

Metonymy as a Universal Cognitive Phenomenon: Evidence from Multilingual Lexicons

Temuulen Khishigsuren (kh.temulen@gmail.com)

National University of Mongolia, Ulaanbaatar, Mongolia

Gábor Bella (gabor.bella@unitn.it)

University of Trento, Trento, Italy

Thomas Brochhagen (thomas.brochhagen@upf.edu)

Universitat Pompeu Fabra, Barcelona, Spain

Daariimaa Marav (daariimaa@num.edu.mn)

National University of Mongolia, Ulaanbaatar, Mongolia

Fausto Giunchiglia (fausto.giunchiglia@unitn.it)

University of Trento, Trento, Italy

Khuyagbaatar Batsuren (khuyagbaatar@num.edu.mn)

National University of Mongolia, Ulaanbaatar, Mongolia

Abstract

Metonymy is regarded as a universally shared cognitive phenomenon; as such, humans are taken to effortlessly produce and comprehend metonymic senses. However, experimental studies on metonymy have been focused on Western societies, and the linguistic data backing up claims of universality has not been large enough to provide conclusive evidence. We introduce a large-scale analysis of metonymy based on a lexical corpus of 20 thousand metonymy instances from 189 languages and 69 genera. No prior study, to our knowledge, is based on linguistic coverage as broad as ours. Drawing on corpus and statistical analysis, evidence of universality is found at three levels: systematic metonymy in general, particular metonymy patterns, and specific metonymy concepts. These findings imply that a shared conceptual structure for these patterns and concepts holds across societies.

Keywords: metonymy; lexical semantics; universals; multilingual lexical resources; conceptual structure

Introduction

Metonymy was considered a figure of speech where one meaning of a word serves as a reference for another meaning of the same word. In the sentence ‘*The chicken was tasty*,’ the animal sense of *chicken* refers to the meat sense. Since the emergence of the theory of conceptual mappings, however, metonymy is more commonly deemed a cognitive phenomenon rather than a mere linguistic expression (Lakoff & Johnson, 2008). Some metonymies appear to be non-arbitrary and show a certain level of systematicity, which has been described as *metonymic patterns*, such as in the preceding example ANIMAL FOR MEAT: we can say ‘*The lamb was tasty*,’ ‘*The fish was tasty*,’ ‘*The turkey was tasty*,’ and so forth. In contrast, Nunberg (1995)’s famous example of ‘*The ham sandwich is at table 7*’ highlights a different kind of metonymy, often called ‘circumstantial,’ that is only understandable within a restricted range of situations. Unlike circumstantial metonymies, systematic metonymies are highly conventionalized and registered in lexicons.

A common theoretical stance on the first, systematic kind of metonymy posits that it is a universal cognitive phenomenon; hence metonymic senses are taken to be

highly accessible and comprehended by humans automatically (Barcelona et al., 2003; Brdar & Brdar-Szabó, 2003; Croft, 2002a; Gibbs, Gibbs, & Gibbs, 1994; Kövecses & Radden, 1998; Lakoff & Johnson, 2008; Panther & Radden, 1999). Some researchers even claim that metonymy is grounded in a basic cognitive tendency that children can acquire without prior experience (Papafragou, 1996). Recent psycho- and neurolinguistic evidence supports these claims. Several studies report that no additional processing cost is required for comprehending systematic metonymy, and late positivity (pragmatic adjustment) is observed only in the processing of circumstantial metonymy (Frisson & Pickering, 2007; Piñango et al., 2017; Weiland, Bambini, & Schumacher, 2014). Recent work further shows that prior experience of particular exemplars is not necessary for children to learn at least one kind of systematic metonymy: PRODUCER FOR PRODUCT as in ‘*I read Shakespeare*’ (Zhu, 2021).

However, these experimental studies were conducted in Western societies, which leaves the question whether these results generalize to non-Western cultures unresolved. One way to avoid this bias in experimental paradigms is to explore cross-linguistic evidence (Youn et al., 2016; J. Jackson, Watts, List, Drabble, & Lindquist, 2020). However, past cross-linguistic surveys on metonymy have also been limited to a small number of oft-studied languages (Barcelona et al., 2003; Janda, 2011; Panther & Thornburg, 1999; Ruiz de Mendoza Ibáñez & Pérez Hernández, 2003; Sweep, 2012; Zhang, Speelman, & Geeraerts, 2015). An important reason for this limitation is that current methodologies in cross-linguistic semantic analysis require serious involvement of language experts (Brdar-Szabó & Brdar, 2003a, 2003b, 2012) or native speakers (Kamei & Wakao, 1992; Slabakova, Cabrelli Amaro, & Kang, 2013; Srinivasan & Rabagliati, 2015), or that they are simply not suitable for metonymy studies; e.g., elicitation techniques (Koptjevskaja-Tamm, Rakhilina, & Vanhove, 2015).

On the other hand, the recent trend of exploiting digitally available lexical resources makes large-scale semantic




Figure 1: The Universal Knowledge Core (UKC) structure.

studies feasible; e.g., the study of the emotion domain in 2,474 languages (J. C. Jackson et al., 2019), concepts from 22 semantic domains over 246 languages (Xu, Duong, Malt, Jiang, & Srinivasan, 2020), or the areal typology of colexification patterns across 4,664 languages (Gast & Koptjevskaja-Tamm, 2018). This method is especially suitable for systematic metonymy, as it is lexically encoded. Building on these efforts, we here used lexico-semantic content of high-quality multilingual lexical databases to build a large-scale metonymy corpus that covers 26 metonymy patterns and 20 thousand metonymy instances (word pairs) in 189 languages, belonging to 69 genera. This new, freely available, online corpus of metonymy examples categorized by patterns is also reusable for future studies.

Due to their broad linguistic coverage, our results considerably strengthen the claim that metonymy is a universal cognitive phenomenon. What is more, they further suggest a particularly strong universality for specific metonymy patterns, e.g., FRUIT FOR PLANT and PLANT FOR FOOD. Yet, even rarer universal patterns are attested across diverse languages from different parts of the world. Finally, on the conceptual level, many concepts appearing in universal patterns turn out to be universal themselves.

In what follows, we start by describing our method to extract our metonymy corpus from a database. Subsequently, we present our results of the metonymy corpus and of the statistical analysis applied on it. Finally, we provide a general discussion.

Methods

Computational account of metonymy Building a large multilingual corpus of metonymies in an efficient, partially automated manner necessitates that metonymy can be described at a level of formality that a computer can successfully exploit in search for metonymies over a given textual database. In order to find such formal criteria, we start by the most basic definitions already announced in the first sentence of our introduction: *a figure of speech where one meaning of a word serves as a reference for another meaning of the same word*. In other terms, we need to identify words having mul-

iple meanings where one meaning ‘refers’ to another. Such references are understood to be systematic: the existing literature on metonymy (Copestake & Briscoe, 1995; Kövecses & Radden, 1998; Klein & Murphy, 2002; Peirsman & Geeraerts, 2006; Lakoff & Johnson, 2008; Dölling, 2020) agrees on *metonymy patterns* being a useful typological device to characterise this systematicity. Metonymy patterns follow the structure of *target for vehicle*, such as ANIMAL FOR MEAT or BODY PART FOR PERSON, where both *target* and *vehicle* are understood to be categories or domains to which the meanings of the word belong.

Thus, the English word *guard* instantiates the ACTION FOR AGENT metonymy pattern through its polysemous senses *person keeping watch* (as a noun) and *to keep watch over* (as a verb) (see Figure 1). However, many researchers argue that the generation of similar meaning pairs through morphological alternation should also be considered as cases of metonymy, on the basis that the same underlying cognitive principles are being applied (Copestake & Briscoe, 1995; Brdar-Szabó & Brdar, 2003b; Janda, 2011). In French, for instance, the first sense is lexicalized as *garde* and the second one as *garder* (*garde* + *-er*). We will refer to the first, polysemy-based, kind of metonymy as *lexical metonymy* and to the second, derivation-based, kind as *morphological metonymy*.¹

Database used in the study Among the various kinds of resources—databases, dictionaries, corpora—that were available to us, *multilingual lexical databases* respond best to the criteria detailed above, namely: the explicit representation of words, their meanings, and their domains in multiple languages, as well as the presence of a cross-lingual alignment of meanings and domains. Our database of choice is the Universal Knowledge Core (UKC)² (Giunchiglia, Batsuren, & Bella, 2017), due to its wide linguistic, lexical, and conceptual coverage (120 thousand word meanings, 2 million words in 1,127 languages). The UKC has been used in several stud-

¹According to François (2008), lexical metonymy is a ‘strict colexification’ while morphological metonymy is a ‘loose colexification.’

²<http://ukc.datascientia.eu/>

ies in computational linguistics and lexical semantics, such as the computation of cognates (Batsuren, Bella, & Giunchiglia, 2019, 2021a) and multilingual morphology (Batsuren, Bella, & Giunchiglia, 2021b).

Metonymy corpus extraction Metonymy patterns are straightforward to model through the three-layered *domain–concept–lexicon* architecture of the UKC (Figure 1). The concept layer represents *supra-lingual meanings* as a hierarchy of concepts based on the standard lexicographic broader–narrower (hypernym–hyponym) relationship. The domain layer of the UKC provides a simple semantic categorization of concepts into domains such as *Animal*. The lexical layer, finally, consists of a separate *lexicon* for each language, each one lexicalizing the supra-lingual concept layer. The French word *poulet* and the English word *chicken* both have two meanings, connected to the corresponding supra-lingual concepts. Due to the presence of such instances, we consider that these concept pairs are *metonymically related*.

The extraction process consisted of three steps:

- 1) *Pattern mapping and selection*: from the metonymy patterns mentioned in the literature, an expert-driven selection of a subset for which the UKC provides data;
- 2) *Concept pair selection*: automatic extraction and expert validation of metonymically related concept pairs;
- 3) *Metonymy instance extraction*: based on the definitive set of metonymically related concept pairs, automated extraction of lexicalizations for all languages in the database.

Pattern Mapping and Selection

The mapping of metonymy patterns was a manual process that consisted of: (a) understanding the meaning of categories inside metonymy patterns found in the literature; and (b) the formalization of these categories by mapping them to one or more existing database domains. For the understanding of categories, our starting point was a set of about a hundred metonymy patterns found in the literature from the last 25 years (Copestake & Briscoe, 1995; Dölling, 2020; Klein & Murphy, 2002; Kövecses & Radden, 1998; Lakoff & Johnson, 2008; Peirsman & Geeraerts, 2006). These metonymy patterns refer to 65 semantic categories in total; however, the same category is often interpreted differently according to the pattern, and different categories can overlap significantly. Thus, *EVENT* in *EVENT FOR PEOPLE* specifically denotes social events while in *PLACE FOR EVENT* it may also involve physical events such as an explosion (*Tchernobyl*).

The domain layer of the UKC categorizes concepts into 45 domains. The mapping between categories from the literature and UKC domains required the creation of new domains, sometimes by splitting apart or fusing existing ones. For example, *POSSESSOR* and *AGENT* were mapped to the domain *Person*. The category *CAUSE* was mapped to two domains corresponding to animate and inanimate causes: *Person* and *Stimulus*. The result of this formalization was a shortlist of 26 patterns successfully mapped to UKC domains, as shown in Table 1.

Concept Pair Selection

The next step towards the extraction of metonymic lexicalizations is the identification of metonymically related concepts, e.g. *domestic fowl* and *flesh of a chicken* as shown in Figure 1. Based on the UKC domains identified for each metonymy pattern, the retrieval of potentially related *candidate* concept pairs can be automated. We considered a concept pair c_1, c_2 as a candidate for metonymic relatedness if both of the following criteria are fulfilled: (a) the concept pair instantiates one of the patterns listed in Table 1, i.e. $c_1 \in d_v$ and $c_2 \in d_t$ such that the domains (d_v, d_t) form a metonymy pattern; (b) c_1 and c_2 have colexifications in at least one language.

Using this algorithm, we automatically extracted over 51 thousand candidate concept pairs. The fact that the UKC is based on Princeton WordNet (Miller, Beckwith, Fellbaum, Gross, & Miller, 1990), which consists of some low-frequency subtle senses, results in a high number of candidate concept pairs. For instance, there are 32 different senses of the verb *fall* resulting in 32 candidate pairs. But only one is an instance of *ACTION FOR TIME* pattern. Similarly, out of 16 different senses of the verb *dress*, only three are metonymic extensions of the noun *dress* ‘a clothing.’ Meanings such as ‘to dress a cake,’ ‘to dress a window,’ and ‘to dress hair’ appear to be metaphors.

In order to filter the candidates to retain only metonymic relationships, two linguists who are the authors of the paper manually annotated them. They started by becoming familiar with the metonymy patterns through illustrative examples in English. Then, they were provided with the extracted data for each pattern from the previous step. The data contains information on concepts defined by gloss description, approximate English words, and language-specific examples. Based on this information, the linguists were asked the following question: ‘Do you consider the concepts c_1 and c_2 metonymically related?’, having to provide a yes/no answer.

Metonymy Instance Extraction

Our linguists annotated over 4,900 concept pairs with metonymic relatedness within a two-month period. After the annotation process, we automatically extracted the corresponding lexicalizations for all languages in the UKC database. If, for example, experts annotated the following two concepts as metonymically related: *person keeping watch* and *to keep watch over*, then the corresponding lexicalizations such as the English *guard* and the French *garde–garder* were automatically extracted from the database (Figure 1). Next, we present the results of this metonymy extraction process.

Analysis of Results

Table 1 reports the statistics of the metonymy corpus³ extracted semi-automatically from the database. Overall, 4,951 concept pairs were annotated as metonymically related, and the corresponding 20,095 metonymy instances were retrieved

³The metonymy corpus is freely accessible as a stand-alone resource at <https://github.com/kbatsuren/UniMet>.

Table 1: Metonymy corpus statistics

Metonymy pattern	Abbrev.	Illustrative example	Met. instances			Langs	Fami-lies	Gen-era
			Met. concepts	Lex.met	Morph.met			
Substance for Artifact	SubArt	He filled the <i>glass</i> with water.	390	1,076	699	110	24	46
Fruit for Plant	FruPla	The gardener watered the <i>lemon</i> .	408	1,934	1,396	114	24	43
Instrument for Action	InsAct	She <i>combed</i> her hair.	617	490	1,593	95	24	40
Community for Place	ComPla	He traveled to the <i>country</i> .	87	503	232	97	22	38
Plant for Food	PlaFoo	<i>Broccoli</i> is delicious.	318	1,244	295	80	19	34
Animal for Meat	AniMea	The <i>chicken</i> is tasty.	156	413	333	85	19	33
Action for Result	ActRes	My thumb has a deep <i>cut</i> .	729	416	1,143	77	17	32
Object for Action	ObjAct	They are well <i>dressed</i> .	546	398	1,255	79	17	31
Substance for Action	SubAct	I <i>milked</i> cows by hand.	242	177	765	78	17	31
Emotion for Cause	EmoCas	You are my <i>joy</i> .	104	305	100	64	14	28
State for Causal agent	StaAge	He was a <i>success</i> .	160	407	238	59	15	27
Food for Action	FooAct	They had <i>breakfasted</i> so early.	51	51	159	55	13	27
Building for People	BuiPeo	<i>Church</i> sang a song.	71	248	136	63	15	26
Possessed for Possessor	PosPos	She married <i>power</i> .	190	379	334	60	14	26
Agent for Action	AgeAct	The sheep will be <i>butchered</i> .	232	158	497	53	14	26
Product for Content	ProCon	The <i>book</i> is interesting.	46	293	253	60	14	25
Body part for Person	BodPer	I saw many new <i>faces</i> today.	156	204	238	61	12	25
Action for Food	ActFoo	They provided a <i>drink</i> at the party.	16	22	68	47	14	22
Animal for Fur	AniFur	She likes to wear <i>mink</i> .	51	133	138	35	14	20
Container for Contained	ConCon	He drank half of the <i>bottle</i> .	88	278	183	42	10	19
Event for People	EvePeo	<i>Party</i> went crazy.	61	166	91	39	11	18
Action for Object	ActObj	A <i>lift</i> fell to the bottom of its shaft.	91	67	86	41	12	17
Action for Agent	ActAge	You may be a <i>help</i> later.	57	31	176	34	12	17
Time for Action	TimAct	We <i>honeymooned</i> in Bali.	25	14	94	36	11	17
Food for Event	FooEve	<i>Dinner</i> took longer than usual.	9	87	45	36	11	17
Action for Time	ActTim	My <i>shift</i> is over this morning.	50	25	29	21	8	14
Total			4,951	9,519	10,576	189	34	69

in 189 languages from the database. However, the number of languages in itself is not a reliable metric of universality, for two reasons. First, the majority of the thousand lexicons in our source database are very small (50% of them have less than 30 words). They thus provide few metonymy instances or none at all (we examine this effect in the next section). Second, a set of unrelated languages is much more representative of universality than a set of closely related ones (e.g. Romance languages).

For this reason, following studies on universals (Croft, 2002b), we examined how many phylogenetically different languages attest metonymy in general and metonymy patterns in particular. Thus, we computed statistics on the number of languages, families, and genera, that we provide in Table 1. We found evidence of metonymy in 189 languages belonging to 34 different families (phyla) and 69 genera. These languages are also geographically stratified (Figure 2). To the best of our knowledge, these results provide the widest linguistic coverage so far on metonymy (the broadest prior study we are aware of is by Hilpert (2007) that reported 39 phylogenetically different languages using *eye* to refer to *vision*).

Out of 20 thousand metonymy instances, we found 9,519 lexical and 10,576 morphological metonymies. As shown in Table 1, all ACTION patterns tend to be expressed by morphological metonymy rather than lexical metonymy, as noun-verb and verb-noun conversions often involve morphological alternations in many languages. Other patterns except for ANIMAL FOR FUR prefer lexical over morphological metonymy.

We identified 62 languages, 15 families, and 27 genera on average per pattern. Although some patterns were identified




Figure 2: The presence of metonymy in world’s languages (the same colors indicate the same family).

in more languages than others, even the least widely covered patterns are attested across phylogenetically diverse languages from around the world. For example, the least diverse pattern, ACTION FOR TIME, is still attested in 21 languages from 14 genera and eight families from Africa, East Asia, the Pacific, Europe, and the Middle East. This result suggests that diverse societies use ACTION FOR TIME metonymies. Based on the number of genera, the most universal pattern is SUBSTANCE FOR ARTIFACT for which we found 110 languages from 46 genera. FRUIT FOR PLANT and INSTRUMENT FOR ACTION are also very widely attested patterns for each we found 43 and 40 genera, respectively.

On the conceptual level, even specific concept pairs appear to be universal: for instance, 49 languages from 22 genera use the concept ‘pear’ to refer to its plant name either through lexical metonymy (31 languages from 15 genera; e.g., Russian группа ‘pear’ and ‘pear tree’) or morphological metonymy (29 languages from 16 genera; e.g., Catalan *pera* ‘pear’ and

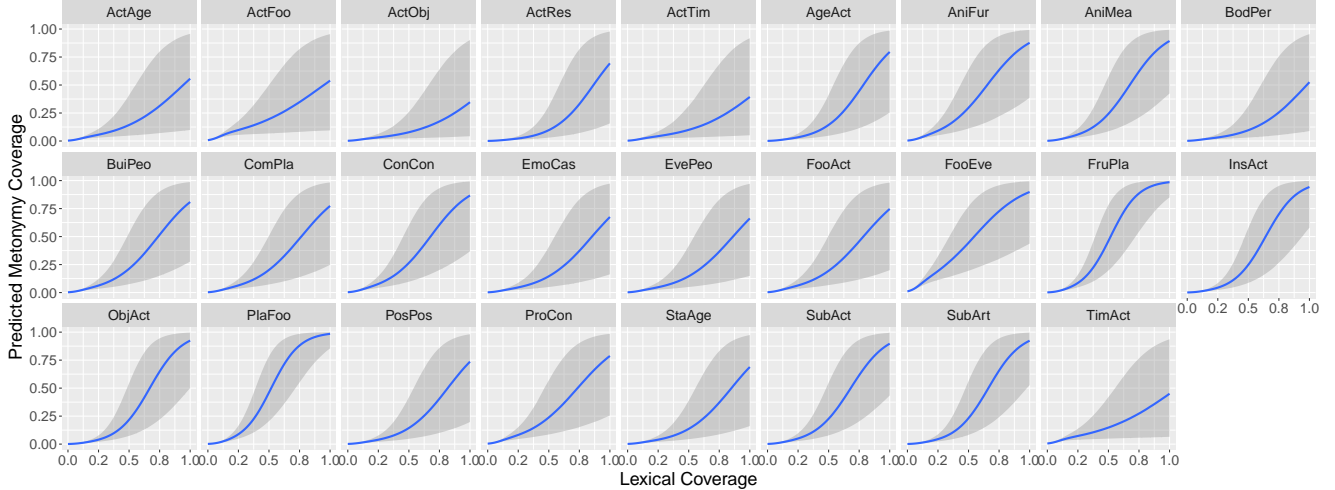


Figure 3: Marginal effect of lexical coverage on metonymy by pattern (blue line: mean effect, gray area: 95% credible intervals).

perera ‘pear tree’). The example of ‘chicken’ in the Introduction is exploited by 35 languages from 18 genera, while the English examples of COMMUNITY FOR PLACE and INSTRUMENT FOR ACTION patterns in Table 1 are also attested in 69 and 39 other languages from 34 and 22 genera, respectively.

A close look at the corpus also reveals the existence of language-specific metonymy instances. In Basque, for example, the word meaning ‘skin’ refers to a naked person as an instantiation of BODY PART FOR PERSON, but it is used to name a skinny person in Mongolian. The word meaning ‘foot’ can mean a disliked woman in Dutch but someone obedient in Malay. By contrast, it refers to a surveillance team member in English. These results reflect different cultural preferences over the same body part term.

In conclusion, we found evidence of universality not only for systematic metonymy in general but also for selected metonymy patterns as well as specific metonymy concepts. These findings imply that there is a shared conceptual structure for at least these particular patterns and concepts across diverse societies. On the concept level, we also explored language-specific idiosyncrasies.

Lexical Coverage Effects

The UKC, our source database, consists of lexicons of widely uneven sizes, from English containing over 100,000 words to the incomplete lexicons of many endangered languages consisting of less than 10 words. In this section we examine the effect of varying lexicon sizes on our results.

Among the 189 languages for which metonymies were found, six provide 47% of all metonymy instances, while the 87 smallest lexicons (with less than a thousand words) account for only 1% of them. The lexicon size of a language is highly correlated with the number of metonymy instances found ($r(189) = 0.88, p < 0.0001$) and also with the number of attested patterns ($r(189) = 0.73, p < 0.0001$). From these results, it seems reasonable to expect that if the lexical cov-

erage of low-resourced languages increases, the number of languages attesting metonymy will also increase. In order to verify this hypothesis more deeply, we performed a Bayesian statistical analysis to evaluate the effect of lexical coverage on metonymy.

All statistical analyses were conducted with *R* v4.1.2 (R Core Team, 2021) and *brms* v2.16.1 (Bürkner, 2021) was used to fit a Bayesian zero-inflated beta model. The zero-inflated beta model was appropriate for our data as we computed proportional measures that contain many zeros. Two kinds of measures were computed on the metonymy corpus and the database: *metonymy coverage* and *lexical coverage*. We define *metonymy coverage* as follows:

$$\text{MetCov}(l, p) = \frac{|\text{Concepts}(l, M_p)|}{|\text{Concepts}(p)|} \quad (1)$$

where $\text{Concepts}(l, M_p)$ is the set of concepts covered by a metonymy set M_p in language l for the given pattern p , and $\text{Concepts}(p)$ denotes all concepts of the pattern p . We define *lexical coverage* as follows:

$$\text{LexCov}(l, p) = \frac{|\text{Concepts}(l, L_p)|}{|\text{Concepts}(p)|} \quad (2)$$

where $\text{Concepts}(l, L_p)$ is the set of concepts covered by a set of lexicalizations L_p in language l for the given pattern p , and $\text{Concepts}(p)$ denotes all concepts of the pattern p .

Metonymy coverage was predicted by lexical coverage as a fixed effect. We added random intercepts and slopes for metonymy patterns to allow variation across patterns. We included random intercepts and slopes for language families to allow for genetic relatedness between languages. Metonymy coverage is positively associated with lexical coverage (logit coefficient = 5.68, 95% credible interval = [4.64, 6.80], probability of the effect $> 0 = 100\%$). We can interpret this result as the probability of new lexicalizations being metonymy instances. The population-level effect shows that at 100%

coverage, for example, the probability of new lexicalizations being metonymy is 81% on average, with credible intervals between 62% and 94%. Therefore, the number of languages (189) attesting metonymy is likely to increase if the lexical coverage of low-resourced languages increases.

We illustrate the marginal effect by pattern in Figure 3. The effect of lexical coverage on metonymy coverage varies with each pattern. The effects are high for FRUIT FOR PLANT and PLANT FOR FOOD with low uncertainty, further strengthening the universality for these patterns. Compared to these two, high mean effects with slightly higher uncertainty were observed in patterns such as FOOD FOR EVENT and OBJECT FOR ACTION. Interestingly, FOOD FOR EVENT is one of the least diverse patterns according to Table 1, but statistical analysis suggests otherwise, provided that the coverage increases. We cannot provide conclusive evidence for most of the remaining patterns due to high uncertainty, although some patterns show higher mean effects (e.g., AgeAct) and some show lower mean effects (e.g., ActTim).

Discussion

Metonymy is often claimed to be a universal cognitive phenomenon. However, heretofore this claim had not been backed by sufficient empirical evidence. Here, we present a new, large-scale multilingual analysis based on a lexical corpus of over 20 thousand metonymy instances in 189 languages. We found that phylogenetically diverse languages attest 26 metonymy patterns known from past literature. Although our analysis suggests a particularly strong universality for patterns like FRUIT FOR PLANT, even the rarest universal patterns are attested across diverse languages and communities. In the future, we plan to extend our study beyond the selected 26 patterns. We also plan to recruit native speakers to judge language-specific instances to further corroborate the quality of our corpus.

The fact that many concepts that participate in universal patterns turn out to be universal themselves raises a question for Srinivasan and Rabagliati (2015)'s conventions-constrained-by-concepts model of polysemy. Their model assumes that polysemy patterns, including metonymy, should be present across languages but that language-specific instantiations must vary substantially. Although this prediction squares well with, e.g., the BODY PART FOR PERSON pattern in our data; others, e.g., FRUIT FOR PLANT were instead found to have low cross-linguistic variation, with universally attested concepts. Shedding further light on this apparent conflict may be a promising venue for future research.

Under the Typological Prevalence Hypothesis (Gentner & Bowerman, 2009), the universality of semantic categories is closely related to the notion of *naturalness* in psychology. Under this hypothesis, patterns that are more common across languages, such as FRUIT FOR PLANT and PLANT FOR FOOD, are more natural to people. Naturalness predicts frequency of use and ease of acquisition. Consequently, more natural patterns are attested more often across languages. Srinivasan and

Rabagliati (2015) also found some patterns to be more natural than others according to native speaker judgments. In particular, while PLANT FOR FOOD was also found to be more natural, SUBSTANCE FOR ACTION was judged to be less natural, contrary to our findings. This contradiction may be the consequence of language coverage (15 vs 189) but it could also be due to our inclusion of morphological metonymy. For example, while they report that Hungarian speakers found SUBSTANCE FOR ACTION unacceptable, we found plenty of evidence for such morphological metonymy in Hungarian; e.g. *só* 'salt' and *sóz* 'to add salt.'

All in all, these findings suggest that considering loose colexification (in our case, morphological metonymy) may be crucial in the study of conceptual relations. We found, in particular, that certain relations are predominantly expressed by morphology rather than polysemy. While Janda (2011) also reported similar results, previous work on conceptual relations often overlooked loose colexifications. Xu et al. (2020)'s study, for example, suggests that more related concepts tend to colexify more often across languages. However, this study only considered strict colexifications, and the picture may change if loose metonymy is factored in. Ultimately, we believe that further research is necessary to investigate conceptual relations both in terms of strict and loose colexifications; and to address the above contradictions. Once lexical databases such as the UKC evolve towards a more balanced coverage of low-resourced languages, they will be able to provide sufficient data covering both strict and loose colexifications to examine, e.g., which patterns are more natural.

In addition, we want to emphasize the importance of lexical coverage measures in resource-based language studies such as ours. To illustrate, a look at CLICS³ (Rzymiski et al., 2020) reveals that *hand-arm* is one of the most common colexifications, attested in 294 languages from 48 families. The lexical coverage of these concepts is high, about two thousand words each. A similar relation of physical contiguity is present in the colexification of *upper back-back*. However, these concepts are colexified in only 29 languages from 13 families. A driving factor for this difference might be that the concept *upper back* has only 60 lexicalizations. We accordingly believe that formally laying out coverage metrics and factoring them into analyses is necessary for cross-linguistic studies based on (often imbalanced) digitally available resources.

Finally, we wish to draw attention to so far overlooked questions in experimental paradigms. For instance, whether there is a difference within systematic metonymies with respect to their processing and their ease of acquisition by children. Experimental studies have focused on differences between systematic and circumstantial metonymies (Piñango et al., 2017). However, our findings, together with previous literature, hint at a plausible difference within the class of systematic ones. Examining this within-class variation may shed light on underlying cognitive mechanisms, explaining why some patterns are more natural than others and, more generally, what principles govern metonymic mappings.

References

- Barcelona, A., et al. (2003). Names: A metonymic “return ticket” in five languages. *Jezikoslovlje*, 4(1), 11–41.
- Batsuren, K., Bella, G., & Giunchiglia, F. (2019). Cognet: A large-scale cognate database. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 3136–3145).
- Batsuren, K., Bella, G., & Giunchiglia, F. (2021a). A large and evolving cognate database. *Language Resources and Evaluation*, 1–25.
- Batsuren, K., Bella, G., & Giunchiglia, F. (2021b). Morphynet: a large multilingual database of derivational and inflectional morphology. In *Proceedings of the 18th sigmorphon workshop on computational research in phonetics, phonology, and morphology* (pp. 39–48).
- Brdar, M., & Brdar-Szabó, R. (2003). in english, croatian and hungarian. *Metonymy and pragmatic inferencing*, 113, 241.
- Brdar-Szabó, R., & Brdar, M. (2003a). The manner for activity metonymy across domains and languages. *Jezikoslovlje*, 4(1), 43–69.
- Brdar-Szabó, R., & Brdar, M. (2003b). Referential metonymy across languages: What can cognitive linguistics and contrastive linguistics learn from each other? *International Journal of English Studies*, 3(2), 85–106.
- Brdar-Szabó, R., & Brdar, M. (2012). The problem of data in the cognitive linguistic research on metonymy: a cross-linguistic perspective. *Language Sciences*, 34(6), 728–745.
- Bürkner, P.-C. (2021). Bayesian item response modeling in R with brms and Stan. *Journal of Statistical Software*, 100(5), 1–54. doi: 10.18637/jss.v100.i05
- Copestake, A., & Briscoe, T. (1995). Semi-productive polysemy and sense extension. *Journal of semantics*, 12(1), 15–67.
- Croft, W. (2002a). The role of domains in the interpretation of metaphors and metonymies. *Metaphor and metonymy in comparison and contrast*, 161–205.
- Croft, W. (2002b). *Typology and universals*. Cambridge University Press.
- Dölling, J. (2020, 11). Systematic polysemy. In (p. 1-27). doi: 10.1002/9781118788516.sem099
- François, A. (2008). Semantic maps and the typology of colexification. *From polysemy to semantic change: Towards a typology of lexical semantic associations*, 106, 163.
- Frisson, S., & Pickering, M. J. (2007). The processing of familiar and novel senses of a word: Why reading dickens is easy but reading needham can be hard. *Language and cognitive processes*, 22(4), 595–613.
- Gast, V., & Koptjevskaja-Tamm, M. (2018). The areal factor in lexical typology. *Trends in Linguistics Studies and Monographs*, 3.
- Gentner, D., & Bowerman, M. (2009). Why some spatial semantic categories are harder to learn than others: The typological prevalence hypothesis. *Crosslinguistic approaches to the psychology of language: Research in the tradition of Dan Isaac Slobin*, 465, 480.
- Gibbs, R. W., Gibbs, R. W., & Gibbs, J. (1994). *The poetics of mind: Figurative thought, language, and understanding*. Cambridge University Press.
- Giunchiglia, F., Batsuren, K., & Bella, G. (2017). Understanding and exploiting language diversity. In *Ijcai* (pp. 4009–4017).
- Hilpert, M. (2007). Chained metonymies in lexicon and grammar. *Aspects of meaning construction*, 77–98.
- Jackson, J., Watts, J., List, J.-M., Drabble, R., & Lindquist, K. (2020). From text to thought: How analyzing language can advance psychological science.
- Jackson, J. C., Watts, J., Henry, T. R., List, J.-M., Forkel, R., Mucha, P. J., ... Lindquist, K. A. (2019). Emotion semantics show both cultural variation and universal structure. *Science*, 366(6472), 1517–1522.
- Janda, L. A. (2011). Metonymy in word-formation.
- Kamei, S.-i., & Wakao, T. (1992). Metonymy; reassessment, survey of acceptability, and its treatment in a machine translation system. In *30th annual meeting of the association for computational linguistics* (pp. 309–311).
- Klein, D. E., & Murphy, G. L. (2002). Paper has been my ruin: Conceptual relations of polysemous senses. *Journal of Memory and Language*, 47(4), 548–570.
- Koptjevskaja-Tamm, M., Rakhilina, E., & Vanhove, M. (2015). The semantics of lexical typology. In *The routledge handbook of semantics* (pp. 450–470). Routledge.
- Kövecses, Z., & Radden, G. (1998). Metonymy: Developing a cognitive linguistic view. *Cognitive linguistics*, 9(1), 37–77.
- Lakoff, G., & Johnson, M. (2008). *Metaphors we live by*. University of Chicago press.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. J. (1990). Introduction to wordnet: An online lexical database. *International journal of lexicography*, 3(4), 235–244.
- Nunberg, G. (1995). Transfers of meaning. *Journal of semantics*, 12(2), 109–132.
- Panther, K.-U., & Radden, G. (1999). *Metonymy in language and thought* (Vol. 4). John Benjamins Publishing.
- Panther, K.-U., & Thornburg, L. (1999). The potentiality for actuality metonymy in english and hungarian. *Metonymy in language and thought*, 333, 359.
- Papafragou, A. (1996). On metonymy. *Lingua*, 99(4), 169–195.
- Peirsman, Y., & Geeraerts, D. (2006). Metonymy as a prototypical category. *Cognitive linguistics*, 17(3), 269–316.
- Piñango, M. M., Zhang, M., Foster-Hanson, E., Negishi, M., Lacadie, C., & Constable, R. T. (2017). Metonymy as referential dependency: psycholinguistic and neurolinguistic arguments for a unified linguistic treatment. *Cognitive science*, 41, 351–378.
- R Core Team. (2021). R: A language and environment for

- statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org>
- Ruiz de Mendoza Ibáñez, F. J., & Pérez Hernández, L. (2003). High-level modal metonymies in english and spanish. *Jezikoslovlje*, 4(1), 103–120.
- Rzyski, C., Tresoldi, T., Greenhill, S. J., Wu, M.-S., Schweikhard, N. E., Koptjevskaja-Tamm, M., ... others (2020). The database of cross-linguistic colexifications, reproducible analysis of cross-linguistic polysemies. *Scientific data*, 7(1), 1–12.
- Slabakova, R., Cabrelli Amaro, J., & Kang, S. K. (2013). Regular and novel metonymy in native korean, spanish, and english: Experimental evidence for various acceptability. *Metaphor and Symbol*, 28(4), 275–293.
- Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124–152.
- Sweep, J. (2012). *Metonymical object changes: a corpus-oriented study on dutch and german*. Netherlands Graduate School of Linguistics.
- Weiland, H., Bambini, V., & Schumacher, P. B. (2014). The role of literal meaning in figurative language comprehension: Evidence from masked priming erp. *Frontiers in Human Neuroscience*, 8, 583.
- Xu, Y., Duong, K., Malt, B. C., Jiang, S., & Srinivasan, M. (2020). Conceptual relations predict colexification across languages. *Cognition*, 201, 104280.
- Youn, H., Sutton, L., Smith, E., Moore, C., Wilkins, J. F., Maddieson, I., ... Bhattacharya, T. (2016). On the universal structure of human lexical semantics. *Proceedings of the National Academy of Sciences*, 113(7), 1766–1771. Retrieved from <https://www.pnas.org/content/113/7/1766> doi: 10.1073/pnas.1520752113
- Zhang, W., Speelman, D., & Geeraerts, D. (2015). Cross-linguistic variation in metonymies for person: A chinese-english contrastive study. *Review of Cognitive Linguistics. Published under the auspices of the Spanish Cognitive Linguistics Association*, 13(1), 220–256.
- Zhu, R. (2021). Preschoolers' acquisition of producer-product metonymy. *Cognitive Development*, 59, 101075.