To appear in *Cognitive Psychology*. This file is a pre-print and may contain errors or omissions not present in the final published version.

Evolution of word meanings through metaphorical mapping:
Systematicity over the past millennium

Yang Xu

Department of Linguistics, Cognitive Science Program, University of California, Berkeley

Barbara C. Malt

Department of Psychology, Lehigh University

Mahesh Srinivasan

Department of Psychology, Cognitive Science Program, University of California, Berkeley

Correspondence concerning this article should be addressed to Yang Xu, 1203 Dwinelle Hall,
Department of Linguistics, University of California, Berkeley, Berkeley, CA 94720. Contact:
yang_xu_ch@berkeley.edu

Abstract

One way that languages are able to communicate an unbounded set of ideas through a finite lexicon is by *compressing* meanings into words, such that individual words carry multiple, related senses of meaning. We propose that overarching communicative and cognitive forces have created *systematicity* in how new *metaphorical* senses have developed from existing word senses over the history of English. To investigate this idea, we used computational models to test whether some domains are more likely to be the starting point of metaphorical extensions (source domains) and others the end-points (target domains) over the past millennium. We found that a compact set of variables including externality and embodiment account for approximately 90% of the directions of recorded metaphorical mappings between domains, dating back 1300 years. These results provide the first large-scale evidence that metaphorical mappings are systematic, and suggest that metaphor provides an efficient device for compressing emerging ideas into an existing lexicon.

Keywords: word meaning, semantic change, polysemy, metaphorical mapping, language evolution

Evolution of word meanings through metaphorical mapping:

Systematicity over the past millennium

Introduction

Every natural language faces the challenge of having to express an unbounded range of ideas through a finite set of words (cf. von Humboldt, 1836; Chomsky, 1965). One way in which languages meet the challenge of maintaining a compact lexicon is by compressing multiple ideas into individual word forms. The most common form of compression in the lexicon, both in English and in other languages, is *polysemy*: Cases where a single word has multiple, related senses of meaning (e.g., Breal, 1897; Brugman, 1988; Geeraerts, 1997; Lakoff & Johnson, 1980; Pustejovsky, 1995; Srinivasan & Rabagliati, 2015; Sweetser, 1991). One prominent mechanism that generates polysemy is *metaphorical mapping* (Lakoff & Johnson, 1980), whereby new word meanings are created by mapping an existing sense of a word from its own *source* domain to another *target* domain, based on structural similarities between the domains. For example, the English word *grasp* originally conveyed a physical action (source), as in “grasp a fruit,” and was later extended to express an abstract sense of understanding (target), as in “grasping an idea”, thus construing ideas as objects that can be held and controlled. In the present study, we test whether there is *systematicity* in how new metaphorical senses have developed from existing word senses over the history of English, shaped by communicative and cognitive pressures.

Consistent with the idea that metaphorical sense extension is systematic, prior research (e.g., Lakoff & Johnson, 1980; Sweetser, 1991) has suggested that mappings between semantic domains tend to be asymmetric, and occur in one direction but not the other (e.g., for *grasp*, it is action→knowledge but not knowledge→action). However, directionality in historical metaphorical mappings has not been evaluated *at scale* against records of historical semantic change, leaving open whether metaphorical mappings reflect predictable patterns. Using an historical corpus of the English lexicon, we seek to predict – at a large scale – the historical directions of metaphorical mapping by testing whether some

domains are more likely to have been the starting point of historical extensions (the source domains) while others have been more likely to be the end-points (target domains). For example, in the case of *grasp*, the “physical action” sense first appeared in English around AD 1300, preceding the abstract sense of “understanding” that emerged around AD 1600 (Christian, Roberts, Samuels, Wotherspoon, & Alexander, 2015a). Our analysis examines the predictability of such historical metaphorical sense mappings in the English lexicon.

Directions of metaphorical mapping need not be predictable: Instead, the evolutionary path of meaning change could be varied and unpredictable across words. Thus, although some words, like *grasp*, may have begun with a relatively concrete sense that was subsequently extended to a more abstract sense, other words may have been extended in the opposite direction, from abstract to concrete senses. For instance, *irritable* initially conveyed an abstract sense of “anger,” as in an “irritable person” (AD 1662; Christian et al., 2015a), but was subsequently used to convey a more concrete, physical meaning, as in “irritable skin” (AD 1791; Christian et al., 2015a). A particular word's path could be shaped by the ever-changing cultural conditions and communicative needs that have pressured words to develop new meanings across history (cf. Aitchison, 2001). The contributing factors may be variable across words and across eras of history, resulting in little predictability on a large scale.

However, an alternative possibility is that metaphorical sense extension is predictable, because it is shaped by enduring functional pressures on language evolution. Recent work in computational cognitive science has suggested that many aspects of language and cross-linguistic variation can be understood in terms of general design principles, such as the need for linguistic structures to support informative communication and language learning (Zipf, 1949; see also Kemp & Regier, 2012; Kirby, Tamariz, Cornish, & Smith, 2015; Piantadosi, Tily & Gibson, 2011; Regier, Kemp, & Kay, 2015). Applying this perspective to polysemy, systematicity in mappings may result from language users minimizing the costs associated with communicating novel ideas and learning a lexicon: Mappings will tend to be driven by their ease of construction, their effectiveness as communication devices, and/or

their learnability by children. They should thus facilitate accurate inferences of novel senses at a broad scale, rendering an *efficient mechanism* for compressing new ideas into an existing lexicon.

But this perspective says little about *what kinds* of metaphorical mappings will be constructed. An independent line of research, from cognitive psychology and linguistics, provides suggestions about the specifics: Namely, that the processes that give rise to metaphorical polysemy may reflect conceptual structure (Lakoff, 1987; Sweetser, 1991). One prominent proposal in this vein is *Conceptual Metaphor Theory* (CMT; Lakoff & Johnson, 1980; Reddy, 1979), which posits that people reason about abstract concepts – such as “understanding” – via metaphorical mappings from knowledge domains that are more concrete and tied to bodily experience – such as “physical action.” Although CMT is controversial as a theory of cognition, other theories also predict that concrete and embodied word senses will be extended to more abstract senses because this may facilitate communication and learning (Murphy, 1997; Srinivasan & Carey, 2010; Thibodeau & Durgin, 2008). For example, it could be easier to communicate about something abstract – for which achieving shared reference is difficult – by analogically referring to a more concrete meaning, for which shared reference is easier. Further, because concrete and embodied meanings tend to be frequent (Hanley et al., 2013; cf. Winter, Thompson, & Urban, 2014), they may be more readily retrieved as sources for meaning extension by speakers who need to communicate a new idea.

Prior research has not explicitly linked the perspective of efficient language design with the study of the cognitive foundations of polysemy or metaphorical mapping. We bridge this gap by analyzing a large set of metaphorical mappings between source and target domains spanning 1300 years, dating from Old English to the present. To our knowledge, this is the first large-scale study to evaluate the predictability of polysemous metaphorical mappings against the record of historical change in a lexicon. We expect that if metaphorical mapping provides an efficient cognitive mechanism for compressing ideas into a communicative and learnable lexicon, the historical development of word senses should be highly predictable, and explained by a compact set of variables.

Material and methods

We identified six candidate variables from the literature that could explain directions of metaphorical mapping based on communicative and cognitive considerations. We used these variables to predict the directions of metaphorical domain mappings recorded in the history of English. In the following, we describe: 1) Each of the candidate variables, 2) The historical dataset of metaphorical domain mappings we used, and 3) A survey through which we elicited ratings of the domains on each of the candidate variable dimensions.

2.1 Candidate variables

Three of the variables we considered as predictors are suggested by Conceptual Metaphor Theory (Lakoff & Johnson, 1980; Reddy, 1979; Sweetser, 1991) and supported by other perspectives suggesting that certain types of meanings should be the sources of metaphorical extensions (Srinivasan & Carey, 2010; Thibodeau & Durgin, 2008), as discussed above.

1. *Concrete* → *Abstract*. Word senses that refer to things perceived through the sensory systems should serve as a source of metaphorical mappings, relative to ones labeling less perceptible referents.

2. *Embodied* → *Disembodied*. Word senses that refer to things that are more directly experienced through our bodies should serve as a source of metaphorical mappings, relative to ones that label referents that are less directly experienced.

3. *External* → *Internal*. Word senses that refer to entities in the external world should serve as a source of metaphorical mappings, relative to ones that label internal, mental entities, such as emotions or feelings (Traugott, 1990).

An additional three variables were identified, based on the idea that metaphorical extensions tend to be communicatively expressive:

4. *Animate* → *Inanimate*. Word senses that refer to animate entities may carry more expressive power than ones referring to less animate entities, and may thus serve as a source of metaphorical mapping to convey salient features of target inanimate entities (Silverstein, 1976; Traugott, 2003).

5. *Less valenced* → *More valenced*. If metaphorically-derived senses arise in part because of their expressive power, we might expect derived senses to be more emotionally valenced than originating senses. This predicts that when word senses from source domains are extended, the resulting senses in the target domain will have stronger valence (whether positive or negative) (Ullmann, 1957; Osgood, Suci, & Tannenbaum, 1964; cf. Hollis & Westbury, 2016).

6. *More Intersubjective* → *Less intersubjective*. Intersubjectivity refers to the degree to which people experience something the same way (Traugott, 1989, 2003). Word meanings that are more intersubjective (e.g. *wooden*) would be easier to understand and establish shared labels for than ones that are less so (e.g., *beautiful*). They might thus serve as a source of metaphorical mappings, relative to less intersubjective meanings.

Some of these variables are likely to be correlated, e.g., many of the same word meanings could be external, concrete, intersubjective, and embodied. However, they are not identical; For example, *arm* and *table* may both be highly concrete, but arms are more directly and viscerally experienced than tables. Our analyses tease apart the relationships among these variables in accounting for metaphorical mappings in the history of English.

2.2 Historical data

We obtained data from the Metaphor Map of English (MME) database (Christian, Roberts, Samuels, Wotherspoon, & Alexander, 2015b), which identifies metaphorical mappings of senses classified by semantic domains over more than a millennium (see *Supplemental Material* for more details on the domains and the database). The MME database is derived from the Historical Thesaurus of English (HTE; Christian et al., 2015a), the only historical thesaurus in the world. The HTE is, in turn, based on the second edition of the Oxford English Dictionary (Simpson & Weiner, 1989).

All word senses in the MME database are classified into 415 distinct semantic domains (e.g. “textiles,” “digestive organs,” “anger,” “pride”) in rough accordance with the standard of semantic classification adopted by the HTE, as taken from the second edition of the Oxford English Dictionary.

All semantic domains from the database are sub-fields from three broader ontological domains: I. The external world; II. The mental world; III. The social world. Section I includes semantic domains that concern readily observable phenomena of the universe, such as the earth and sea. Section II includes domains that concern cognitive processes such as perception and emotion. Section III includes domains that concern social systems such as the law and morality.

Each specific semantic domain defines a category of meaning that word senses can evolve from (source domain) or towards (target domain), and a single domain can include multiple but distinct word senses. The MME database records metaphorical mappings that have occurred between domains and provides a sample of words that exemplify these domain mappings. Our analyses examined historical metaphorical mappings at the level of semantic domains, and examined whether there has been systematicity in the directionality of domain mappings.

The data set we analyzed presents approximately 6000 recorded metaphorical sense mappings among 400 semantic domains (see *Supplemental Material* for details on these domains and the data), where each attested domain-domain mapping reflects the historical extension of word senses between two domains. Together, these mappings span an 1100 year period, from the Old English period around AD 800 through to the present era, around AD 1950. Within the MME database, metaphorical mappings are summarized separately for the Old English period (i.e., before AD 1100), and in 50-year steps for the subsequent 800 years, providing 18 unique historical time points for our analysis.

The database indicates the historical direction of sense extension among each pair of semantic domains. Three types of directions are recorded: 1) $A \rightarrow B$, i.e. A is source domain and B is target domain; 2) $A \leftarrow B$, i.e. A is target and B is source; 3) $A \leftrightarrow B$, i.e. a bidirectional mapping. Our analyses examined whether these relationships among semantic domains could be predicted from empirical ratings of the domains along the six candidate variable dimensions, based on an on-line survey that we describe next.

2.3 Survey for rating elicitation

We elicited ratings of the 400 semantic domains along the six candidate variables through an on-line survey (see *Supplemental Material* for experimental instructions). For each variable, participants rated each of the 400 domains on a seven-point scale. Participants first read a definition of each domain (e.g., “Plant - A living thing that grows in the ground, usually has leaves or flowers, and needs sun and water to survive”), assembled by consulting the Merriam-Webster Dictionary. Participants then rated the domain on one of the six variables. For example, for concreteness, participants selected a number between 1 and 7, where 1 represented “highly abstract,” 7 “highly concrete,” and 4 “intermediate.” Because of the large number of ratings we needed to elicit (400 domains x 6 variables), each participant only rated a random sample of 40 domains along a single variable dimension.

Our research was approved by the Institutional Review Board at the University of California, Berkeley. Data were collected from 1439 participants who self-reported as native English speakers, using the online Qualtrics experimental platform (<http://www.qualtrics.com/>), disseminated via Amazon’s Mechanical Turk (<https://www.mturk.com/mturk/>). Following standard protocol (Oppenheimer, Meyvis & Davidenko, 2009), responses were excluded from participants who did not respond correctly to three “catch” trials assessing attention to the task ($n = 94$). On average, we obtained 18 ratings ($SD = 2$) for each of the domain-variable questions. The average standard deviation of ratings was 1.6 ($SD = 0.32$), indicating a high level of inter-rater agreement (see *Supplemental Material* for the mean variable ratings for all 400 domains).

Results

To evaluate the hypothesis that metaphorical mapping in history is systematic, we performed three computational analyses: 1) We tested whether directions of metaphorical extensions recorded in the historical data set can be predicted by ratings along the six variable dimensions described above, either singly or in combination; 2) We then examined whether the same ratings can be used to identify

the semantic domains that have served as sources vs. targets of metaphorical extensions; 3) Finally, we examined the best predictor of metaphorical mappings across epochs to see if its explanatory power is consistent over time. Additional details of our computational models and methods can be found in the *Supplemental Material*.

3.1 Analysis I: Predicting directions of metaphorical mappings over history

For each of the candidate variables, we constructed a parameter-free model that predicts the direction of mapping between a pair of semantic domains based on average ratings of these two domains along a variable of interest (see *Supplemental Material* for more details). For example, for the animacy variable, we predicted the direction of mapping to be Animate $\rightarrow$ Inanimate, such that the domain that was rated as more animate on average would serve as the source, and the domain that was rated as less animate on average would serve as the target. Our models also predicted bidirectional mappings (which accounted for 6.76% of the available mappings) if the average ratings of the two domains were equal. Since the valence predictor is polarized (i.e. a domain can be either positive or negative), we calculated the absolute value of its ratings, and for this model predicted the direction of mapping between domains to follow Less valenced $\rightarrow$ More valenced. Thus, a domain rated ‘7’ (highly positive) would be treated as equally-valenced as a domain rated ‘1’ (highly negative). Table 1 specifies the predictions of each model along with example mappings from the historical database that the models successfully predicted.

Table 1. *Examples of metaphorical sense mappings predicted by each model (ordered alphabetically).*

Model	Source Domain (rating) $\rightarrow$ Target Domain (rating)	Sample Words	Attested Period
Animacy	Male person(6.5) $\rightarrow$ Courage(2.0)	manly	Old English
<i>Animate $\rightarrow$ Inanimate</i>	Reptiles(6.4) $\rightarrow$ Bad(1.9)	serpent	1350-1400

Concreteness	Reflection (of light) (6.0) → Virtue(1.6)	clear	1350-1400
<i>Concrete → Abstract</i>	Plants(6.8) → Pride(2.5)	flourish	1350-1400
Embodiment	Sense & speech organs(6.5) → Rivers & streams(2.6)	mouth	Old English
<i>Embodied → Disembodied</i>	Strength(5.9) → Politics(2.2)	stalwart	1850-1900
Externality	Light(6.4) → Thought(1.1)	reflect	1550-1600
<i>External → Internal</i>	Textiles(6.7) → Belief & opinion(1.0)	spin	1950-2000
Intersubjectivity	Textiles(5.7) → Difficulty(3.3)	mesh	1500-1550
<i>Agreeable → Disagreeable</i>	Birth(5.4) → Faith(3.3)	reborn	1550-1600
Valence	Relative position(4.1) → Excitement(6.4)	up	1300-1350
<i>Neutral → Valenced</i>	Semi-fluidity(4.3) → Moral evil(1.3)	slime	1550-1600

Note. Each row shows 1) a type of model, 2) examples of metaphorical domain-domain mappings consistent with the predictions of that model, 3) sample words that have senses within each of the domains, and 4) the first attested time period in which senses exemplifying the domain-domain mapping emerged.

Each model made a prediction about the direction of extension (the source-target relation) at different historical epochs for a given pair of domains A and B, out of three possibilities: $A \rightarrow B$, $A \leftarrow B$, or $A \leftrightarrow B$. To establish a baseline, we used a baseline model that predicted the direction of mappings by random guessing. The results appear in Figure 1a. Overall, all models predicted directional change above chance. In particular, externality (accuracy = 68.1%) and concreteness (67.6%) were roughly equivalent to one another, with each exceeding chance-level accuracies (33.3%) by twofold. These were followed by intersubjectivity, valence, and embodiment with accuracies all above 50%, while animacy (48.8%) explained the data the least.

Figure 1. Summary of results on predicting historical directions in metaphorical sense mapping:

a) Predictive performance (illustrated by the relative proportion of counts of successfully predicted domain-domain mapping directions in black bars with respect to total counts of recorded domain pairs in white bars) from individual variables at each available epoch in history. “OE” stands for Old English, b) Correlations among average empirical ratings of the variables, c) Unique variance explained by the variables.

To verify that the specific predictive directions we proposed for these models (e.g., Concrete → Abstract) were indeed more accurate than their opposite (e.g. Abstract → Concrete), we also ran these models by reversing their predictive arrows. In each of these cases, we found that the predictive accuracy was substantially closer to chance: animacy (44.2%), concreteness (25.4%), embodiment (40.5%), externality (25.0%), intersubjectivity (35.9%), valence (37.0%). Together, these results indicate that historical metaphorical mappings are systematic and predictable, such that words are more likely to be extended from some domains toward other domains, compared to the reverse.

To better interpret these findings, we next tested whether the six variables act in complementary ways and account for unique variance in explaining directions of metaphorical extension. Figure 1b

shows the Pearson correlations between all pairs of variables. As can be seen, externality and concreteness were most strongly inter-correlated ($r = .84, p < .001$). For this reason, we expected these variables to make similar predictions about mapping directions and for one to make little independent contribution to explaining the directionality of metaphorical mappings over the other. Intersubjectivity was most strongly correlated with both externality ($r = 0.44, p < .001$) and concreteness ($r = .47, p < .001$) among the remaining pairwise correlations. Animacy, embodiment, and valence were generally less correlated with other variables.

To take into account these correlations, we applied an iterative residual fitting procedure to analyze the total variance explained from all of these variables as follows. First, we found the variable that best predicted the directions of all available metaphorical mappings. We then iteratively searched for the variable that best predicted the remaining mappings until there existed no further variables. Figure 1c summarizes the total variance explained via this procedure. In total, the variables accounted for over 90% of all available mappings.¹ Externality stood out as the single most accurate predictor, explaining 68% of mappings. Embodiment (14%) and valence (5%) explained the majority of the remaining data not predicted by externality, while animacy accounted for 2.3%. Finally, concreteness and intersubjectivity each explained no more than 1% of data. This is not to suggest that these latter variables play no role in predicting metaphorical mapping, but only that their power in explaining the residual variance was likely subsumed by the externality variable, with which they were highly inter-correlated.

Together, our findings indicate that externality, embodiment, and valence define a relatively orthogonal space that can help explain the directionality of the majority of metaphorical mappings in

¹ Our models account for about 90% of the variance in metaphorical mappings, but this is actually a conservative estimate of their explanatory power. The majority of mappings that the models did not explain were bidirectional mappings (which account for over 76% of the unexplained variance), and in fact our models explained about 98% of the unidirectional mappings. The difficulty with bidirectional mappings might stem from the fact that our models involved minimal assumptions and could only posit bidirectional mappings in a very restricted set of cases (i.e., when the average ratings for the two domains were identical).

the historical dataset. These findings provide the first large-scale evidence for the idea that metaphorical senses have developed across history in systematic and predictable ways, constrained by a relatively compact set of psychological dimensions.

3.2 Analysis II: Identifying the source and target domains of metaphorical mapping

Our analyses so far have focused on metaphorical mappings at different points in history, but the outcome of these diachronic changes should also be reflected in more global, time-independent asymmetries as to which domains tend to be the sources of metaphorical extension, and which the targets. We next examined the extent to which the candidate variable dimensions explain the source-target asymmetry of specific semantic domains.

To address this question, we first specified a quantifiable indicator of source-target asymmetry for each semantic domain. To do so, we created a source-target asymmetry index, formally defined as the difference between two probabilities: $p(\text{source}) - p(\text{target})$. Here, $p(\text{source})$ is the probability of a domain serving as a source of metaphorical mappings across all time points in history (i.e., a synchronic summary of “source” strength), and $p(\text{target})$ is the probability of a domain being a target of metaphorical mappings across time points. A high positive value for the asymmetry index indicates a strong source domain, whereas a low negative value indicates a weak source (or strong target) domain (see *Supplemental Material* for more details on the asymmetry index). Table 2 lists the strongest source and target domains in the dataset based on this measure.

Table 2. *Strongest source and target domains across history*

Strongest sources	Strongest targets
Textiles	Excitement
Supernatural	Pride
Digestive organs	Anger
Hardness	Hatred and hostility

Softness	Bad
Ruminants	Behaviour and conduct
Cultivated plants	Money
Wetness	Literature
Darkness	Fear
Solidity and density	Vigorous action and degrees of violence

We then computed Pearson correlations between each domain's asymmetry index and its mean empirical rating for each of the variable dimensions. Figure 2a summarizes the results. Ratings for all variable dimensions except for those of animacy ($r(400) = .09, p > .250$) correlated significantly with the asymmetry indices ($p < .01$ for the five predictive variables). Specifically, externality ($r(400) = .46, p < .001$) and concreteness ($r(400) = .42, p < .001$) showed the highest strengths of correlation, consistent with their leading roles in predicting directions of metaphorical mappings. Intersubjectivity ($r(400) = .22, p < .001$) and embodiment ($r(400) = .13, p < .009$) were also moderately correlated with asymmetry indices. These results suggest that domains that are external, concrete, embodied, and share a conceptualization across people, tend to serve as sources of metaphorical mappings across history. Finally, valence negatively correlated ($r(400) = -0.19, p < .001$) with asymmetry, suggesting that more valenced domains tend to be targets of metaphorical mappings, also consistent with our previous findings.

Figure 2. Summary of results on identifying source and target domains of metaphorical sense mapping:

a) Predicting the source-target asymmetry of domains from individual variables, b) The relative weighting of these variables as estimated by multiple linear regression.

Again taking into account the inter-correlated nature of the variables, we performed a multiple linear regression to estimate the relative contributions of the six variables. We expected the most important (and orthogonal) predictor ratings to be weighted the highest when regressed against the asymmetry indices. Figure 2b shows the relative contributions of these variables as reflected by their weights. The regression fit had a Pearson $r(393) = 0.42$, $p < .001$. The candidate variable with the largest absolute weight was externality ($p < .001$, $\beta = .007$, 95% CI [0.0037, 0.0093]), followed by embodiment ($p < .09$, $\beta = .0019$, 95% CI [-0.0002, 0.0040]) and valence ($p > .250$, $\beta = -.001$, 95% CI [-0.0026, 0.0007]), although neither of the latter two variables contributed significantly beyond externality. This result is largely consistent with the results from our residual predictive analysis, which found that the same set of variables accounted for most of the variance in explaining metaphorical mapping directions during different historical epochs.

Figure 3 provides two concrete examples illustrating the power of the externality variable in

explaining the direction of metaphorical mappings. “Water” was rated as a highly external domain (mean externality rating = 6.2 out of 7) and as shown, almost all metaphorical mappings involving this domain in our dataset consisted of projections from this domain to other domains, exemplifying the principle that external domains tend to act as sources of metaphorical mappings across history. In contrast, “Mind” was rated as a highly internal domain (mean externality rating = 1.4) and as shown, almost all mappings involving this domain were projections from other domains to this domain, exemplifying the principle that internal domains tend to be the targets of metaphorical mappings.

Taken together, our diachronic and synchronic analyses suggest that metaphorical sense extensions are highly systematic and constrained by a compact set of variable dimensions.

Figure 3. Visualization of historical metaphorical mappings in two domains: “water” (Panel A) and “mind” (Panel B). Arrows between domains indicate historical directions of metaphorical mapping. The domains in red have served either as sources or targets to the domain in question. (The year in parentheses indicates the time period of the first appearance of the domain-domain mapping in history.) The domains labeled in black are randomly chosen reference domains, a subset of the 400 domains which are shown as gray dots projected onto a low-dimensional space via principal components analysis using the empirical ratings for the six variables. The two principal components correspond roughly to a concrete-to-abstract dimension (as read from right to left in the horizontal dimension), and an embodied-to-disembodied dimension (as read from top to bottom in the vertical dimension).

3.3 Analysis III: Assessing externality across time

Our results so far have indicated that externality is the best predictor of the direction of historical metaphorical mappings. As discussed before, more external domains could tend to provide a source of metaphorical mappings if they are the most cognitively available and accessible to extend further. To further understand if the External → Internal mapping is a privileged direction of semantic change, we tested whether external domains have remained a strong source of metaphorical mappings throughout history. An alternative possibility is that externality emerged as a strong predictor only at the very earliest points in history (e.g., because external ideas tend to be lexicalized earlier), and thus serve as sources of metaphorical mappings during only these periods.

To distinguish between these two possibilities, we performed a temporal analysis that examined the strength of external domains as sources over times. Specifically, we divided the available domains in our data set into “external” and “internal” depending on whether the ratings of externality for each domain were greater or less than “4” (the middle scale value between highly external “7” and highly

internal “1”; our results obtained regardless of the exact threshold used to separate external and internal domains). Then, we computed the source probability (described earlier) for each of the two groups of domains at each of the available historical time points provided in the corpus. If externality serves as a strong predictor for metaphorical mappings mainly due to the lexicalization of external concepts at the earliest points in history, the source probability for external domains should on average be higher than for internal domains, but this difference should decrease over time. On the other hand, if External → Internal is a privileged direction of semantic change, external domains should have greater source probabilities than the internal domains across time.

Our analyses support the second possibility (see Figure 4). Specifically, external domains had significantly higher source probabilities than internal domains throughout the 1300-year period (joint $p < .001$, $n = 18$ time points with Fisher's method; time point with minimal $p < 7.71\text{e-}011$, $t(396) = 6.57$; time point with maximal $p < .035$, $t(396) = 1.81$). Also, source probabilities for the external domains did not differ significantly between the initial and end periods ($p > .250$, $t(476) = -.29$).

source probabilities over time for domains partitioned along three variable dimensions: *externality*, *embodiment*, and *valence*. The vertical bars represent standard errors above and below the mean. The dots and stars above the bars indicate significant differences between high and low groups (at a given time point) at $p < .01$ (uncorrected) and $p < .01$ (Bonferroni-corrected for multiple comparisons across all time points).

Figure 4 shows a comparison of external domains against embodied (vs. less-embodied) and valenced (vs. less-valenced) domains, results for which were obtained via parallel analyses to those used for externality.² In contrast with externality, we found that embodied and valenced domains did not serve as sustained sources for metaphorical mappings across time. Specifically, more embodied domains tended to facilitate metaphorical sourcing earlier in history (possibly because more embodied domains were lexicalized earlier in time), whereas more valenced domains acted as weak sources of metaphorical mappings, but only later in historical time. Together, these observations are consistent with the idea that External → Internal is a privileged basis for metaphorical mapping, with concepts from more external domains providing a means for communicating about concepts from more internal domains.

Discussion and conclusions

The present study focused on one of the key mechanisms through which words gain new meanings: metaphorical mapping. By drawing on an historical record of semantic change in English, we evaluated – on a large scale – whether new metaphorical word senses have been created in systematic and predictable ways. We discovered that a compact set of variables including externality and embodiment account for approximately 90% of the directions of recorded metaphorical mappings between domains, dating back 1300 years. Furthermore, we found that externality – the strongest predictor – explained the direction of metaphorical mappings throughout the history of English, and not just for an initial period early in English vocabulary development.

Our findings lend strong support to the idea that the evolution of metaphorical polysemy proceeds in systematic ways. The six predictor variables that we explored were motivated by communicative and cognitive considerations, drawing upon theories of cognition and language. Our critical assessment of their contributions in accounting for historical sense extension illuminates the

² We found concreteness to be highly similar to externality, and we found no statistical significance for animacy and intersubjectivity. For conciseness, we omit presenting these results.

nature of metaphorical language. In particular, the fact that metaphorical sense extension can be predicted by a compact set of variables suggests that metaphor provides an efficient way of compressing emerging meanings into an existing lexicon, without requiring the construction of word forms *de novo*. Our work thus extends previous studies about communicative and cognitive constraints on synchronic features of language such as word length (Piantadosi, Tily & Gibson, 2011; Zipf, 1949) and semantic structures (Kemp & Regier, 2012; Regier, Kemp & Kay, 2015) to explain the evolution of metaphorical polysemy in the lexicon (cf. Geeraerts, 1997; Blank, 1999).

In many respects, our finding – that historical metaphorical sense extensions are predictable – is surprising. Due to the ever-changing cultural conditions and communicative needs that have pressured words to develop new meanings (cf. Aitchison, 2001), metaphorical sense extensions could have in principle been varied and unpredictable across words. It is clear that cultural conditions and communicative needs have shaped metaphorical extension in many cases. For instance, *mouse* would never have been extended to a computer accessory if the invention of the computer had not come about, nor would *file* or *folder* have been extended from objects held in the hand to virtual ones without specific cultural developments. Our data demonstrate that, in spite of the effects of complex and dynamic cultural conditions, there is an underlying regularity to metaphorical sense extension driven by measurable communicative and cognitive principles.

Our results leave open exactly why variables such as externality were reliable predictors of metaphorical mappings across history. On the one hand, these results could be taken to support prominent accounts of cognition such as Conceptual Metaphor Theory (Lakoff & Johnson, 1980; Reddy, 1979). By these accounts, ideas tied to human experience and sensation provide anchors for thinking about more abstract ideas. This predicts that more external (and embodied) word meanings will be well understood, and as such, will provide a source of mappings for word meanings that are more difficult to understand. Alternatively, our results may not reveal that abstract concepts are understood in terms of external and embodied concepts, *per se*, but only that external and embodied

word meanings are easier to achieve shared reference for (Srinivasan & Carey, 2010; Thibodeau & Durgin, 2008), and are more frequent and thus easier to retrieve for speakers who wish to communicate new idea (Hanley et al., 2013; cf. Winter, Thompson, & Urban, 2014). Future work should evaluate these accounts of metaphorical extension further.

Our work also creates many questions for future research in computational approaches to semantic change and polysemy. For example, our analysis examined sense change at the level of semantic domains, but further analyses could assess the systematicity of semantic change at the individual word level, which would require a different source of data. Future work could also explore other aspects of lexical semantic change such as semantic chaining (Lakoff, 1987; Xu, Regier & Malt, 2015), the socio-cultural factors that explain the rise or fall in communicative needs for different semantic domains, cognitive machineries other than metaphor that support polysemous extension, and polysemous extension in other languages. The current work provides an empirical approach for in-depth explorations into the evolution and nature of meaning.

Acknowledgements

We thank Eve Sweetser and George Lakoff for helpful input and Jessie Salas and Elizabeth Kaplan for assistance with data collection. This project was funded by NSF award SBE-1041707 to the Spatial Intelligence and Learning Center (SILC) and NSF award SBE-16302040 to MS.

References

- Aitchison, J. (2001). *Language change: Progress or decay?* Cambridge: Cambridge University Press.
- Breal, M. (1897). *Essai de semantique: Science des significations*. Paris: Hachette.
- Blank, A., & Koch, P. (1999). *Historical semantics and cognition*. Walter de Gruyter.
- Brugman, C. M. (1988). *The story of over: Polysemy, semantics, and the structure of the lexicon*. New York: Garland Press.
- Chomsky, N. (1957). *Syntactic Structures*. Mouton, The Hague, Paris.
- Christian, K., Roberts, J., Samuels, M., Wotherspoon, I., & Alexande, M. (2015a). *The Historical Thesaurus of English, version 4.2*. Glasgow: University of Glasgow
(<http://historicalthesaurus.arts.gla.ac.uk/>)
- Christian, K., Roberts, J., Samuels, M., Wotherspoon, I., & Alexander, M. (2015b). Metaphor Map of English. In *Mapping Metaphor with the Historical Thesaurus*. Glasgow: University of Glasgow. (<http://mappingmetaphor.arts.gla.ac.uk>)
- Geeraerts, D. (1997). *Diachronic prototype semantics*. Oxford: Oxford University Press.
- Hanley, J. R., Hunt, R. P., Steed, D. A., & Jackman, S. (2013). Concreteness and word production. *Memory & Cognition*, 41, 365–377.
- Hollis, G., & Westbury, C. (2016). The principals of meaning: Extracting semantic dimensions from co-occurrence models of semantics. *Psychonomic Bulletin & Review*, 23, 1-13.
- Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science*, 336, 1049–1054.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: The University of Chicago Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*.

Chicago: University of Chicago press.

Murphy, G. L. (1997). Reasons to doubt the present evidence for metaphoric representation.

Cognition, 62, 99-108.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks:

Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867-872.

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1964). *The measurement of meaning*. Urbana: University of Illinois Press.

Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Science*, 108, 3526–3529.

Pustejovsky, J. (1995) *The Generative Lexicon*. Cambridge: MIT Press.

Reddy, M. J. (1979). The conduit metaphor: A case of frame conflict in our language about language. *Metaphor and Thought*, 2, 164–201.

Regier, T., Kemp, C., Kay, P. (2015). Word meanings across languages support efficient communication. In *The Handbook of Language Emergence* (pp. 237-263). Hoboken: Wiley-Blackwell.

Silverstein, M. (1976). Hierarchy of features and ergativity. In R. Dixon (Ed.), *Grammatical categories in Australian languages*.

Simpson, J., & Weiner, E. S. (1989). *Oxford English Dictionary, Second Edition*. Oxford: Oxford University Press.

Srinivasan, M., & Carey, S. (2010). The long and the short of it: on the nature and origin of functional overlap between representations of space and time. *Cognition*, 116, 217–241.

Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124-152.

Sweetser, E. (1991). *From etymology to pragmatics: Metaphorical and cultural aspects of*

semantic structure. Cambridge: 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, 521-540.

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. (1990). From less to more situated in language: the unidirectionality of semantic change. In *5th International Conference on English Historical Linguistics* (pp. 497-517).

Traugott, E. C. (2003). From subjectification to intersubjectification. *Motives for Language Change*, 124–139.

Ullmann, S. (1957). *The principle of semantics*. Glasgow: Jackson, Son & Co.

von Humboldt, W. (1836). *Über die Verschiedenheit des menschlichen Sprachbaues und ihren Einfluss auf die geistige Entwicklung des Menschengeschlechts*. F. Dummler, Berlin.

Winter, B., Thompson, G., & Urban, M. (2014). Cognitive factors motivating the evolution of word meanings: Evidence from corpora, behavioral data and encyclopedic network structure. In *Evolution of Language: Proceedings of the 10th International Conference*, 353-360.

Xu, Y., Regier, T., Malt, B. C. (2015). Historical semantic chaining and efficient communication: The case of container names. *Cognitive Science*. doi: 10.1111/cogs.12312.

Zipf, G. (1949). *Human Behavior and the Principle of Least Effort*. New York: Addison-Wesley.