This file is a preprint and may contain errors or omissions not present in the final published version.

Why is semantic change asymmetric?

The role of concreteness and word frequency and metaphor and metonymy

Bodo Winter^{1*} & Mahesh Srinivasan²

¹ University of Birmingham, Dept. of English Language & Linguistics

² University of California, Berkeley, Dept. of Psychology

* corresponding author: bodo@bodowinter.com

Abstract

Metaphors and other tropes are commonly thought to reflect asymmetries in concreteness, with concrete sources being used to talk about relatively more abstract targets synchronically. Similarly, originating senses in diachronic semantic change have often been argued to be more concrete than extended senses. In this paper, we use a dataset of cross-linguistically attested semantic changes to empirically test the idea that asymmetries in figurative language are predicted by asymmetries in concreteness. We find only weak evidence for the role of concreteness and argue that concreteness is not a helpful notion when it comes to describing changes where both originating and extended senses are highly concrete (e.g., skin > bark, liver > lungs). Moreover, we find that word frequency data from English and other languages is a stronger predictor of these typologically common semantic changes. We discuss the implications of our findings for metaphor theory and theories of semantic change.

Keywords: asymmetry; unidirectionality; taboo; figurative language; language evolution; word frequency

1. Introduction

Concepts do not live in isolation but are embedded in complex webs of relations. Metaphor and metonymy are two common ways through which concepts are related to each other. For example, perception verbs often acquire new cognitive meanings (Evans & Wilkins, 2000; Ibarretxe-Antuñano, 2008; San Roque et al., 2018; Sweetser, 1990), as exemplified by the use of English “I see” to mean “I understand.” This evidences a historical change from ‘vision’ to ‘knowledge.’ Such metaphoric changes have taken place for many of the linguistic structures observed in present-day languages, as exemplified by words for abstract grammatical and spatial relations, which have often been derived from concrete body part terms (Dingemanse, 2009; Heine et al., 1991).

One of the most dominant frameworks in metaphor research is conceptual metaphor theory (Gibbs, 1994; Kövecses, 2002; Lakoff & Johnson, 1980; Littlemore, 2019), including its more recent extensions (Barcelona, 2003; de Mendoza Ibáñez & Masegosa, 2011; Grady, 1997; Lakoff, 2008, 2012, 2014). While some theories of metaphor allow mappings to be symmetrical (e.g., Blending Theory; Fauconnier & Turner, 1998), conceptual metaphor theory views metaphors as involving asymmetrical mappings, which extend from a ‘source’ to a ‘target’ domain. This asymmetry between source and target domains does not only characterize synchronic language use (speakers use words associated with source domains to talk about target domains more than the reverse), but is also reflected diachronically, with concrete originating senses developing more abstract extended meanings over

the course of language change. For example, prepositions for abstract spatial relations have been developed from concrete body part terms (Heine, 1997; Heine et al., 1991), not vice versa.¹

In the present study, we present a series of analyses of an existing dataset of cross-linguistically attested semantic changes involving source-target asymmetries. We test the dominant view—and one of the central tenets of conceptual metaphor theory (Lakoff & Johnson, 1980, p. 207)—that asymmetries are based on differences in concreteness between the source and target domains of mappings (see discussion in Dunn, 2015; Porat & Shen, 2017). Dancygier and Sweetser (2014, p. 62) write that “the idea that metaphor treats abstract concepts in terms of concrete ones is a recurring theme in metaphor studies.” The notion that metaphoric asymmetry is rooted in concreteness is also codified in textbooks, as when Kövecses (2002, p. 6) writes that “our experiences with the physical world serve as a natural and logical foundation for the comprehension of more abstract domains.” According to Kövecses, this asymmetry in concreteness is “why in most cases of everyday metaphors the source and the target domains are not reversible.”

The role of concreteness in theories of figurative language is often presumed, but rarely tested empirically (but see Dunn, 2015; Reijnierse et al., 2019; Xu et al., 2017). It is possible that there are other factors that lie at the root of metaphoric

¹ We use ‘source’ and ‘target’ as synonymous with ‘originating sense’ and ‘extended sense’ throughout this paper, but we do not use these terms to imply that semantic changes necessarily reflects intentional, novel extensions of words on the part of speakers.

asymmetries, beyond concreteness (see also Murphy, 1997). For example, Bottini and Casasanto (2013) characterize asymmetry in terms of “accessibility,” stating that the source domain may be relatively more familiar, perceptually available, imageable, memorable, and so on. Inherent in this characterization is an openness to considering that there may be multiple different dimensions that may characterize the distinction between sources and targets. In a similar vein, Radden and Kövecses (1999) provide a long list of different motivations for asymmetries in metonymy, of which concreteness is only one dimension among others, such as importance, immediacy, stereotypicality, etc.

As another alternative source of asymmetries, consider Dancygier and Sweetser (2014), who reconceptualize concreteness in terms of the notion of intersubjective accessibility. More intersubjectively accessible concepts are defined as those concepts that can be accessed by multiple speakers, as well as shared between them. Concrete objects that are manifested in the external world are perceptually accessible to a person’s individual visual system, and what can be seen by one person can also be seen by others, allowing for the establishment of joint attention (Tomasello & Farrar, 1986) on the part of multiple individuals (San Roque et al., 2015). Accordingly, Dancygier and Sweetser (2014) suggest that metaphors may commonly be based on physical source domains not because of the concreteness of these domains, *per se*, but instead because the physical world often allows for more shared intersubjective access than more abstract or internal concepts.

Similarly, Traugott (1989) and Sweetser (1990) argue that in the development of words for mental concepts, the directionality of metaphor goes from the external to the internal and subjective mental world (as when using “I see” to mean “I understand”) because things within the external world that can be perceived by our sensory systems are also more intersubjectively accessible (see also San Roque et al., 2015; for empirical evidence see Xu et al., 2017). Relatedly, Thibodeau and Durgin (2008, p. 531) suggest that when describing emotion concepts, which are internal and subjective, language often employs metaphoric source domains that involve “shared concrete referents that can stand in for private ones.”

Concreteness also appears to be an insufficient explanatory concept when it comes to explaining mappings between conceptual domains that are similarly concrete. For example, Blust (2005) points out that in many Austronesian languages, lungs are frequently talked about in terms of livers, as in the Lakalai language, which uses the expression *la hate malailai* ‘withering liver’ to describe lungs, or in the Lau language, which uses the expression *sae-ŋisu* ‘spittle-covered liver/frothy liver.’ Blust’s evidence clearly points to a linguistic asymmetry between livers and lungs, however, it is hard to see how this asymmetry could be explained with reference to concreteness, as livers and lungs appear to be similarly concrete. Relatedly, while notions of embodiment and concreteness may help to explain why body part terms are often used to talk about abstract concepts (e.g., Heine & Kuteva, 2002), they cannot easily explain asymmetries in metonymic extensions between body parts,

such as the fact that expressions denoting the concept ‘hand’ are often extended to denote ‘arm’, rather than the reverse (Wilkins, 1996).

Another set of changes that are difficult to explain by invoking concreteness are those involving perception verbs. Based on a detailed analysis of cross-linguistic conversation data, San Roque et al. (2018) suggest that extensions between different sensory modalities (such as from vision to smell) may be mediated by more abstract cognitive meanings. In the case of semantic shifts between sensory modalities, both sources and targets are sensory (Strik Lievers, 2017, p. 114); thus, concreteness may not be an appropriate reference concept for explaining the motivation behind these shifts.

1.1 The present study

The preceding discussion calls into question whether concreteness explains metaphoric and metonymic asymmetries, and whether other factors might better explain them. Recent work has begun to address this question empirically. For example, Winter (2019) provides evidence that in perceptual metaphors, emotionally valenced adjectives are more likely to be used as sources (e.g., “sweet melody,” “sour note” etc.), compared to relatively more neutral adjectives. Littlemore et al. (2018) find that for a set of ‘A is B’ type nominal metaphors (such as “lawyers are sharks,” “education is a stairway”), concreteness differences between source and target are less predictive of metaphor goodness judgments than differences in word frequency and emotional valence. Finally, Xu et al. (2017) find that concreteness does

not explain asymmetries in how English words have gained new metaphoric senses over time beyond other predictors (externality, embodiment, and emotional valence).

Building on these previous case studies—which have focused on a restricted set of semantic domains (e.g., perception) and languages (English)—we ask how well concreteness explains asymmetric semantic changes attested in a typologically balanced sample of 149 different languages, involving nominal concepts from multiple semantic domains (Urban, 2011). This dataset includes a large number of changes that are metaphoric, metonymic, or both. Our primary analyses assess whether concreteness judgments from English speakers (Brysbaert et al., 2014) can be used to distinguish sources and targets for these semantic changes.

Our investigation concerns not only semantic changes that are metaphoric, but also within-domain or ‘within-frame’ changes characterized by contiguity relations or associations within the same semantic frame, specifically, metonymy (Littlemore, 2015; Srinivasan & Rabagliati, 2015). Contemporary cognitive linguistic theory suggests that many metaphors are based on metonymies (Barcelona, 2003; Kövecses, 2013), and factors such as concreteness have also been hypothesized to motivate metonymic mappings, the same way that they have been used to motivate metaphoric mappings (see discussion in Littlemore, 2015). Our analysis seeks principles that unify metaphor and metonymy.

In addition to assessing the role of concreteness in predicting cross-linguistically attested asymmetries, we also explore the role of word frequency by

testing whether meanings that are more frequent—as indexed by English word frequency metrics—are more likely to be cross-linguistic sources than targets. It is well-known that words that are more frequent are also more likely to carry multiple senses of meaning (Piantadosi et al., 2012; Zipf, 1945), opening the possibility that a more frequent word meaning will be more likely to be a source of subsequent extension. This could be because more frequent word forms are particularly easy for speakers to access when communicating new meanings (Harmon & Kapatinsky, 2017). Alternatively, more frequent words are also more familiar (Noble, 1954) and widely known (Keuleers et al., 2015), providing ideal sources for meaning extension. Thus, frequency could be a proxy of intersubjective accessibility. Our purpose here was to assess whether word frequency predicts cross-linguistic asymmetries in meaning extension, leaving open the precise mechanism for future work.

Our analyses employ predictive measures from English (and English speakers) as opposed to measures based on other languages, due to the accessibility of large-scale data. Although it can be problematic to rely on English-based measures in research involving other languages (Wierzbicka, 2014), the semantic changes that our analyses focus concern a set of languages that is typologically balanced; thus, English is no more similar or different to the languages in our dataset than any other language we could have chosen. This makes it unlikely that we would obtain a different pattern of results by relying on measures from another language.

2. Dataset of cross-linguistically attested semantic changes

In this paper, we explore cross-linguistically attested, asymmetrical semantic changes. Our analyses address how well concreteness ratings and word frequency measures explain the difference between the sources and targets of these attested semantic changes. Importantly, we focus on concepts that are already semantically related. We consider, for example, the concept pair ‘skin ~ bark’. In many different languages, the concept ‘bark’ (of a tree) is a morphologically complex expression that is based on the concept ‘skin’ (Urban, 2011). Obviously, skin and bark are semantically similar as they both refer to outer layers of bodies. Our account is not focused on asking why it is specifically skin that is used to talk about bark (Xu et al., 2020), which requires understanding what it is about skin and bark that makes the two concepts mappable. Instead, we ask the question: given a semantically related conceptual pair, why is one concept a more likely source (‘skin’) than the other (‘bark’)?

The dataset we analyze comes from Urban (2011), who compiled dictionary entries for 160 different nominal concepts in a typologically balanced sample of 149 languages. The word list was constructed with inspiration from existing lists of basic vocabulary, with the goal of getting a large number of concepts that can be assumed to be familiar in most parts of the world today to maximize cultural neutrality and that are likely going to be represented in the dictionaries of the respective languages (cf. Urban, 2012, pp. 47–48 for more detail). The languages were sampled to construct a variety sample that includes as many distinct language families as

possible and at least one language per language family from Dryer's (2005) classification of the World Atlas of Language Structures (Urban, 2011; for more extensive justification of the sampling method, see Urban, 2012, pp. 48–53).

A set of 47 concept pairs was discovered that exhibit strong evidence for asymmetry, such as the concept pair 'skin ~ bark' discussed above. Using this concept pair as an example, Urban (2011) used the following criteria to operationalize asymmetry (e.g., to conclude that skin is a source and bark is a target):

- (1) At least two languages from different language families need to have a lexical expression that means both 'skin' and 'bark' (= polysemy, co-lexification).
- (2) At least two languages from different language families need to have a lexical expression where 'bark' is morphologically marked, using lexical material from the source 'skin.'
- (3) No language in the sample has a morphologically complex form for 'skin' that is based on 'bark.'

Criterion (1) assures that the two concepts are related. Criterion (2) and (3) assure that they are related in an asymmetrical fashion. For example, the Yuki language (an Native American language that is now extinct) had an expression for 'bark' which was *ʔol šil*, where *ʔol* meant 'tree' and *šil* meant 'skin' (Sawyer &

Schlichter, 1984; discussed in Urban, 2011). Thus, the compound literally translates to ‘tree skin,’ a morphologically complex (marked) expression that uses ‘skin’ to talk about ‘bark.’ In Urban’s (2011) typological sample, there is no language that has an analogous expression where the concept ‘bark’ is used to talk about ‘skin,’ akin to ‘human bark.’ By inference then, we can conclude that for those languages where the concepts ‘skin’ and ‘bark’ are co-lexified (as in polysemy), ‘skin’ must have been the primary meaning, which was later extended to ‘bark.’ One historical pathway through which such polysemy can emerge is via a morphologically complex intermediate stage (‘tree skin’), which is subsequently simplified by dropping one morpheme (‘skin’), ultimately resulting in the word form for both concepts being the same (Urban, 2011).

It is worth pointing out that for polysemy, Urban’s argument rests on an inference: for polysemous forms, the directionality of change is presumed to go from skin to bark *because this mirrors the asymmetry that is observed in the morphologically complex forms* (‘tree skin’ is observed, but not ‘human bark’). However, even if we do not accept this indirect inference as evidence for directionality in the development of such polysemy, the dataset also directly speaks to semantic changes even if we only consider the morphologically complex forms, which always involve descriptions of target concepts in terms of their sources (‘tree skin’), but never the reverse. In other words, the morphologically complex forms by themselves are a form of change, as they provide evidence that one form was derived from the other in word formation.

Note that for the sample considered in Urban (2011), the concept pairs are asymmetrical without exception. Winter et al. (2015) define “directionality” as the case where sources influence targets, but *never* the reverse, contrasting it with “asymmetry,” defined as when the source merely exerts a stronger influence onto the target, which allows for the possibility of reverse mappings. Although there are no reverse mappings in the Urban (2011) dataset (directionality based on the above distinction), we use the term asymmetry to describe these semantic changes as it is possible that examples from other languages may be found that present exceptions to the directionality observed within the Urban (2011) dataset. Crucially, even if a new analysis of additional languages were to uncover individual languages contravening the patterns established here, the focus on patterns without exceptions in Urban’s sample makes it unlikely that the overall asymmetries within these conceptual pairs would be overturned.

Since the original publication, Matthias Urban has added an additional 24 concept pairs to the list (personal communication). Table 1 presents a full list of the 71 conceptual pairs that form the basis of our analysis, with the cross-linguistic source listed to the left of the cross-linguistic target.

cloud~fog/mist	breast~milk	tongue~flame	feather~pen
		road/street/way~Milky	
sun~moon	mouth~lip	Way	feather~beard
grass~straw/hay	belly/stomach~womb	bed~nest	star~meteoroid
smoke~fog/mist	heart~belly/stomach	egg~testicle	river~Milky Way
steam~fog/mist	milk~nipple	sun~clock	mountain~valley
animal~bird	liver~lungs	seed~testicle	puddle~swamp

lake~swamp	car~train	shadow~mirror	sky~rain
smoke~dust	heart~kidney	mouth~estuary	river~spring
smoke~cloud	mirror~glasses	bird~airplane	river~river bed
tree~branch	heart~lungs	foam~lungs	moon~star
ashes~embers	(molar) tooth~jaw	ball~testicle	pus~brain
tree~forest	belly/stomach~navel	day~clock	eyebrow~eyelash
day~dawn	cheek~buttocks	bone~needle	eyelid~eyelash
flower/blossom~bud	mouth~cheek	day~cloud	finger nail~finger
day~noon	skin~bark	tooth~beak	eyeball~pupil
sun~noon	mouth~beak	seed~bud	
honey~wax	saliva/spittle~foam	dew~fog	
bone~horn	house~nest	soil/earth~dust	
river/stream~flood	mouth~estuary	flame~embers	

Table 1. 71 concept pairs that exhibit cross-linguistic asymmetries in semantic

change based on Urban (2011) and a subsequent extension of this list by Matthias

Urban; sources are listed to the left of targets. Translation equivalents for a particular concept are separated by a '/' (e.g., 'fog/mist').

The concept pairs are drawn from a number of different conceptual domains. Some involve *intra-field* associations (cf. Matisoff, 1978), e.g., among nature-related concepts, as in the pairs 'sun~moon,' 'animal~bird' and 'honey~wax.' Other intra-field associations are among human-related concepts, as in the pairs 'breast ~ milk,' 'mouth ~ lip,' or 'milk ~ nipple.' There are also *inter-field* concept pairs, e.g., where concepts from the human sphere are mapped onto nature-related concepts, as in 'skin ~ bark,' 'mouth ~ beak' or 'house ~ nest,' or where concepts are mapped from the natural sphere to the human sphere, e.g., 'bird ~ airplane,' 'foam ~ lungs' and 'eggs ~ testicle.'

Traugott and Dasher (2001, p. 79) say that “metaphor involves similarity, mapping of one domain onto another, in what is essentially a conceptual mode of analogy, paradigmatic choice, and iconicity,” whereas “metonymy involves contiguity, syntagmatic relations, and indexicality, and shifts within the same domain.” The dataset that we worked with was not collected with applications to metaphor and metonymy in mind (and in fact, Urban (2011) does not focus on figurative language). Given this, there are a number of different kinds of semantic relations represented, including taxonomic relations such as ‘animal ~ bird.’

However, a quick look at the list reveals that the majority of cases involve mappings based on similarity, contiguity, or both. For example, the concept pair ‘skin ~ bark’ passes the ‘is like’ test to identify metaphors (Gibbs, 2017), e.g., it makes sense to say that bark is like skin, given that both are outer layers of organisms. Other similarity-based mappings include ‘eggs ~ testicle,’ ‘sun ~ moon,’ and ‘tongue ~ flame’ (which sometimes have similar shapes). Other cases pass the ‘stands for’ test (Gibbs, 2017), or alternatively, Urban’s (2012) ‘has something to do with’ test, which helps to identify metonymies. For example, ‘breast’ has something to do with ‘milk’ (‘breast ~ milk’), ‘spring’ is where the ‘river’ comes from (‘river ~ spring’), and ‘clock’ and ‘day’ are related via causal/temporal connection (‘clock ~ day’). Other cases appear to have both a metaphoric and a metonymic component, as suggested by the fact that they pass both tests, e.g., mappings between different body parts, such as ‘liver ~ lungs’ and ‘heart ~ kidney.’ For these conceptual pairs, there is an element of similarity (the body organs have similar shapes and colors, and also perform similar

functions, namely, bodily functions); and there is an element of contiguity as the body parts are physically connected.

The diversity of relations among the conceptual pairs we analyzed is a strength of our study, as it presents a challenge for identifying a unifying principle that explains asymmetry in semantic change. For example, one could invoke differences in embodiment to explain some asymmetries (such as ‘skin ~ bark,’ where a human body part is extended to a plant), but not others (such as ‘feather ~ beard,’ where a concept relating to animals is extended to a human body part). Similarly, one could invoke taboo avoidance (Allan & Burridge, 2006) to explain examples of asymmetries like ‘eggs ~ testicle,’ but this explanation would not account for many of the other examples. Another interesting comparison is between ‘breast ~ milk’ and ‘river ~ spring.’ In the former, the actual, real-world source of milk is also the cross-linguistic source; in the latter case, the real-world source appears as the target (i.e., rivers emerge from springs). Thus, it is hard to find unifying generalizations for these examples based on semantic analysis alone. Our analyses address whether concreteness is a possible correlate of the source ~ target asymmetries observed in these different kinds of semantic mappings, and if not, whether there are other correlates.

In the following sections, we assess how useful concreteness is in characterizing the source ~ target asymmetry, and we subsequently assess the role of word frequency.

3. Analysis

3.1. Assessing the role of concreteness

A quick look at the concepts shown in Table 1 would already suggest that concreteness may not provide an adequate way of distinguishing sources from targets. Intuitively, skin, feathers, and eggs (attested sources) are not more or less concrete than bark, beards, or testicles (attested targets). To more rigorously assess the role of concreteness, we used a large-scale and widely used dataset of concreteness ratings of English words collected from human subjects by Brysbaert and colleagues (Brysbaert et al., 2014). In using these concreteness ratings, we also adopted a clear operational definition of concreteness. Dunn (2015, p. 260) states that although concreteness is crucial to conceptual metaphor theory, “a principled definition of abstractness versus concreteness is never provided.” In Brysbaert et al. (2014, p. 906), participants who provided ratings were told that words for concrete concepts “refer to things or actions in reality, which you can experience directly through one of the five senses,” whereas words for abstract concepts “refer to meanings that cannot be experienced directly.” Thus, by using Brysbaert et al. (2014) concreteness ratings, we defined concreteness/abstractness in terms of accessibility to the senses (see Dunn, 2015 for other interpretations of concreteness). Brysbaert and colleagues’ dataset contains 40,000 English words rated on a five-point scale from 1 (most abstract) to 5 (most concrete).

For our first analysis, we began by matching the English concept names given in Urban (2011) with their corresponding concreteness ratings. Clearly, meanings do

not translate perfectly across languages, but some degree of translational equivalency is already assumed in the dataset by Urban (2011), i.e., if a dictionary lists that the meaning of a word is 'ball' in one language, then this meaning is assumed to have considerable overlap with the word for 'ball' in the dictionary of another language. Moreover, it is likely that the coarse meanings we were interested in here are relatively stable across languages with respect to the concrete/abstract dimension. For example, the English concepts 'egg' and 'liver' received high concreteness ratings in the Brysbaert et al. (2014) data. While individual languages may differ with respect to the perceived abstractness or concreteness of these concepts, it seems unlikely that a large number of the 149 languages in the typological sample collected by Urban (2011) would judge these concepts to be very abstract.

For source or target concepts that had multiple translational equivalents (such as fog/mist and road/street/way), we computed the average concreteness ratings and used these averaged ratings to quantify the concreteness of the corresponding source or target concepts. For 132 out of the 142 individual concepts (71 pairs, sources and targets together), concreteness data was available (93%). All analyses were conducted in R (R Core Team, 2018), using the "tidyverse" package version 1.2.1 for data processing (Wickham, 2017), the "effsize" package 0.7.6 for computing effect size (Torchiano, 2019), and the "brms" package 2.9.0. for Bayesian regression (Bürkner, 2017). The data and code for all of the analyses presented below are

publicly accessible under the following Open Science Framework repository:

<https://osf.io/qukgh/>

Across the conceptual pairs from Urban (2011), sources and targets were fairly similar to each other with respect to concreteness. On average, source words had slightly higher concreteness ratings (mean rating = 4.74, $SD = 0.28$) than target words ($M = 4.68$, $SD = 0.43$). Effect size measures (paired Cohen's d) revealed that this difference was small ($d = 0.20$, 95% confidence interval: [0.11, 0.51]).

Furthermore, there was no evidence for the concreteness difference between source and target words being statistically reliable ($t(61) = 1.32$, $p = 0.19$). When comparing concreteness ratings within each pair, only 56% of all pairs had more concrete sources than targets (35 out of 62 pairs for which concreteness data exists for both sources and targets). In sum, the source ~ target asymmetries in the Urban (2011) data are not well characterized by concreteness differences, at least as operationalized via English concreteness ratings.

3.2. Assessing the role of word frequency

As we did for the concreteness ratings, we matched translational equivalents from Urban (2011) with word frequency data, taken from the spoken section of the Corpus of Contemporary American English (COCA; Davies, 2010), which uses transcripts of unscripted conversations from more than 150 different TV and radio programs (<https://www.english-corpora.org/coca/>). For concepts that had multiple

translational equivalents (e.g., road/street/way) we computed average frequency across these terms.

Overall, word frequencies patterned much more strongly with the observed source ~ target asymmetries than concreteness ratings. English words for concepts corresponding to cross-linguistic sources were more frequent (mean log₁₀ frequency of 3.29, $SD = 1.03$) than English words corresponding to cross-linguistic target concepts ($M = 2.41$, $SD = 1.13$). Cohen's d shows this to be a large-sized effect ($d = 0.81$, 95% confidence interval: [0.47, 1.14]). In raw frequencies, source words occurred on average 12,078 times, compared to target words, which occurred on average 1,277 times. When comparing frequencies within each pair, 70% of all pairs had more frequent sources than targets (50 out of 71), which was a considerably higher percentage than what we found for concreteness (56%). These results were surprisingly robust across different registers, including other subsections of the COCA and of the British National Corpus (cf. McEnery et al., 2017), ranging from a lowest effect size of $d = 0.75$ for COCA news to the highest for BNC spoken $d = 0.86$.²

² At a reviewer's suggestion, we assessed how well the cross-linguistic asymmetry was predicted by a measure of word prevalence (Brysbaert et al., 2018), since prevalence of a word within a community may be a good measure of how much that meaning is in the common ground (and thus how intersubjectively accessible it is; Keuleers et al., 2015). We also assessed the role of other variables that could potentially be relevant for psycholinguistic processing, including comprehension speed from the English Lexicon Project and British Lexicon Project (Balota et al., 2007; Keuleers et al., 2012), sensory experience ratings (Juhasz & Yap, 2013), sensorimotor ratings (Lynott et al., 2019), emotional valence ratings (Warriner et al., 2013), age of acquisition ratings (Kuperman et al., 2012), and contextual diversity from the COCA (Davies, 2010) and BNC (McEnery et al., 2017). Computing effect size for each of these variables with Cohen's d revealed that word frequency

As noted before, it is unlikely that our reliance on English-based measures of concreteness and frequency biased the outcome of our analyses, which involved data from a typologically-balanced set of languages (such that English is no more likely to be an “outlier” than any other language). Still, to address whether word frequency measures derived from other languages might also explain the observed cross-linguistic asymmetries, we leveraged Google Books, for which frequency data is available for Chinese, English, French, German, Hebrew, Italian, Russian, and Spanish.

Some caveats are worth mentioning when considering analyses based on Google Books. First, prior work suggests that inferences about cultural change are difficult to make from this corpus, due to the changing composition of the registers included over time (e.g., the sudden inclusion of more scientific texts; Pechenick et al., 2015). However, since we are not interested in making inferences about cultural change over time, the concerns raised by Pechenick et al. (2015) are less relevant to the current study. Second, because the same books are likely to be hosted in the corpora for different languages (in translated form), frequency counts for these different languages are unlikely to be independent. These limitations aside, these corpora provide some of the best available comparable word frequency data for languages beyond English.

differences (log10 differences) most strongly separated cross-linguistic sources from targets (see online supplementary materials: <https://osf.io/qukgh/>).

We looked up the Google Books frequencies of the translation equivalents for all of the concepts shown in Table 1. Concepts were translated automatically using Google Translate and we discarded translations that did not successfully back-translate (e.g., if the translation from English → Chinese → English resulted in a different word). The frequency data used by Google Ngram is *relative* frequency data (normalized by the entire corpus), which were further log10-transformed and z-scored within language so that the different languages are on comparable metrics. Google Ngram includes data from the beginning of the 19th century up to modern times, however, older data tends to be very sparse and unreliable for languages other than English, which is why we only consider frequencies since 1980.

We fit a mixed logistic regression to this data, with the fixed effect “log relative frequency” (z-scored) and random intercepts ‘concept pair’ (e.g., ‘skin ~ bark’ would be one pair, with two data points, one for ‘skin’ and one for ‘bark’) and for ‘language.’ We also included by-concept pair and by-language varying random slopes for the frequency effect; in other words, the model does not assume that concept pairs or languages are equal with respect to frequency differences between sources and targets. The model was estimated as a Bayesian model.³

³ The model formula was: `brm(SourceVsTarget ~ Frequency + (1 + Frequency | ConceptPair) + (1 + Frequency | Language))`; MCMC sampling was performed with 8,000 iterations for 4 chains (4,000 warmup samples). We set a weakly informative prior on the beta coefficient (a normal distribution centered at 0 with SD = 1) (cf. Lemoine, 2019). Convergence was successful (all Rhat values = 1.0).

Overall, there was a strong effect of frequency: for each increase in relative log frequency by one standard deviation, the log odds of being unmarked rose by +2.78 ($SE = 0.46$, odds: 16.12 to 1). The fact that the Bayesian 95% credible interval of this coefficient is far away from and does not include zero [1.95, 3.75] suggests that the effect is reliable. The posterior probability of the effect of frequency on the source ~ target asymmetry being below zero is: $p = 0$. As the model controls for variation in this effect across languages (via the by-language random slope term), this result can be seen as generalizing across all the eight languages in the Google Ngram sample. The frequency effect (log odd coefficient from the logistic regression model) is shown for all languages in Figure 1. As can be seen, none of the 95% Bayesian credible intervals for the respective languages overlap with zero (in fact, the posterior probability of the frequency effect being below zero is $p = 0$ for all languages).

Figure 1. Logistic regression coefficients (positive means more frequent = more likely source) from the mixed effects model, separately for each language; as can be seen, for all languages the posterior means of the frequency effect are well above zero; the 95% credible interval for each coefficient does not include zero for any of the languages

4. General discussion

4.1. Summary of findings

Why are some concepts more likely to be the sources of semantic change, and others more likely to be the targets? Asymmetries in metaphoric thought and language are generally characterized in terms of concreteness, with concrete sources

being used to communicate about abstract targets. This is one of the main reasons why conceptual metaphor theory is thought to be an embodied theory of language and cognition, as metaphor is seen as a way of grounding abstract thinking in concrete conceptual domains that directly relate to the body, and to our senses. Yet this idea that metaphoric mappings go from the concrete to the abstract has rarely been tested empirically or compared with alternative explanations (Dunn, 2015; Xu et al., 2017).

Here, we investigated a dataset of cross-linguistically attested semantic changes collected by Urban (2011), who assembled a list of 71 concept pairs which exhibit strong asymmetries between source and target domains. These asymmetries cover both metaphoric mappings and metonymic mappings, among other conceptual connections. Our first key finding is that concreteness was *not* a good predictor of these asymmetries. Although the concepts in our dataset exhibited strong cross-linguistic asymmetries, they were each relatively concrete. This finding points to the limits of concreteness in explaining asymmetries in metaphor, metonymy and semantic change.

This finding is consistent with other work that does not find concreteness to be a primary factor in explaining metaphoric asymmetries (Xu et al., 2017). For example, just as ‘skin’ and ‘bark’ are equally concrete in our cross-linguistic dataset, perceptual metaphors (“rough taste,” “loud color,” “sweet melody” etc.) involve mappings between two perceptual—and thus relatively concrete—domains (Strik Lievers, 2015; Ullmann, 1959; Williams, 1976; Winter, 2019). Likewise, metonymic

semantic changes among different body part terms (Blust, 2005; Wilkins, 1996) involve mappings among concrete concepts. Together with the results of our analysis, the existence of these cases suggests that as a sole explanatory concept, concreteness is too limiting.

Our second key finding is that concreteness differences between cross-linguistic sources and targets were dwarfed by source-target differences in word frequency. Specifically, we found that English words corresponding to the sources of cross-linguistic semantic change were more frequent than English words corresponding to the targets of semantic change. Similar results were obtained for frequency measures from other majority languages spoken in industrialized cultures (Chinese, Hebrew, German, French, Spanish, Italian, Russian). The fact that asymmetries in word frequencies in English (and the other languages we examined) corresponded so well to a large cross-linguistic sample of asymmetric semantic changes was striking. This finding extends results from Littlemore and colleagues (Littlemore et al., 2018), who found that for nominal “A is B” type metaphors, such as “lawyers are sharks” or “education is a stairway,” frequency asymmetries better predicted metaphoric goodness judgments than concreteness asymmetries. In the next section, we explore potential explanations for why frequency may predict source-target asymmetries.

4.2. Why word frequency?

Here, we consider multiple alternative explanations for the correlation we found between the frequency of a word and whether its associated meaning is the cross-linguistic source or target of semantic change. Our discussion focuses on intersubjective accessibility (§4.2.1), cognitive accessibility (§4.2.2), and referential abundance and cultural salience (§4.2.3). These different accounts reflect a shift away from focusing on what asymmetry reveals about concepts and embodiment toward considering what it might reveal about pragmatics and communication (§4.2.1), the dynamics of lexical access (§4.2.2), and the impact of the environment on lexical structure (§4.2.3). The explanations themselves are not mutually exclusive and further research is needed to disentangle their relative contributions (§4.2.4).

4.2.1. Intersubjective accessibility

In all communication, interlocutors need to track that what is within the common ground, and this applies to figurative uses of language like metaphor and metonymy. Thus, when it comes to generating new metaphoric or metonymic meanings, communication may be more successful if interlocutors use source domains that are more familiar, known and sharable (e.g., objects in the shared physical environment), as opposed to source domains that are less intersubjectively accessible (Dancygier & Sweetser, 2014; cf. San Roque et al., 2015).⁴ From this

⁴ Just as the demands of communication may constrain semantic change, so may the demands of language acquisition. The process of learning a new word for an abstract or subjective concept could be simpler if children have already learned a prior, meaning for that word that is more inter-subjective (Srinivasan & Carey, 2010; Starr,

perspective, it makes sense that words that are more frequent are more likely to be cross-linguistic sources because word frequency is correlated with perceived familiarity (Noble, 1954) and how many people know a word (Keuleers et al., 2015).

As discussed in the Introduction, an appeal to intersubjective accessibility helps make sense of a variety of cases of semantic extension. One notable example is the fact that vision is a common source for semantic change (Sweetser, 1990), as in the case with perception verbs (San Roque et al., 2018; Viberg, 1983). For example, in the Australian aboriginal language Walpiri, the verb *nyani* ‘to see, look at’ is used as the basis for ‘to hear, to listen’ (*purda-nyanyi*) and ‘to smell’ (*parnti-nyanyi*) (Evans & Wilkins, 2000, p. 557). In contrast to vision, taste and smell are more subjective, with much research attesting to the intersubjective variability of perception in these latter two modalities (Winter, 2016). Being so intersubjectively variable, it makes sense that taste and smell are less likely than vision to serve as sources for semantic extension in the domain of perception verbs. Of course, this doesn’t mean that taste and smell are never used as sources for metaphors across cultures (see, e.g., O’Meara & Majid, 2020). For example, smell is used in English metaphors for suspicion. Here too, however, there is a relative difference in intersubjective accessibility between source

Cirolia, Tillman & Srinivasan, 2020). For example, if a child has already learned the spatial sense of *high* – which denotes a relatively inter-subjective concept – this could help them understand what *long* might mean when it is metaphorically used to refer to time. This predicts that metaphors that employ intersubjective concepts as source domains will be preferred because they scaffold the acquisition of otherwise difficult-to-learn word meanings.

and target domains: When somebody says “something smells fishy” to describe their suspicion (Sweetser, 1990), this ‘smell’ is more intersubjectively shared than the relatively more subjective and internal idea of suspicion.

4.2.1. Cognitive accessibility

A different explanation for why word frequency may predict source-target asymmetries is that, when a speaker communicates a new meaning, they will be more likely to select a word form that is easy to access. Critically, the emphasis here shifts away from how a speaker integrates information about what is in common ground with a listener (the intersubjectivity account discussed above) toward what is easiest for the speaker to produce. Prior research shows that word frequency is a strong predictor of psycholinguistic measures of accessibility, including how likely a word will be recalled in a memory task (for review, see Gregg, 1976). Thus, it is possible that when speakers come up with novel expressions for concepts, including metaphoric and metonymic ones, more frequent words will come to mind more quickly; accordingly, more frequent words may be more likely to provide sources for semantic extension than less frequent words.

Consistent with this proposal, Harmon and Kapatsinski (2017) performed a series of experiments which showed that participants were more likely to semantically extend a more frequent word form than a less infrequent word form. The authors interpreted this effect with reference to cognitive accessibility (p. 30): “The retrieval demands inherent to production can push the speaker to produce the

most accessible form to express a meaning, even when a less accessible form would be a better cue to the meaning for the same speaker.”

4.2.3. Referential abundance and cultural salience

A final possibility is that words that are more frequent correspond to referents in the environment that are also more abundant. These words that denote abundant kinds of referents might also be expected to develop early in a language’s history—and serve as a source of semantic extension.

There is some support for the first component of this proposal (that more frequent words reflect more abundant or culturally salient referents). Clearly there are many things we frequently encounter without also verbalizing them more often, but all else being equal, several studies have found correlations between real world abundance and word frequency. For example, Ladle et al. (2019) show that the scientific names of birds with wider geographic ranges that also overlap with urban settlements are more likely to be mentioned on the internet. And Regier et al. (2016) show that regions with colder climates in the US are more likely to mention the English words *ice* and *snow*.

There is also support for the second component of this proposal, according to which words for abundant referents develop earlier in a language. For example, the study by Ladle et al. (2019) suggests that birds that are more frequently encountered by humans (based on their geographic range and proximity to urban areas) are also more likely to have been scientifically described and named earlier. A similar result

was obtained for indigenous Zapotec bird names by Abbott and Kemp (2020), who showed that more abundant birds in Mexico were also more likely to have folk names (see also Gosler, 2017). These results could help make sense of our findings. Words for more abundant referents—corresponding to more frequent words—could be more likely to develop first. And concepts with less abundant referents may be more likely to have derived names based on those concepts that have been named first.

Clearly however, word frequency cannot be reduced to real-world abundance (e.g., the nightingale is a relatively infrequent bird that is culturally very salient and hence frequently talked about). It has been shown that factors other than abundance determine word frequency, such as in the case of birds, how visually salient they are (Ladle et al., 2019). And within the anthropological and ethnobiological literature, there is active discussion about how other factors, such as the utility of particular concepts (such as whether birds can be hunted or eaten) contribute to cultural salience (Hunn, 1982).

Indeed, a utility-based argument for cultural salience has been put forth by Blust (2005) for one of the concept pairs in our dataset, ‘liver ~ lungs.’ As discussed in the Introduction, Blust (2005) observed that between these two concepts, complex words for lungs are often based on the word for liver but never the reverse in the Austronesian language family. Blust suggests that this asymmetry may be grounded in the cultural importance of livers, as opposed to lungs. For example, livers are eaten in many cultures, more so than lungs.

4.2.4. A complex web of causal relations

As should be apparent from the discussion above, word frequency is embedded in a complex web of causal relations (cf. Ladd et al., 2015). Word frequency could be a *cause* of the observed asymmetries — compatible with both an intersubjective accessibility (§4.2.1) and cognitive accessibility (§4.2.2) account. Alternatively, word frequency could be a *consequence* of the factors that ultimately drive asymmetries in semantic change, as in the case of real-world abundance or cultural salience (§4.2.3).

Experiments are needed to disentangle these factors. For example, Harmon and Kapatsinski (2017) modulated the frequency of word forms seen by participants, which influenced subsequent semantic extensions. This task could be expanded to include a communicative condition (word frequencies are shared between participant and an experimental confederate) in contrast with a non-communicative control condition (manipulating word frequencies only for the participant). In addition, there is need for a deeper understanding of what causes word frequency distributions more generally (cf. Piantadosi, 2014), including how different factors such as extralinguistic abundance and various measures of salience interact in driving this measure. Regardless of what explains the relationship between word frequency and asymmetry in semantic change, the strength of this effect (i.e., ‘large’ effect size based on Cohen’s *d*, accounting for 70% of all cases) calls out for further investigation.

Finally, it has to be pointed out that an important limitation of the current work is its reliance on English or, in the case of the Google Ngram data, on other languages from large industrialized societies. We strongly believe that the close match in word frequency data taken from these languages and the directionality of change here is suggestive of deeply rooted commonalities in shared experience with respect to the particular set of concepts considered in this study, but future work needs to look at word frequency data beyond English. In addition, concreteness rating studies have so far have almost exclusively been conducted in languages from industrialized cultures (e.g., Brysbaert et al., 2014; Guasch et al., 2016; Imbir, 2016). It is important to conduct concreteness ratings studies with languages from less industrialized cultures spoken by smaller communities for a whole range of concepts going beyond the set of words considered here to assess the degree to which the concept of ‘concreteness’ carries over across languages and cultures.

5. Conclusion

Our results provide a test of the classic idea that concreteness is the key factor that drives asymmetry in figurative language. We have advanced beyond existing research by not taking concreteness as a self-evident notion (cf. Dunn, 2015) and operationalizing it in a objective and reproducible manner using concreteness ratings (Brysbaert et al., 2014). Using data from Urban (2011), we show that concreteness asymmetries do not help in explaining the asymmetries observed in a large set of cross-linguistically attested metaphoric and metonymic semantic

changes. Surprisingly, we find that word frequency is a much stronger correlate of cross-linguistic asymmetries, suggesting that this is an under-explored factor in cognitive linguistic research on metaphor and metonymy.

Acknowledgements

We thank Matthias Urban and Graham Thompson for help with earlier versions of this project (Winter et al., 2013). We thank Márton Sóskuthy for helping us out with the Google Translate API. We thank the Evolang 2013 conference organizers for awarding the James Hurford prize to this work. MS is supported by awards from the National Science Foundation (SBE-16302040) and the James S. McDonnell Foundation.

References

- Abbott, J. T., & Kemp, C. (2020). Birds and Words: Exploring environmental influences on folk categorization. *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*, 744–749.
- Allan, K., & Burridge, K. (2006). *Forbidden words: Taboo and the censoring of language*. Cambridge University Press.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., & Treiman, R. (2007). The English lexicon project. *Behavior Research Methods*, 39(3), 445–459.

- Barcelona, A. (2003). On the plausibility of claiming a metonymic motivation for conceptual metaphor. In A. Barcelona (Ed.), *Metaphor and metonymy at the crossroads: A cognitive perspective* (pp. 31–58). Mouton de Gruyter.
- Black, M. (1979). More about metaphor. In A. Ortony (Ed.), *Metaphor and Thought* (pp. 19–43). Cambridge University Press.
- Blust, R. (2005). Liver and lungs: A semantic dyad in Austronesian languages. *Oceanic Linguistics*, 44(2), 537–543.
- Bottini, R., & Casasanto, D. (2013). Space and time in the child’s mind: Metaphoric or ATOMIC? *Frontiers in Psychology*, 4, 803.
- Brysbaert, M., Mander, P., McCormick, S. F., & Keuleers, E. (2018). Word prevalence norms for 62,000 English lemmas. *Behavior Research Methods*, 1–13.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904–911.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1), 1–28.
- Dancygier, B., & Sweetser, E. (2014). *Figurative language*. Cambridge University Press.
- Davies, M. (2010). The Corpus of Contemporary American English as the first reliable monitor corpus of English. *Literary and Linguistic Computing*, 25(4), 447–464.
- de Mendoza Ibáñez, F. J. R., & Masegosa, A. G. (2011). Going beyond metaphonymy: Metaphoric and metonymic complexes in phrasal verb interpretation. *Language Value*, 3, 1–29.
- Dingemanse, M. (2009). The selective advantage of body-part terms. *Journal of Pragmatics*, 41(10), 2130–2136.
- Dryer, M. S. (2005). Genealogical language list. *The World Atlas of Language Structures*, 601.

- Dunn, J. (2015). Modeling abstractness and metaphoricity. *Metaphor and Symbol, 30*(4), 259–289.
- Evans, N., & Wilkins, D. (2000). In the mind's ear: The semantic extensions of perception verbs in Australian languages. *Language, 76*(3), 546–592.
- Fauconnier, G., & Turner, M. (1998). Conceptual integration networks. *Cognitive Science, 22*(2), 133–187.
- Gibbs, R. W. (1994). *The poetics of mind: Figurative thought, language, and understanding*. Cambridge University Press.
- Gibbs, R. W. (2017). *Metaphor wars*. Cambridge University Press.
- Gosler, A. G. (2017). The human factor: Ecological salience in ornithology and ethno-ornithology. *Journal of Ethnobiology, 37*(4), 637–662. <https://doi.org/10.2993/0278-0771-37.4.637>
- Grady, J. (1997). Theories are buildings revisited. *Cognitive Linguistics, 8*(4), 267–290.
- Gregg, V. (1976). Word frequency, recognition and recall. In J. Brown (Ed.), *Recall and recognition* (pp. 183–215). Wiley Blackwell.
- Guasch, M., Ferré, P., & Fraga, I. (2016). Spanish norms for affective and lexico-semantic variables for 1,400 words. *Behavior Research Methods, 48*(4), 1358–1369. <https://doi.org/10.3758/s13428-015-0684-y>
- Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology, 98*, 22–44.
- Heine, B. (1997). *Cognitive foundations of grammar*. Oxford University Press.
- Heine, B., Claudi, U., & Hünemeyer, F. (1991). *Grammaticalization: A conceptual framework*. University of Chicago Press.

- Heine, B., & Kuteva, T. (2002). *World lexicon of grammaticalization*. Cambridge University Press.
- Hunn, E. (1982). The utilitarian factor in folk biological classification. *American Anthropologist*, 84(4), 830–847. <https://doi.org/10.1525/aa.1982.84.4.02a00070>
- Ibarretxe-Antuñano, I. (2008). Vision metaphors for the intellect: Are they really cross-linguistic? *Atlantis*, 15–33.
- Imbir, K. K. (2016). Affective norms for 4900 Polish words reload (ANPW_R): Assessments for valence, arousal, dominance, origin, significance, concreteness, imageability and, age of acquisition. *Frontiers in Psychology*, 7, 1081. <https://doi.org/10.3389/fpsyg.2016.01081>
- Juhász, B. J., & Yap, M. J. (2013). Sensory experience ratings for over 5,000 mono-and disyllabic words. *Behavior Research Methods*, 45(1), 160–168.
- Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods*, 44(1), 287–304. <https://doi.org/10.3758/s13428-011-0118-4>
- Keuleers, E., Stevens, M., Mandera, P., & Brysbaert, M. (2015). Word knowledge in the crowd: Measuring vocabulary size and word prevalence in a massive online experiment. *The Quarterly Journal of Experimental Psychology*, 68(8), 1665–1692.
- Kövecses, Z. (2002). *Metaphor: A practical introduction*. Oxford University Press.
- Kövecses, Z. (2013). The metaphor–metonymy relationship: Correlation metaphors are based on metonymy. *Metaphor and Symbol*, 28(2), 75–88.
- Kuperman, V., Stadthagen-Gonzalez, H., & Brysbaert, M. (2012). Age-of-acquisition ratings for 30,000 English words. *Behavior Research Methods*, 44(4), 978–990. <https://doi.org/10.3758/s13428-012-0210-4>

- Ladd, D. R., Roberts, S. G., & Dediu, D. (2015). Correlational studies in typological and historical linguistics. *Annual Review of Linguistics*, 1(1), 221–241.
<https://doi.org/10.1146/annurev-linguist-030514-124819>
- Ladle, R. J., Jepson, P., Correia, R. A., & Malhado, A. C. (2019). A culturomics approach to quantifying the salience of species on the global internet. *People Nat*, 1, 524–532.
<https://doi.org/10.1002/pan3.10053>
- Lakoff, G. (2008). The neural theory of metaphor. In R. W. Gibbs (Ed.), *Cambridge handbook of metaphor and thought* (pp. 17–38). Cambridge University Press.
- Lakoff, G. (2012). Explaining embodied cognition results. *Topics in Cognitive Science*, 4(4), 773–785.
- Lakoff, G. (2014). Mapping the brain’s metaphor circuitry: Metaphorical thought in everyday reason. *Frontiers in Human Neuroscience*, 8, 958.
<https://doi.org/10.3389/fnhum.2014.00958>
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.
- Lemoine, N. P. (2019). Moving beyond noninformative priors: Why and how to choose weakly informative priors in Bayesian analyses. *Oikos*, 128(7), 912–928.
<https://doi.org/10.1111/oik.05985>
- Littlemore, J. (2015). *Metonymy*. Cambridge University Press.
- Littlemore, J. (2019). *Metaphors in the mind: Sources of variation in embodied metaphor*. Cambridge University Press.
- Littlemore, J., Sobrino, P. P., Houghton, D., Shi, J., & Winter, B. (2018). What makes a good metaphor? A cross-cultural study of computer-generated metaphor appreciation. *Metaphor and Symbol*, 33(2), 101–122.
<https://doi.org/10.1080/10926488.2018.1434944>

- Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2019). The Lancaster Sensorimotor Norms: Multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods*, 1–21.
<https://doi.org/10.3758/s13428-019-01316-z>
- Matisoff, J. A. (1978). *Variational semantics in Tibeto-Burman: The “organic” approach to linguistic comparison*. ISHI Publications.
- McEnery, T., Love, R., & Brezina, V. (2017). Introduction: Compiling and analysing the Spoken British National Corpus 2014. *International Journal of Corpus Linguistics*, 22(3), 311–318.
- Murphy, G. L. (1997). Reasons to doubt the present evidence for metaphoric representation. *Cognition*, 62(1), 99–108.
- Noble, C. E. (1954). The familiarity-frequency relationship. *Journal of Experimental Psychology*, 47(1), 13–16.
- O’Meara, C., & Majid, A. (2020). Anger stinks in Seri: Olfactory metaphor in a lesser-described language. *Cognitive Linguistics*. <https://doi.org/10.1515/cog-2017-0100>
- Pechenick, E. A., Danforth, C. M., & Dodds, P. S. (2015). Characterizing the Google Books Corpus: Strong Limits to Inferences of Socio-Cultural and Linguistic Evolution. *PLOS ONE*, 10(10), e0137041. <https://doi.org/10.1371/journal.pone.0137041>
- Piantadosi, S. T. (2014). Zipf’s word frequency law in natural language: A critical review and future directions. *Psychonomic Bulletin & Review*, 21(5), 1112–1130.
<https://doi.org/10.3758/s13423-014-0585-6>
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3), 280–291.
<https://doi.org/10.1016/j.cognition.2011.10.004>

- Porat, R., & Shen, Y. (2017). Metaphor: The journey from bidirectionality to unidirectionality. *Poetics Today*, 38(1), 123–140.
- R Core Team. (2018). *R: A language and environment for statistical computing*.
- Radden, G., & Kövecses, Z. (1999). Towards a theory of metonymy. In K.-U. Panther & G. Radden (Eds.), *Metonymy in language and thought* (Vol. 4, pp. 17–60). John Benjamins.
- Regier, T., Carstensen, A., & Kemp, C. (2016). Languages support efficient communication about the environment: Words for snow revisited. *PloS One*, 11(4).
<https://doi.org/10.1371/journal.pone.0151138>
- Reijnierse, W. G., Burgers, C., Bolognesi, M., & Krennmayr, T. (2019). How polysemy affects concreteness ratings: The case of metaphor. *Cognitive Science*, 43(8), e12779.
<https://doi.org/10.1111/cogs.12779>
- San Roque, L., Kendrick, K. H., Norcliffe, E., Brown, P., Defina, R., Dingemanse, M., Dirksmeyer, T., Enfield, N. J., Floyd, S., & Hammond, J. (2015). Vision verbs dominate in conversation across cultures, but the ranking of non-visual verbs varies. *Cognitive Linguistics*, 26, 31–60. <https://doi.org/10.1515/cog-2014-0089>
- San Roque, L., Kendrick, K. H., Norcliffe, E., & Majid, A. (2018). Universal meaning extensions of perception verbs are grounded in interaction. *Cognitive Linguistics*, 29(3), 371–406. <https://doi.org/10.1515/cog-2017-0034>
- Sawyer, J. O., & Schlichter, A. (1984). *Yuki vocabulary*. Berkeley University Press.
- Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124–152.
- Strik Lievers, F. (2015). Synaesthesia: A corpus-based study of cross-modal directionality. *Functions of Language*, 22(1), 69–95.

- Strik Lievers, F. (2017). Figures and the senses. *Review of Cognitive Linguistics*, 15(1), 83–101.
- Sweetser, E. (1990). *From etymology to pragmatics*. Cambridge University Press.
- Thibodeau, P., & Durgin, F. H. (2008). Productive figurative communication: Conventional metaphors facilitate the comprehension of related novel metaphors. *Journal of Memory and Language*, 58(2), 521–540.
- Tomasello, M., & Farrar, M. J. (1986). Joint attention and early language. *Child Development*, 1454–1463.
- Torchiano, M. (2019). *effsize: Efficient effect size computation*.
<https://doi.org/10.5281/zenodo.1480624>
- Traugott, E. C. (1989). On the rise of epistemic meanings in English: An example of subjectification in semantic change. *Language*, 31–55.
- Traugott, E. C., & Dasher, R. B. (2001). *Regularity in semantic change* (Vol. 97). Cambridge University Press.
- Ullmann, S. (1959). *The principles of semantics*. Jackson, Son & Co.
- Urban, M. (2011). Asymmetries in overt marking and directionality in semantic change. *Journal of Historical Linguistics*, 1(1), 3–47.
- Urban, M. (2012). *Analyzability and semantic associations in referring expressions: A study in comparative lexicology* [PhD Thesis]. Leiden University Centre for Linguistics (LUCL), Faculty of the Humanitas, Leiden University.
- Viberg, Å. (1983). The verbs of perception: A typological study. *Linguistics*, 21(1), 123–162.
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4), 1191–1207.

- Wickham, H. (2017). *tidyverse: Easily install and load the “tidyverse.”* <https://CRAN.R-project.org/package=tidyverse>
- Wilkins, D. P. (1996). Natural tendencies of semantic change and the search for cognates. In M. Durie & M. Ross (Eds.), *The comparative method reviewed: Regularity and irregularity in language change* (pp. 264–304). Oxford University Press.
- Williams, J. M. (1976). Synaesthetic adjectives: A possible law of semantic change. *Language*, 461–478.
- Winter, B. (2016). Taste and smell words form an affectively loaded and emotionally flexible part of the English lexicon. *Language, Cognition and Neuroscience*, 31(8), 975–988.
- Winter, B. (2019). *Sensory linguistics: Language, perception, and metaphor*. John Benjamins.
- Winter, B., Marghetis, T., & Matlock, T. (2015). Of magnitudes and metaphors: Explaining cognitive interactions between space, time, and number. *Cortex*, 64, 209–224.
- Winter, B., Thompson, G., & Urban, M. (2013). Cognitive factors motivating the evolution of word meanings: Evidence from corpora, behavioral data and encyclopedic network structure. In E. Cartmill, H. Lyn, & H. Cornish (Eds.), *10th International Conference on the Evolution of Language* (pp. 353–360). World Scientific.
- Xu, Y., Duong, K., Malt, B. C., Jiang, S., & Srinivasan, M. (2020). Conceptual relations predict colexification across languages. *Cognition*, 201, 104280. <https://doi.org/10.1016/j.cognition.2020.104280>
- Xu, Y., Malt, B. C., & Srinivasan, M. (2017). Evolution of word meanings through metaphorical mapping: Systematicity over the past millennium. *Cognitive Psychology*, 96, 41–53. <https://doi.org/10.1016/j.cogpsych.2017.05.005>
- Zipf, G. K. (1945). The meaning-frequency relationship of words. *The Journal of General Psychology*, 33(2), 251–256. <https://doi.org/10.1080/00221309.1945.10544509>