycholinguistic studies of sentence interpretation have found this pattern both across word classes (nouns vs. verbs) and within the noun class (entity nouns vs. relational nouns). For example, when partici-

pants were asked to paraphrase semantically strained sentences in which the noun did not meet the argument specification of the verb (e.g., The car laughed), their paraphrases were far more likely to preserve the meanings of the nouns than of the verbs (Gentner & France, 1988) (e.g., ‘The automobile sputtered and refused to start’). Further evidence comes from studies testing recognition memory of nouns and verbs (Kersten & Earles, 2004). Verbs were recognized better if found in the same context as at encoding, but nouns were recognized equally well whether in the same context or in a new context. Kersten and Earles suggested that this difference stemmed from the greater mutability of verbs (Gentner, 1981). Because the meanings of verbs adjust more to their contexts than do the meanings of nouns, a verb may be interpreted as having very different meanings at encoding and at test. This made it difficult for participants to recognize that the same word was used in both cases.

There is also evidence that relational nouns have greater mutability than entity nouns. Using a similar paradigm to the one used by Kersten and Earles (2004), Asmuth and Gentner (2005) gave participants conceptual combinations consisting of one relational noun and one entity noun (e.g., *a truck limitation, a barrier peanut*) and later tested their recognition memory for the individual nouns, which were either presented in the old context (the same context as at encoding) or in a new context. Overall, recognition of entity nouns was better than recognition of relational nouns. Additionally, recognition for relational nouns was more impaired by a shift to a new context than was recognition of entity nouns. This is consistent with the mutability claim that relational nouns are encoded in a context-dependent manner.

To summarize, the evidence that relational words are more mutable than entity words suggest that they should have greater metaphoric potential. If a word’s meaning readily adjusts to its context, this can result in metaphoric extensions that go beyond the word’s basic or standard meaning. Since relational words are more mutable than entity words, they should be more likely to be extended in this way.

We test this prediction both across and within word classes. Comparing across word classes, we predict that verbs will have greater metaphoric potential than nouns. Comparing within word class,

² A common test for relational nouns involves the use of genitive *of* (Barker and Dowty, 1993). For example, *friend* is a relational noun, and “friend of John” and “John’s friend” are both grammatical and interchangeable. The *of* form is not grammatical for non-relational nouns (e.g., John’s truck, *truck of John) (Barker, 1995).

³ Goldwater and Willits found that relational nouns were most frequently followed by *of*, but distributional approaches could be extended to other common role-bearing terms.

we predict that relational nouns will have greater metaphoric potential than entity nouns.

There is already evidence for the predicted difference between word classes: metaphorical uses of verbs have been found to be more common than metaphorical uses of nouns in poetry (Brooke-Rose, 1958), in classroom discourse (e.g., Cameron, 2003), and across various spoken and written genres (e.g., Shutova & Teufel, 2010; Steen, Dorst, Herrmann, Kaal, Krennmayr, & Pasma, 2010).

In the studies that follow, we tested our predictions using data collected from a corpus search. We randomly sampled uses of verbs, relational nouns, and entity nouns and collected novelty and metaphoricity ratings for each of these uses. We were particularly interested in the pattern among conventional metaphorical uses, which are the most challenging to identify with NLP methods.

2 Experiment 1

2.1 Materials

The materials consisted of 20 uses each of nine entity nouns, eighteen relational nouns⁵, and nine verbs. The entity and relational nouns were selected based on data from a previous rating task (the same as provided Goldwater and Willits with their sample). The entity nouns we selected were rated as conveying an entity meaning to a higher degree than the relational nouns, and vice versa, all $ps < .001$. The nine verbs were selected to match the frequencies of the nouns, using data from the Corpus of Contemporary American English (COCA) (Davies, 2008-). There were no differences in the frequencies of the word types, $F(3, 32) = .857, p = .474$.

We collected a random sample of 20 uses of each of the 36 words from COCA, with an equal number from the spoken, fiction, magazine, news, and academic registers. We used the following criteria to determine whether a word use would be included in the sample. First, the word had to be used as a noun or verb, depending on its preselected word type⁶ (e.g., for the verb *talk*, we only

⁵ There were two different kinds of relational nouns (*schema* nouns, which refer to relational systems, e.g., *party*, and *role* nouns, which refer to roles within such systems, e.g., *guest*), but we do not distinguish them in the analyses that follow.

⁶ We reserve the use of the term *word class* for accepted syntactic distinctions (e.g., nouns vs. verbs) and use the term

collected uses in which *talk* was used as a verb). Second, the word had to be used in a full sentence or phrase, so as to give sufficient context to determine how metaphorical the use was. Third, the sentence had to be a statement, not a question. Finally, the word could only appear once in the sentence.

2.2 Rating Task

The three raters were Ph.D. students of English or Comparative Literature. They were chosen because they had extensive experience identifying figurative language and would be able to identify metaphors that may not have been particularly salient to average readers.

The raters were given sets of sentences with the key terms bolded and underlined: e.g.,

The human mind, the only **device** capable of traveling through time, tends to want to stay in its own time.

A smartphone or other technological **device** used during worship also can be a distraction.

They were instructed to rate the metaphoricity and the novelty of the indicated words, on two separate scales from 1 (not at all novel/metaphorical) to 6 (very novel/metaphorical). The separate ratings were used to ensure that raters were not conflating novelty and metaphoricity. For each item, we calculated the average of the individual novelty and metaphoricity scores assigned by the raters.

2.3 Results

As predicted, we found that the metaphor ratings differed across word types. We conducted an ANOVA with the average metaphor ratings for each use as the dependent variable. The type of word rated (Entity Noun, Relational Noun, or Verb) was the independent variable and the specific word was a random effect. This resulted in a marginally significant effect of word type, $F(2, 684) = 3.22, p = .053$. Post-hoc Tukey HSD tests revealed that uses of verbs ($M = 1.77, SD = 0.76$) and relational nouns ($M = 1.75, SD = 0.91$) were rated as more metaphorical than uses of entity nouns ($M = 1.27, SD = 0.61$), $ps < .001$.

The difference between word types was also

word type to differentiate the kinds of words compared in these experiments.

marginal when analyzing only conventional uses (i.e., those rated “1” for novelty), $F(2, 437) = 2.80$, $p = .075$. Uses of verbs ($M = 1.36$, $SD = 0.33$) and relational nouns ($M = 1.27$, $SD = 0.41$) were rated as more metaphorical than uses of entity nouns ($M = 1.16$, $SD = 0.29$), $ps < .01$. No differences in metaphoricality were observed between word types for novel uses (i.e., those rated higher than “1” for novelty), $F(2, 213) = 1.49$, $p = 0.24$.

While these results are only marginally significant, they provide encouragement that relationality might influence the metaphoric potential of a word.

2.4 Concreteness, imageability and metaphoric potential

One concern regarding Experiment 1 is that we did not control for concreteness or imageability of the words. Previous research is conflicted about what effects concreteness and imageability should have on a word’s metaphorical potential. On the one hand, it has been argued that greater concreteness (e.g., Katz, 1989) and imageability (e.g., Goatly, 2011) should result in greater metaphoricality. However, previous work by Gentner and Asmuth (2008) has shown that relational words, which we found to have greater metaphorical potential, tend to be less concrete and imageable than entity words. In accordance with these findings, we found that concreteness and imageability of the words (using data from MRC Psycholinguistic Database - Coltheart, 1981) varied across word types (concreteness: $F(2, 26) = 27.36$, $p < .001$; imageability: $F(2, 26) = 15.71$, $p < .001$). Entity nouns were more concrete and imageable than relational nouns and verbs, all $ps < .001$.

The entity nouns were more concrete and imageable than the relational words, but their uses were less metaphorical. Thus in our sample, the relationship between concreteness, imageability, and metaphoricality was the opposite of that predicted by Katz (1989) and Goatly (2011)⁷: more concrete words were rated as less metaphorical, $r(27) = -.43$, $p = .021$. (The relationship between imageability and metaphoricality was not significant ($r(27) = -.30$, $p = .117$)). Because of these findings, in the next experiment we controlled for concreteness and imageability across word classes.

⁷ It is possible that Katz’s and Goatly’s predictions drew on different contexts of use from those in our corpus.

3 Experiment 2

3.1 Materials

The materials consisted of 20 uses each of eight entity nouns, sixteen relational nouns, and eight verbs. The words were selected in the same manner as those in Experiment 1, except that in addition to controlling for frequency across word types, we also controlled for concreteness and imageability. In the resulting sets, there were no differences in the concreteness of the word types (Coltheart, 1981), $F(2, 29) = .745$, $p = .484$, in the imageability of the word types (Coltheart, 1981), $F(2, 29) = .043$, $p = .958$, nor in the frequencies of the word types (using frequency data from COCA), $F(2, 29) = .144$, $p = .867$.

3.2 Rating task

The raters were three Ph.D. students of English or Comparative Literature who had not participated in the first experiment. The raters received the same instructions and followed the same procedure as in the first experiment.

3.3 Results

Overall, as predicted, we found that uses of relational words were rated as more metaphorical than uses of entity words. An ANOVA like that used in Experiment 1 showed that the average metaphor ratings differed across word types, $F(2, 608) = 3.77$, $p < .05$. Uses of verbs ($M = 2.35$, $SD = 1.69$) were rated as more metaphorical than uses of entity nouns ($M = 1.47$, $SD = 1.08$) and relational nouns ($M = 1.67$, $SD = 1.21$), $ps < .001$. However (in contrast to the first study) relational nouns were not significantly more metaphorical overall than entity nouns ($p = .17$).

The pattern was stronger when we looked only within conventional uses⁸ (i.e., those rated “1” for novelty), $F(2, 513) = 4.22$, $p < .05$. Conventional uses of verbs were rated as more metaphorical ($M = 2.19$, $SD = 1.59$) than uses of entity nouns ($M = 1.16$, $SD = 0.56$) and relational nouns ($M = 1.42$, $SD = 0.96$), all $ps < .001$. Uses of relational nouns were also rated as more metaphorical than uses of entity nouns, $p < .05$.

As in the first study, there were no differences

⁸ Conventional uses made up 85% (545/640) of the sample.

between word types for the novel uses (i.e., those rated higher than “1” for novelty), $F(2, 70) = 1.22$, $p = 0.31$.

4 Discussion

The results of two experiments provide support for the hypothesis that relational words have greater metaphoric potential than entity words. In the first experiment, verbs and relational nouns were rated as marginally more metaphorical than entity nouns. In the second experiment, in which concreteness and imageability were equated across the word types, verbs were rated more metaphorical than nouns. Within conventional uses, verbs were rated as more metaphorical than nouns, and relational nouns were rated more metaphorical than entity nouns.

4.1 Relationality and language change

These findings accord with our prediction that relational words have greater metaphorical potential than entity words, and that this pattern is stronger for conventional uses. Why is this the case? We conjecture that there may be two paths at work here. First, metaphor conventionalization may result in words acquiring relational senses. According to the Career of Metaphor hypothesis (Bowdle & Gentner, 1999, 2005; Gentner, Bowdle, Wolff, & Boronat, 2001; Gentner & Wolff, 1997) when a word is used in a novel metaphoric way, the use is processed by aligning the literal target and base of the metaphor in order to abstract their common structure. If the base term is repeatedly paired with other similar target terms, the structure abstracted through alignment may become another conventional meaning of the base term. Providing some initial support for this idea, Zharikov and Gentner (2002) traced the meanings of current relational words (e.g., *bridge*), and found that their relational meanings had evolved from earlier concrete, entity meanings. This idea also fits with accounts that argue that metaphoric use is one of the mechanisms precipitating semantic shifts in meaning (e.g., Traugott, 2004).

A second conjecture is that because relational words often identify deep relations among their arguments, uses of relational words across domains should result in more apt and relevant metaphors, which may therefore be more readily accepted.

This means that metaphorical uses of relational words may be more likely to become conventionalized than metaphorical uses of entity words.

4.2 Applications to NLP

The metaphoric uses we found in our experiments were in general not high-salient, striking figures of speech. Moreover, these metaphoric uses co-existed with literal uses of the same word (conventional non-metaphoric language was still the most common form). It is these unstriking, conventional metaphorical uses that pose a challenge for NLP (Shutova, Tuefel, & Korhonen, 2012).

How can NLP models of metaphor utilize our results? One possibility is to use differences in metaphoric potential among word types to inform searches for metaphoric language in corpora. Such an application would naturally also require a method for identifying relational words. Of course, since verbs appear to have a higher metaphoric potential than nouns, just using grammatical category information that is already available can make a substantial gain. Moreover, the distributional findings discussed earlier offer a potential way to distinguish relational from entity nouns (Goldwater & Willits, 2010). Assuming their results generalize, it might be feasible to distinguish relational nouns from entity nouns using distributional information.

In sum, we believe that taking into account the relationality of words has the potential to improve NLP models of metaphor. We look forward to the research that would come from uniting this psychological research with current research in NLP.

Acknowledgments

The research described here was conducted in the Spatial Intelligence and Learning Center (SILC), a center funded by NSF grant SBE-0541957. We thank the six raters who participated in our rating tasks and the Cognition and Language Lab.

References

- Asmuth, J. & Gentner, D. (2005). Context sensitivity of relational nouns. *Proceedings of the Twenty-seventh Annual Meeting of the Cognitive Science Society*, 163-168.
- Barker, C. (1995). *Possessive Descriptions*. Stanford, CA: CSLI Publications.
- Barker, C., & Dowty, D. (1993). Non-verbal thematic proto-roles. *Proceedings of the North East Linguistic Society 23*, 49-62.
- Brooke-Rose, C. (1958). *A grammar of metaphor*. London: Seeker & Warburg.
- Bowdle, B., & Gentner, D. (1999). Metaphor comprehension: From comparison to categorization. *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society*, 90-95.
- Bowdle, B., & Gentner, D. (2005). The career of metaphor. *Psychological Review*, 112(1), 193-216.
- Cameron, L. (2003). *Metaphor in educational discourse*. New York: Continuum.
- Coltheart, M. (1981). The MRC Psycholinguistic Database. *Quarterly Journal of Experimental Psychology*, 33A, 497-505.
- Davies, M. (2008-) The Corpus of Contemporary American English (COCA): 425 million words, 1990-present. Available online at <http://corpus.byu.edu/coca/>.
- Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, 4, 161-178.
- Gentner, D., & Asmuth, J. (2008). Can relationality be distinguished from abstractness in noun mutability? *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, 863-868.
- Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (2001). Metaphor is like analogy. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 199-253). Cambridge, MA: MIT Press.
- Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 343-382). San Mateo, CA: Morgan Kaufmann.
- Gentner, D., & Kurtz, K. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman & P. W. Wolff (Eds.), *Categorization inside and outside the lab*. (pp. 151-175). Washington, DC: APA.
- Gentner, D., & Wolff, P. (1997). Alignment in the processing of metaphor. *Journal of Memory and Language*, 37, 331-355.
- Goatly, A. (2011). *The language of metaphors* (2nd ed.). London: Routledge.
- Goldwater, M. B., Markman, A. B., Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, 118, 359-376.
- Goldwater, M. B., & Willits, J. (2010). *The Linguistic Distribution of Relational Categories*. Poster presented at the 32nd annual conference of the Cognitive Science Society, Portland, Oregon.
- Kersten, A.W., & Earles, J.L. (2004). Semantic context influences memory for verbs more than memory for nouns. *Memory & Cognition*, 32, 198-211.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: The University of Chicago Press.
- Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Intelligence*, 13, 329-358.
- Reddy, M. J. (1979). The conduit metaphor: A case of frame conflict in our language about language. In A. Ortony (Ed.), *Metaphor and thought*, 1st ed. (pp. 284-324). Cambridge, England: Cambridge University Press.
- Shutova, E., & Teufel, S. (2010). Metaphor corpus annotated for source-target domain mappings. *Proceedings of LREC 2010*, 3255-3261.
- Shutova, E., Teufel, S. & Korhonen, A. (2012). Statistical metaphor processing. *Computational Linguistics*, 39(2), 1-53.
- Steen, G.J., Dorst, A.G., Herrmann, J.B., Kaal, A.A., Krennmayr, T., & Pasma, T. (2010). *A method for linguistic metaphor identification: From MIP to MIPVU*. Philadelphia: John Benjamins.
- Traugott, E. C. (2004). Semantic Change. In K. Brown (Ed.) *Encyclopedia of Language and Linguistics* (2nd ed.). Oxford: Elsevier.
- Zharikov, S., & Gentner, D. (2002). Why do metaphors seem deeper than similes? *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, 976-981.