COGNITIVE FACTORS MOTIVATING THE EVOLUTION OF WORD MEANINGS:

EVIDENCE FROM CORPORA, BEHAVIORAL DATA AND ENCYCLOPEDIC NETWORK STRUCTURE

BODO WINTER[1], GRAHAM THOMPSON[1] & MATTHIAS URBAN[2]

[1] Cognitive and Information Sciences,
University of California, Merced, 5200 North Lake Rd.
Merced, 95340, U.S.A.

[2]Leiden University Centre for Linguistics, University of Leiden
Postbus 9515, 2300 RA Leiden, the Netherlands

Recent linguistic work suggests that the meaning of some words may evolve in a directional fashion: for instance, words for 'skin' may develop the meaning 'bark' more easily than the other way round. Here, we investigate the underpinnings of proposed directional semantic changes by looking at synchronic data. We show that words that have been identified as candidates for the origin of semantic change (such as 'skin') are more frequent in English, have more word senses in Webster, more associations in free word association data and more connections in a network model of Wikipedia. These findings highlight key cognitive principles in the evolution of word meaning, ultimately showing that directional semantic change may be highly motivated.

1. Semantic Change

Language evolution research frequently focuses on the evolution of structural aspects of language, such as compositionality. Perhaps because semantic change often seems haphazard and unpredictable, the evolution of word meanings is less explored.

In a cross-linguistic study on 149 different languages, Urban (2011) found 45 concept pairs that suggest predictable regularity. Take, for example, the concept pair 'bark' and 'skin'[1]. In the cross-linguistic sample, the word for 'skin' is frequently used to express 'bark' (e.g., a language might have a word that can roughly be translated to mean 'tree skin'), but

[1] We use single quotation marks for concepts and italics for words.

not the other way round, i.e., there is no language that has an expression such as ‘human bark’. Urban (2011) suggests that such synchronic asymmetries in how concepts are expressed may be explained by directional trends in semantic change. This proposal is supported by an analysis of attested cases in Indo-Aryan languages (Urban, 2011: 15-24).

Table 1 lists a subset of the 45 asymmetrical concept pairs. In the sample, the word to the right of each arrow was frequently expressed by a morphologically complex form that contains the word on the left. We will call the word on the left “unmarked” and the word on the right “marked,” with the diachronic interpretation that the unmarked concept may be a frequent source and the marked concept a frequent target of semantic change.

Table 1. Some of the concept pairs proposed as developing directionally. For full list and discussion, see Urban (2011: 9-13). Arrows point from cross-linguistically unmarked to marked concepts.

Nature > Nature	***Human > Human***	***Human > Nature***	***Nature > Human***
animal > bird	breast > milk	skin > bark	egg > testicle
sun > moon	mouth > lip	mouth > beak	sun > clock
grass > straw	belly > womb	saliva > foam	shadow > mirror
cloud > fog	heart > belly	house > nest	bird > airplane
honey > wax	liver > lungs	tongue > flame	foam > lungs

Most of the concept pairs in Table 1 are characterized by some form of relation, such as similarity (e.g., ‘cloud’~‘fog’), contiguity (‘honey’~‘wax’) or taxonomy (‘animal’~‘bird’). But mere association does not explain why the direction of change only goes one way, i.e., it does not predict that ‘animal’> ‘bird’ is more likely than ‘bird’>‘animal’. What motivates directional patterns in word formation? *Why* do speakers use the word for ‘skin’ to express the meaning of ‘bark’ but not vice versa?

In this paper, we explore these questions by looking at synchronic correlates of the observed cross-linguistic asymmetries. Following the lead of Jäger and Rosenbach’s “diachronic priming hypothesis” (2008), we look at how asymmetries in *synchronic* mental associations may be useful for predicting diachronic change. To this end, we use word association data from behavioral experiments, expecting to see that words for marked concepts tend to be associated *to* words for unmarked concepts—and less so the other way round. We explore a related idea by looking at associations between concepts in encyclopedias. Using a network model of Wikipedia, where articles correspond to nodes and hyperlinks to connections, we expect to see articles for unmarked concepts to have more connections.

Association data and associations in encyclopedias can be seen as direct synchronic reflections of the trends observed by Urban (2011). To further understand where these association asymmetries come from, we look at English corpora and behavioral data from reaction time studies. We expect unmarked concepts to be more cognitively accessible. That is, we expect them to be more frequent in corpora, and we expect them to elicit faster reaction times. Here, the idea is that when new expressions are formed from old ones, these expressions are likely to be based on more accessible and more frequent concepts.

Finally, we test an old idea from linguistics: It has been proposed that words become semantically extended by virtue of being used in diverse contexts (see, already, Zipf, 1949: 19-31; see also Calude & Pagel, 2011: 1106). Based on this idea, we further expect to see unmarked concepts to occur in more different textual contexts, and, we expect them to have more senses listed in dictionaries.

In all of our analyses, we look at correlations between the cross-linguistic data and a single well-studied language, English. We focus on English for this first case study; future work will need to explore other languages.

2. Results

2.1. *Mental associations*

Nelson, McEvoy and Schreiber (1998) asked over 6,000 participants to list the first word that came to mind in response to cues such as “BOOK ____.” For a given concept pair A,B this produces a forward association probability *P*(A,B) (the proportion of people responding “B” after seeing “A”) and a backward association probability *P*(B,A). The word *animal*, for example, has a forward association of 0.016 to the word *bird*, and a backward association of 0.04. This means that more people responded *animal* when cued by *bird*. Conversely, fewer people responded *bird* when cued by *animal*.

Naturally, the English free association data does not contain exact matches for all of the 45 concept pairs in Urban (2011). In fact, we found only 8 cases in which both words of a pair came up as cue and target in the database (such as *animal* and *bird* discussed above). Probably due to sparse data, the difference between unmarked and marked concepts is only approaching significance ($t(7)=2.01$, $p=0.084$). The numeric trends show that marked concepts are likely leading *towards* words for unmarked concepts (Fig. 1a).

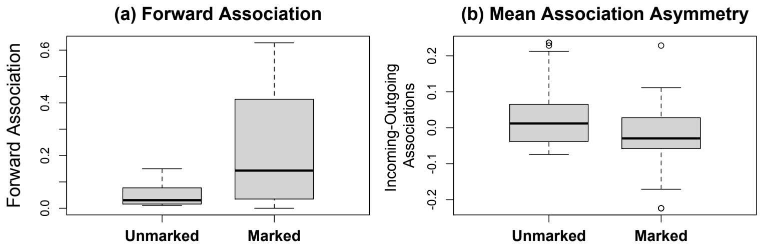

Figure 1. (a) English words for marked concepts have higher forward associations to words for unmarked concepts. (b) English words for unmarked concepts have more incoming than outgoing associations.

A measure less affected by sparseness is the mean forward association and the mean backward association, averaged across all concepts. Take, for example, the concept pair 'animal'>'bird'. The word *animal* has forward associations to *dog* (0.293), *cat* (0.12), *zoo* (0.022), as well as to many other words. It also has "incoming" backward associations from *dog* (0.026), *cat* (0), *zoo* (0.649) (here, 0 means that no participant responded *animal* after seeing *cat*). We can take the mean of all backward associations and subtract the mean of all forward associations, giving us an index of association asymmetry (cf. Hill, Korhonen & Bentz, 2013). A higher asymmetry index indicates more incoming than outgoing associations. The index is higher for words corresponding to unmarked concepts ($t(75)=2.68$, $p<0.01$), indicating that they have more "attracting" associations from other words (see Fig. 1b).

2.2. *Encyclopedic associations*

Previous research has shown that a network model of Wikipedia produces a close fit to human similarity judgments (Witten & Milne, 2008) and human free association tasks (e.g., Thompson, Kello & Montez, 2013). In line with the preceding discussion on word associations, we might thus expect asymmetries in network structure that correspond to the cross-linguistic asymmetries in Urban (2011). We can calculate an asymmetry index by subtracting the number of outgoing links (*to* other articles) from the number of incoming links (*from* other articles). In correspondence to the association data above, unmarked concepts have a positive asymmetry index of about 2,958 links (more incoming than outgoing), as opposed to minus 29 for marked concepts (more outgoing than incoming). However, this difference did not reach significance ($t(42)=0.15$, $p=0.88$), presumably because it was driven by a few articles with a lot of incoming links, such as the article for 'animal'. Articles corresponding to unmarked concepts do, however, have overall more links ($t(43)=2.62$, $p<0.02$)[2], suggesting that they are more "hub-like" in the network structure. Thus, it appears that unmarked concepts are overall more connected in knowledge networks.

2.3. *Frequency and cognitive accessibility*

Frequency data from the English Lexicon Project ("ELP", Balota et al., 2007) reveals that concepts that are cross-linguistically unmarked have more frequent English words associated with them (Fig. 2a). This is significant for the frequency measures from Lund and Burgess (1996), as well as from Kučera and Francis (1967) ($t(43)=5.84$, $p<0.001$; $t(39)=3.11$, $p<0.004$).

Figure 2b furthermore shows that response times (taken from a lexical decision task conducted with 816 participants) were shorter for words associated with unmarked concepts ($t(42)=3.2$, $p<0.003$). This suggests that unmarked concepts may be cognitively more accessible.

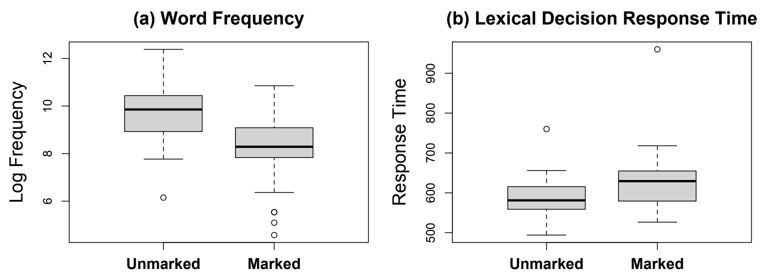

Figure 2. (a) English words corresponding to unmarked concepts have higher word frequencies (L&B data), and are responded to more quickly by English speakers (ELP data).

---

[2] We used the logarithm of link number because of skew. For example, the article for 'animal' has 127,523 connections with other articles, with the next-highest 'moon' only having 6,340. The overall result still holds if 'animal' is excluded ($t(42)=2.38$, $p<0.03$) or if a non-parametric test is used (Wilcoxon signed rank test, $V = 714$, $p<0.01$).

2.4. *Textual diversity*

Following up on the idea that words acquire more shades of meaning by virtue of being used in diverse contexts (see discussion above), we look at contextual diversity, another measure taking from the ELP. Contextual diversity is calculated based on a 10.2 million word corpus of 1,867 British films. On average, words associated with unmarked concepts occurred in 23% of the films, whereas words associated with marked concepts occurred in only 6% of them ($t(43)=4.77$, $p<0.0001$). A look at the number of distinct word senses listed on Merriam-Webster Online (www.merriam-webster.com) reveals that words corresponding to unmarked concepts also tended to have more different senses ($t(44)=2.96$, $p<0.005$), supporting the idea that words occurring in diverse contexts tend to become semantically extended.

3. Discussion

What motivates the direction of semantic change? The preceding discussion has highlighted a number of different variables that correlate with the cross-linguistic data from Urban (2011). These variables are themselves highly correlated with each other, so it is difficult to tease apart individual contributions. One way to look at the variables together is to use conditional inference trees (Hothorn, Hornik & Zeileis, 2006). This approach shows the frequency variable to be most predictive of the cross-linguistic asymmetry (frequency is the only split in the tree).

In addition, we can compare models using Aikaike's Information Criterion (AIC). Using the standard cut-off criterion of ΔAIC=6 (Richards, Whittingham & Stephens, 2011), word frequency, context diversity and free word association asymmetry have about equal support from the data. Lexical decision response times, the number of Webster word senses and the Wikipedia link sum have significantly less support. Hence, these variables are less strongly associated with the cross-linguistic data.

Table 2. Aikaike's Information Criterion for models with the different variables.

Variable	*AIC*
Kučera & Francis Frequency	-26.68
Lund & Burgess Frequency	-25.27
Context Diversity	-25.65
Lexical Decision RTs	-16.35
No. of Word Senses	-6.13
Association Asymmetry	-25.58
Wikipedia Link Sum	-2.68

The analyses suggest that concepts that are more frequent and cognitively accessible serve as the basis for other concepts. This idea has a long tradition in research on grammaticalization and metaphor. Here, we provide empirical evidence that unmarked concepts such as 'skin' have more frequent and more accessible words associated with them in English. The asymmetry in the behavioral free word association data might be a reflection of this difference in cognitive accessibility. It thus appears to be the case that some words, by virtue of coming to mind more quickly, are more frequently re-used to express other concepts.

It is also relevant that we additionally found effects for context diversity and number of word senses. This sub-result points to an additional mechanism by which words that occur in more different contexts acquire more different senses. That this mechanism is very likely another important factor in semantic change is also shown by the fact that contextual diversity is one of the most predictive variables.

Note that these mechanisms are more general explanations than any attempt to form generalizations on the nature of the semantic relations seen in Table 1. In Urban's set of proposed candidates for directional change, unmarked and marked concepts are characterized by many different semantic relations, some of which are part/whole, some of which appear to be metaphorical extensions, some of which are based on contiguity. Thus, there is a diversity of different semantic relations and no straightforward unity. We showed that across these different semantic associations between concepts, there are consistent quantitative patterns that are reflected in synchronic English data.

More generally, the close correspondence of word frequency data, dictionary data, human association data, response times, and encyclopedic network structure with the cross-linguistic results in Urban (2011) shows that we can indeed use synchronic data to understand diachronic patterns (see also Jäger & Rosenbach, 2008). The evolution of word meaning—as evidenced by cross-linguistic patterns in the world's languages—thus becomes a phenomenon that is not always unpredictable and haphazard, but one that has directional tendencies that are cognitively motivated. This also points to directionalities of lexical semantic change with respect to the origins of language: From the perspective of this study, we would expect the earliest words to be for more frequent, accessible and more general concepts; and we would expect the evolution of word meaning to start from there.

References

Balota, D.A., Yap, M.J., Cortese, M.J., Hutchison, K.A., Kessler, B., Loftis, B., Neely, J.H., Nelson, D.L., Simpson, G.B., & Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39, 445-459.

Calude, A. S., & Pagel, M. (2011). How do we use language? Shared patterns in the frequency of word use across 17 world languages. *Phil. Trans. R. Soc. B.*, 366, 1101-1107.

Hill, F., Korhonen, A., & Bentz, C. (2013). A quantitative empirical analysis of the abstract/concrete distinction. *Cognitive Science*, ahead of press, 1-16.

Hothorn, T., Hornik, K., & Zeileis, A. (2006). Unbiased Recursive Partitioning: A Conditional Inference Framework. *Journal of Computational and Graphical Statistics*, 15, 651-674.

Jäger, G., & Rosenbach, A. (2008). Priming and unidirectional language change. *Theoretical linguistics*, 34, 85-113.

Kučera, H., & Francis, W. (1967). *Computational analysis of present day American English*. Providence, RI: Brown University Press.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28, 203-208.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. http://w3.usf.edu/FreeAssociation/ [accessed August 22, 2013].

Richards, S. A., Whittingham, M. J., & Stephens, P. A. (2011). Model selection and model averaging in behavioural ecology: the utility of the IT-AIC framework. *Behavioral Ecology and Sociobiology*, 65, 77–89.

Thompson, G. W., Kello, C. T., & Montez, P. (2013). Searching Semantic Memory as a Scale-Free Network: Evidence from Category Recall and a Wikipedia Model of Semantics. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 3533-3538). Austin, TX: Cognitive Science Society.

Urban, M. (2011). Asymmetries in overt marking and directionality in semantic change. *Journal of Historical Linguistics*, 1, 3-47.

Witten, I., & Milne, D. (2008). An effective, low-cost measure of semantic relatedness obtained from Wikipedia links. In *Proceeding of AAAI Workshop on Wikipedia and Artificial Intelligence: an Evolving Synergy* (pp. 25-30). Chicago: AAAI Press.

Zipf, G. K. (1949). *Human behaviour and the principle of least effort: An Introduction to Human Ecology*. New York: Hafner.