

Words with more diverse semantic networks are more readily extended to novel meanings

Qiawen Liu (qliu295@wisc.edu)^{1,2}, Gary Lupyan (lupyan@wisc.edu)^{1*},
Mahesh Srinivasan (srinivasan@berkeley.edu)^{3*}, Bill D. Thompson (wdt@berkeley.edu)^{3*}

¹Department of Psychology, University of Wisconsin-Madison

²Department of Computer Science, Princeton University

³Department of Psychology, University of California, Berkeley

* These authors share senior authorship (listed alphabetically)

Abstract

Why do we say “grasp an idea” but not “hold an idea,” or “small talk” but not “little talk”? Although these word pairs share similar meanings, they differ in their tendency to be metaphorically extended. Prior research has focused on how well source concepts map onto target meanings to explain why some words are more readily extended to novel metaphorical contexts. However, much less attention has been paid to how the structural properties of a word’s semantic network contribute to its patterns of metaphorical extension. Using a large-scale dataset of semantic networks in English (De Deyne et al., 2019), we explored whether these structural properties predict selective extension in a metaphor rating study with English speakers, who were presented with Mandarin metaphorical expressions that lack direct English equivalents. Our findings revealed that English speakers systematically favor certain conceptual mappings from Chinese, suggesting that some metaphorical mappings may be universally recognized. However, when controlling for conceptual mappings by analyzing synonym pairs, we found that the structural features of a word’s semantic network significantly predict rates of extension. Specifically, words embedded in densely interconnected neighborhoods showed stronger resistance to extension, while words bridging diverse semantic communities demonstrate greater mutability.

Keywords: metaphor; language evolution; semantic extension;

Introduction

Two centuries before radio waves crackled through the air, English farmers walked their fields with bags of grain, scattering seeds in wide arcs. They called this act “broadcasting,” a term that would later be recycled in the age of telecommunications, to metaphorically describe signals dispersed through the airwaves to countless receivers. Many words carry a hidden history of metaphorical semantic extension—transformations where language stretches to capture the shape of new ideas. Today, metaphors are woven into our thoughts and conversations. We “grasp” concepts, “wrestle” with ideas, and “construct” arguments without realizing these are metaphors, remnants of meanings left in our lexicon, like a reef where each branch once grew and calcified before becoming part of a larger whole.

Yet not all words undergo such transformation. While we readily “grasp” ideas, we rarely “hold” or “grip” them. This selectivity appears across languages, revealing both universal tendencies and culturally specific variations. Many languages describe personality in terms of temperature (e.g., “warm personality”), yet Japanese evokes roundness (marui

seikaku, “round personality”). Similarly, while languages like English and Dutch describe pitch using vertical space (e.g., high/low), others such as Farsi, Turkish, and Zapotec instead use metaphors of thickness (e.g., thin/thick) (Shayan et al., 2011). Even within the same culture, semantically similar words can diverge: “small talk” is common whereas “little talk” is not; “big name” is natural while “large name” is less so. These patterns of selectivity raise an intriguing question: what determines which words will successfully make these metaphorical leaps?

Two complementary forces may shape the development of metaphorical senses. The first force is the conceptual mapping between the source concept (the original meaning of the word) and the target concept (the new, extended meaning), which is influenced by both communicative and cognitive pressures. Communicative pressures require that metaphorical extensions be readily understood in interaction, drawing on mutual knowledge of the source domain among speakers (E. V. Clark & Clark, 1979) and on salient, readily accessible features that create common ground for comprehension (H. H. Clark & Gerrig, 1983). As a result, words tied to external, concrete, bodily experiences frequently extend to internal, abstract, disembodied concepts (Srinivasan & Carey, 2010; Xu et al., 2017), with extensions usually revolving around salient features of source domain concepts (H. H. Clark & Gerrig, 1983; Zhu & Malt, 2014). Cognitive pressures further constrain these mappings by favoring extensions that reduce mental effort. For example, extensions that are conceptually related to existing meanings are more readily comprehended, learned, and recalled (Karjus et al., 2021; Rodd et al., 2012; Thibodeau & Durgin, 2008), with relatedness measurable through associative networks (De Deyne et al., 2019), word embeddings (Mikolov et al., 2013), or alignment on semantic dimensions such as valence and arousal (Osgood, 1957; Osgood et al., 1975). Consequently, senses develop via incremental links to closely related existing ones (Lakoff, 2008; Malt et al., 1999), a pattern corroborated by cross-linguistic colexification patterns (Xu et al., 2020) and computational analyses of historical semantic change (Hamilton et al., 2016; Ramiro et al., 2018).

However, some words are more readily extended than others, even when both could conceptually map to the same target concept. For example, why do we speak of “small talk” but not “little talk”? The second, underexplored force in-

fluencing metaphorical extension may lie in a word’s structural position within an association-based semantic network. A word’s network position can influence its adaptability and likelihood of extension. Specifically, words with more diverse associations may be more flexible in meaning and thus more likely to extend metaphorically. In contrast, words that are tightly tied to specific contexts or communities may be less adaptable.

To better understand the roles of conceptual mapping and words’ structural position in metaphorical extension, we designed an experiment using Mandarin metaphorical expressions that do not exist in English. By presenting English speakers with these novel expressions, we can test whether certain metaphorical extensions are universally constrained/favored by forces independent of language or culture (e.g., whether naive English speakers prefer “small/little heart” as in Mandarin over “big/large heart” to describe being careful). Next, by comparing how English speakers rate different synonyms that could equally map to the same meaning (e.g., “small” vs. “little”, “big” vs. “large”), we sought to isolate the effect of a word’s structural position in the semantic network while controlling for conceptual mapping. Ultimately, our experiment examines both universal conceptual constraints on metaphorical extension and the language-specific structural factors that determine which words, among conceptually suitable candidates, are most likely to develop new metaphorical senses.

Methods

Participants. A total of $N = 110$ undergraduate students from UW–Madison and UC Berkeley participated in this study for course credit. Eleven participants reporting Mandarin language experience were excluded from analysis. Responses faster than two standard deviations below the mean were excluded.

Materials & Stimuli We identified 36 Mandarin metaphorical expressions that lack conventional English equivalents. To mask the Mandarin origin and prevent potential cultural biases, we presented these expressions as coming from “Daxy,” a recently discovered language. For each expression, participants received the intended metaphorical meaning in English, a context sentence with the metaphorical term marked as “X”, and four candidate English words. They were asked to rate how natural it is to fit each of these words in the specific context of the sentence.

For example, Mandarin uses “small/little heart” (xiǎoxīn) to describe careful, attentive action. Participants saw the definition “In Daxy, X heart refers to being cautious and careful” with the context “He held the antique vase with an X heart, afraid of breaking it.” They rated four words: two synonyms translating the Mandarin term (small, little) and their antonyms (big, large). This design allows us to examine effects at two levels. At a broad level, comparing translations with their antonyms reveals whether English speakers share the conceptual mappings that are lexicalized in Mandarin but




Figure 1: Correlations between independent variables in the English-Mandarin extension study. The diagonal displays variable distributions, while the lower triangles show Pearson correlations.

not English (e.g., associating smallness rather than bigness with carefulness). We expect these polarity effects to be large and robust. At a finer level, comparing synonyms within each polarity (small vs little, big vs large) lets us examine whether and how network properties influence word choice when conceptual mappings are held constant, because e.g. *small* is embedded in a different semantic network than is *little*.

In addition, synonyms often carry subtle emotional differences that could influence their metaphorical extension. For instance, in English, “little” tends to convey more warmth or endearment than “small” (compare “my little friend” vs. “my small friend”). To ensure our findings about metaphorical extension are not biased by focusing on only positive or negative contexts, we tested each synonym pair in contexts with contrasting emotional implications. For example, the Mandarin expression xiǎoxīn (“small/little heart”) describes careful, attentive action—a positive trait. In contrast, xiǎo cōng míng (“small/little wisdom”) describes foolishness or shortsightedness—a negative trait. Table 1 below for additional example stimuli.

Variables & prediction Our analysis examined seven variables that may influence metaphorical extension, falling into two broad categories: (1) measures of semantic and conceptual properties and (2) measures of words’ network properties (Figure 1 showed correlations between all variables).

First, we examined five measures that capture why certain conceptual mappings are preferred, according to previous literature: semantic relatedness, alignment on core affective dimensions: valence and arousal, concreteness and frequency. We quantified each word’s semantic relatedness using BERT-based language models (Devlin, 2018), computing the log probability of each candidate appearing in the context sentence. This measure controls for existing English expressions that might capture similar meanings. For example, while a

Definition (“In Daxy,...”)	Context	Metaphor Source (Chinese)	Antonyms	Context Valence
X can describe a person who is not concerned with wealth, fame, or material gains.	He has a X view on wealth and fame, and always focusing on his inner peace.	bland/ mild	flavorful/ spicy	positive
X can describe a distant relationship.	As time passed, their emotional connection became X, and they were no longer as close as before.	bland/ mild	flavorful/ spicy	negative
X refers to having a strong, rough quality that is natural and not too refined or polished.	The style of this sculpture is X and powerful, showcasing the artist’s bold, unrefined aesthetic.	thick/ wide	thin/ slim	positive
X describes a tendency towards carelessness or neglect.	He is always very X-hearted in his work, often making small mistakes.	thick/ wide	thin/ slim	negative
X can refer to knowledge that is obscure.	I’ve recently become obsessed with a podcast that shares X knowledge that most people don’t know, and I learn something new every day.	cold/ chilly	heated/ hot	positive
X describes a place or environment that is quiet or lacking in liveliness and activity.	Despite its prime location, the new restaurant remained X most evenings, with only a handful of diners scattered across its many empty tables.	cold/ chilly	heated/ hot	negative

Table 1: Example stimuli used in the experiment.

forgiving attitude is not typically described as “broad” in English, the existence of expressions like “open-minded” might influence preferences for “broad/wide” over “narrow/thin.” We predicted that words with higher BERT-predicted semantic probability would receive higher naturalness ratings. Then, we measured alignment between words and contexts on two basic semantic dimensions: valence and arousal (following the procedure described in Warriner et al. (2013)). This tests whether English speakers prefer words that match the intended meaning’s emotional and psychological qualities (e.g., preferring low-potency words like “small/little” for expressing carefulness). We predicted that words better aligned with contexts’ affective dimensions would be judged as more natural. Furthermore, previous work also suggests concreteness and frequency play a role in semantic extension. (Lakoff & Johnson, 1980; Winter & Srinivasan, 2022) Therefore, we obtained concreteness ratings from Brysbaert et al. (2014) and word frequency from the wordfreq Python library (Speer, 2022). While conceptual metaphor theory typically predicts that more concrete terms are more likely to be extended compared to abstract ones (e.g., journey vs. life), this prediction usually applies to metaphorical mappings across distinct domains. However, when comparing extension tendencies for near-synonyms within a single domain (e.g., small/little), we hypothesized that the relationship might be reversed. Words with greater contextual independence and mutability might be conceived as less concrete due to semantic bleaching, a process where frequent metaphorical use gradually reduces a word’s concreteness (Gentner, 1981; Gentner & France, 1988; King & Gentner, 2022). We also predicted that higher word frequency would facilitate semantic extension, as previous works showed more frequent words are more likely to be sources of metaphorical extension (Harmon & Kapatsinski, 2017; Piantadosi et al., 2012; Winter & Srinivasan, 2022).

To examine whether a word’s semantic network predict semantic extension, we computed two specific network properties for each word using the Small World of Words association network (SWOW) (De Deyne et al., 2019). First,

we calculated *cross-clique connectivity* using the centiserve package (Jalili, 2017) to enumerate every maximal clique in the undirected semantic-association network. For each word w we counted the number of distinct maximal cliques containing w . Because a word can appear in several overlapping cliques, the same pairwise edge may be “double-counted.” We retain this redundancy intentionally: a word that participates in two partially overlapping cliques (A–B–C and A–B–D, for instance) is genuinely acting as a bridge between two otherwise separate fully connected groups, and the duplicate counting reflects that bridging role. Second, we computed the *clustering coefficient*, which measures the degree to which a word’s immediate semantic neighbors are interconnected. This measure captures how tightly a word is “locked-in” to a specific semantic neighborhood, representing a system of associations where the word’s neighbors are themselves strongly associated with one another.

These network properties allow us to assess two specific predictions about semantic extension. First, words that bridge different semantic communities should show greater extensibility. Second, words with less densely clustered neighborhoods should extend more easily than those “locked-in” to dense semantic clusters.

Results

English speakers predict Mandarin metaphors

We first coded for each context whether each of the candidate words represented the translation for the original Mandarin expression or not (e.g., for “X heart” meaning careful, small/little = 1, big/large = 0), allowing us to examine whether English speakers are more likely to endorse extensions that align with the cross-linguistic metaphorical pattern.

We fit a linear mixed-effects model with lme4 in R (Bates, 2010), predicting 1-to-7 naturalness ratings from a binary fixed effect—whether the expression is attested in Mandarin. The model included random intercepts for participants and for questions. Expressions that exist in Mandarin were judged more natural ($M = 3.53$) than their unattested antonyms ($M =$

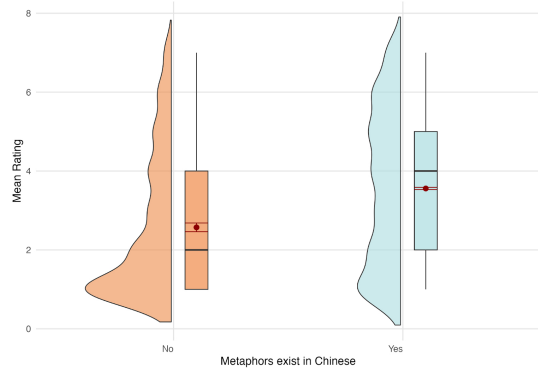


Figure 2: English speakers' naturalness ratings for expressions that do not exist versus exist in Chinese.

2.52). This difference was significant ($b = 1.01$, 95% CI = [0.95, 1.07], $p < .001$; see Figure 2). The effect size (Cohen's $d = 0.61$) is considered medium-to-large and comparable to meaningful contrasts in studies using acceptability judgments (Marty et al., 2020). This suggests that native English speakers, despite having no exposure to Chinese, have intuitions that match conventional expressions in Chinese.

To better understand the factors driving these shared intuitions, we used a generalized linear mixed-effects model to predict which expressions exist in Chinese. The expression existence was significantly predicted by 1) the word-context alignment on affective dimensions, 2) BERT log probability, 3) clustering coefficients, and 4) cross-clique connectivity. These fixed effects explained 33% of the variance in whether an expression exists in Chinese. When including random effects to account for variation across individual items and participants, the model's total explained variance increased to 90%, suggesting substantial systematic variation at both the expression and participant levels.¹

What factors predict metaphorical extension?

To analyze how different factors predict word naturalness ratings, we first established a base model including BERT probability and affective norms (valence and arousal) as predictors. Due to moderate to strong correlations ($r > 0.3$) among network properties, concreteness, and word frequency, we examined these predictors *separately* to avoid multicollinearity. Specifically, we created four additional models by individually adding concreteness, word frequency, cross-clique

¹At the item level, certain questions tend to received higher or lower ratings across all adjective choices. For instance, "x can describe a sneaky action" was judged fairly natural under every adjective frame—large, big, small, and little (all $M > 3.3$; overall $M = 4.1$). By contrast, "x can describe an insightful person" was judged unnatural in three of its four frames—short ($M = 1.53$), low ($M = 1.58$), tall ($M = 2.41$), and high ($M = 3.59$)—yielding an overall mean of 2.3. This item-level variance may reflect unmeasured qualities such as phonological balance (how the phrase sounds) and imageability (how easily the phrase evokes a vivid mental image). At the participant level, we observed individual differences in rating tendencies, with the most lenient rater averaging 3 points above the mean and the strictest 1.7 points below, indicating personal variation in tolerance for novel metaphorical language.

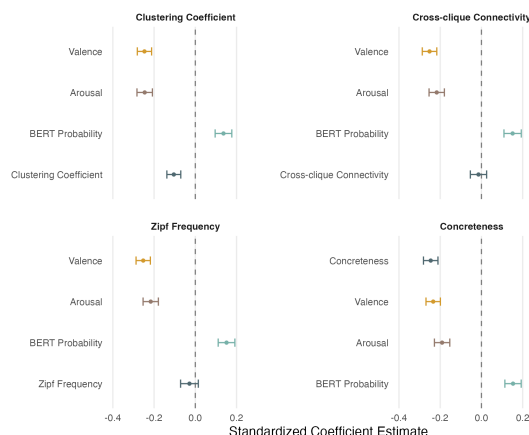


Figure 3: Effects of independent variables in modeling naturalness of extension, considering all four options in a context. Four separate models due to moderate to strong significant correlations among them. Within each panel, variables are sorted based on absolute magnitude of effect.

connectivity, or clustering coefficient to the base model. All models included random intercepts for participants and questions.²

Key predictors of naturalness Figure 3 shows coefficient estimates and confidence intervals for all four models. Semantic relatedness (as measured by log probability of a candidate word within the sentence using the BERT model) was a significant predictor across all models: words with higher semantic probability in a given context were rated as more natural extensions, regardless of which network property was included in the model. Affective alignment between words and contexts also predicted naturalness ratings. Smaller differences in both valence ratings and arousal ratings predicted higher naturalness ratings.

Abstract words were extended more readily Contrary to what conceptual metaphor theory might predict, concreteness emerged as a negative predictor of naturalness ratings ($\beta = -.25$, 95%CI = [-.28, -.21], $p < .001$). More concrete words were less likely to be extended to new contexts. This finding aligns with the mutability hypothesis (Gentner & France, 1988; King & Gentner, 2022), which predicted a positive correlation between a word's semantic flexibility and its level of abstractness - words that appear in diverse contexts tend to be more abstract. This correlation could operate in either direction: abstract words may be more adaptable to new contexts, or words may be viewed as more abstract through diverse us-

²To account for the ordinal nature of our Likert scale data, we conducted parallel analyses using ordinal mixed-effects models for all analyses reported in this paper. In all cases, the key predictors remained significant with similar effect sizes to those reported in the main text, including affective alignment, BERT probability, concreteness, clustering coefficient, and cross-clique connectivity where applicable.

age. Our current data cannot distinguish between these possibilities.

Local clustering coefficient predicted naturalness judgments Our analysis of network properties revealed that local neighborhood structure has significant influence on semantic extension. Words embedded in tight semantic clusters showed strong resistance to being extended to new contexts, as indicated by the significant negative effects of clustering coefficient ($\beta=-0.1$, 95%CI = [-0.14, -0.07], $p<0.001$). This suggests that dense local connectivity constrains semantic flexibility. Cross-clique connectivity ($\beta=-0.01$, 95%CI=[-0.05, 0.03], $p=.5$) and Zipf frequency ($\beta=-0.02$, 95%CI=[-0.07, 0.01], $p=.2$) were not significant predictors of extensibility.

Taken together, these results reveal a hierarchy of factors influencing metaphorical extension. When analyzing naturalness ratings across all candidate words, measures of conceptual mapping - such as affective alignment and BERT-based semantic relatedness - emerged as the strongest predictors. In contrast, the effects of network properties (like clustering coefficient) are smaller and, in some cases (like crossclique connectivity), not significant.

What predicted differences within pairs of synonyms?

To provide further insight into the role of network structure in semantic extension while controlling for strength of conceptual mapping, we conducted a focused analysis on synonym pairs. Our analysis sought to examine whether participants exhibited systematic preferences within pairs of near-synonymous words in the same context (e.g., small/little, big/large). For each pair, we calculated differential scores for each variable by subtracting the second word’s value from the first word’s value (e.g., small minus little). Specifically, we subtracted values for affective norms, BERT probabilities, Zipf frequencies, concreteness ratings, cross-clique connectivity scores, and clustering coefficients. We used these differentials to predict differences in naturalness ratings between synonyms. Again, due to collinearity among concreteness, frequency, and network properties, we created *separate* models testing each of these predictors while including all other variables. All models included random intercepts for participants and questions.

Differences in network structure predict differences in extension between synonyms. Our analysis revealed that network properties were the strongest predictors of differences in naturalness ratings between synonym pairs (Figure 4). The *clustering coefficient* measure of local semantic density was the strongest predictor ($\beta = -0.25$, 95% CI = [-0.33, -0.17], $p < .001$), suggesting that words deeply embedded within specific semantic networks showed reduced flexibility in adapting to new contexts. This finding aligns with our analyses above which showed that local density is a significant predictor of naturalness ratings among all words in our experiment.

To see why local semantic density can be an important

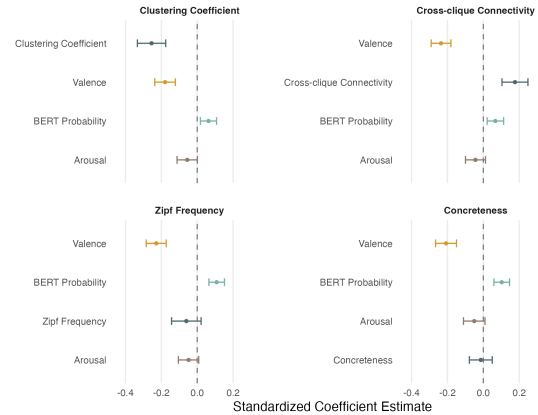


Figure 4: Effects of independent variables in modeling naturalness of extension, considering difference within pairs of synonyms. Variables are sorted based on absolute magnitude of effect.

predictor, consider the contrast between “thick” and “wide” when describing carelessness or neglect, as in “He is always very X hearted in his work.” Participants rated “thick” as more appropriate than “wide,” even though “wide” has higher word frequency and better semantic alignment with the intended meaning (indicated by higher BERT probability scores and more aligned valence ratings). This surprising preference for “thick” exemplifies how network structure can override other predictors of selectivity. As shown in an illustrative diagram (Figure 5), while “wide” exists in a densely interconnected neighborhood where all its neighbors (broad, large, vast, space) are mutually connected, forming multiple triangles that lock it into spatial meanings. In contrast, “thick” connects to diverse, barely interlinked neighbors spanning multiple domains (strong, stupid, fat, solid), with minimal triangulation between them. This less constrained network position gives “thick” greater freedom to extend to novel contexts, allowing it to overcome its disadvantages in frequency and semantic fit.

In a separate model, we also found that *cross-clique connectivity* was a significant positive predictor of differential naturalness ratings between synonyms ($\beta = 0.18$, 95% CI = [0.1, 0.24], $p < .001$). This contrasts with our previous analysis, which failed to detect any effect of this aspect of network structure. Among pairs of synonyms, words that connect more diverse semantic neighborhoods than their synonym showed greater potential for extension.

Valence differentials (but not arousal or concreteness) predict differences between synonyms. Among affective dimensions, valence was a consistent predictor of naturalness differentials, suggesting that synonyms whose valence better aligns with context are more likely to be extended. Arousal showed minimal differentiation between synonyms, likely because word pairs like small/little and big/large are similar in arousal ratings. Concreteness ($\beta = -.01$, 95% CI = [-.07, .05],




Figure 5: Schematic diagram comparing low-clustered (thick) versus high-clustered (wide) network structures

$p = .6$) and Zipf frequency ($\beta = -.06$, 95% CI = $[-.14, .02]$, $p = .15$) did not predict naturalness ratings, likely for similar reasons. Finally, semantic proximity as captured by log probability of the English word in context under a BERT-based model positively predict differentials in semantic extension: words that were more semantically similar to the contextual meaning showed greater potential for extension.

To summarize, when comparing words with similar conceptual mappings, network constraints—especially clustering coefficient—emerge as the most influential predictors, outweighing semantic relatedness and alignment on core affective dimensions in predicting which words are more likely to be metaphorically extended.

Discussion & Conclusion

Traditional accounts of metaphorical extension have emphasized features like semantic attributes, usage frequency, and concreteness as key factors that determine how words acquire additional meanings. Our findings support these insights, but also underscore a crucial, often overlooked element: the structure of a word’s existing network of semantic associations. Specifically, we found that the structural diversity of a word’s semantic neighbors significantly predicted its potential for metaphorical extension, even when controlling for other conceptual and semantic properties.

By asking English speakers to evaluate novel metaphorical expressions adapted from Chinese, we found that factors like valence, arousal, and BERT-based semantic probability strongly predict metaphorical extension when words differed significantly in how well their meanings aligned with the target concept. However, when comparing near-synonymous words, such as “small” and “little,” which shared conceptual mappings, network structure was a key factor in determining whether a novel metaphorical meaning would be endorsed. Specifically, words embedded in tightly connected, homogeneous semantic networks were less likely to be extended metaphorically, while words that bridge diverse semantic communities were more likely to be extended metaphorically.

This dual mechanism emphasizes the complementary roles of conceptual mapping and network structure in semantic extension. A parallel analogy from biological evolution suggests two complementary pathways through which traits—and, by analogy, words—can be co-opted for new functions. The first pathway focuses on features that are preadapted for new functions, independent of an entity’s re-

lationships with others. For example, swim bladders evolved into lungs because the tissue in swim bladders was already well-suited for gas exchange. Similarly, some words acquire new meanings because they are semantically close to or have conceptual properties align well with emerging communicative needs. The second pathway centers on structural relationships, where an entity’s position within a larger network influences its adaptability. In biological systems, elements that bridge different functional modules often emerge as key drivers of new functions (Koentges, 2008). Likewise, words that bridge diverse semantic domains—connecting disparate concepts—are more likely to adopt novel metaphorical meanings. By contrast, words embedded in tightly interconnected semantic clusters are less flexible and resistant to change. This is akin to “evolutionary lock-in,” a concept where tightly integrated systems resist change as modifications would disrupt essential functions (Wagner, 2011). Just as evolutionary lock-in prevents certain genes from evolving without disrupting broader systems, words in tightly-knit semantic neighborhoods resist changes in meaning, as shifts could incur processing costs or reduce communicative informativity.

Our findings also refine and challenged existing theories of semantic change. *Less* concrete words were more likely to be extended into a novel context. This finding is consistent with a “mutability through semantic bleaching” account: words that have developed more flexible usage patterns over time tend to be more abstract. However, this pattern disappeared when looking at synonyms, indicating additional complexities that remain unexplained. Additionally, word frequency did not significantly predict extension, suggesting that being a more familiar term alone cannot account for the observed differences between semantically close words. Meanwhile, factors such as affective alignment with context (particularly valence) remained powerful predictors of extension judgments.

Our study has several limitations that suggest directions for future research. First, our analyses focused primarily on concrete nouns and adjectives in metaphorical extension. Further research could investigate whether similar network effects operate in other word classes and other types of semantic change, such as grammaticalization or pejoration. Second, stronger causal tests could involve longitudinal studies tracking how network position predicts actual semantic change over time. Such studies could examine whether words’ changing network positions precede or follow semantic innovation, helping clarify the causal relationship between network structure and semantic extension.

Our results emphasize that theories of linguistic innovation must integrate both traditionally emphasized conceptual properties of the word itself (e.g. semantic attributes, frequency, concreteness) and the structural properties of a word’s existing network of relations. This perspective may help to explain puzzling discrepancies between the evolutionary trajectories of closely related words, shedding new light on the interactions between functional pressures and structural dependencies in semantic change.

References

- Bates, D. M. (2010). Lme4: Mixed-effects modeling with R. *Behavior research methods*, 42, 137–145.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior research methods*, 46, 904–911.
- Clark, E. V., & Clark, H. H. (1979). When nouns surface as verbs. *Language*, 55, 767–811.
- Clark, H. H., & Gerrig, R. J. (1983). Understanding old words with new meanings. *Journal of verbal learning and verbal behavior*, 22(5), 591–608.
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “small world of words” English word association norms for over 12,000 cue words. *Behavior research methods*, 51, 987–1006.
- Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Gentner, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, 13(1), 56–83.
- Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In *Lexical ambiguity resolution* (pp. 343–382). Elsevier.
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016). Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change. *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing, 2016*, 2116–2121.
- Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, 98, 22–44.
- Jalili, M. (2017). *Centiserve: Find graph centrality indices* [R package version 1.0.0].
- Karjus, A., Blythe, R. A., Kirby, S., Wang, T., & Smith, K. (2021). Conceptual similarity and communicative need shape colexification: An experimental study. *Cognitive Science*, 45(9), e13035.
- King, D., & Gentner, D. (2022). Verb Metaphoric Extension Under Semantic Strain. *Cognitive Science*, 46(5), e13141.
- Koentges, G. (2008). Evolution of anatomy and gene control. *Nature*, 451(7179), 658–663.
- Lakoff, G. (2008). *Women, fire, and dangerous things: What categories reveal about the mind*. University of Chicago press.
- Lakoff, G., & Johnson, M. (1980). The metaphorical structure of the human conceptual system. *Cognitive science*, 4(2), 195–208.
- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40(2), 230–262.
- Marty, P., Chemla, E., & Sprouse, J. (2020). The effect of three basic task features on the sensitivity of acceptability judgment tasks. *Glossa: a journal of general linguistics (2021-...)*, 5(1), 72.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Osgood, C. E. (1957). *The measurement of meaning*. Urbana: University of Illinois Press.
- Osgood, C. E., May, W. H., & Miron, M. S. (1975). *Cross-cultural universals of affective meaning* (Vol. 1). University of Illinois Press.
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3), 280–291.
- Ramiro, C., Srinivasan, M., Malt, B. C., & Xu, Y. (2018). Algorithms in the historical emergence of word senses. *Proceedings of the National Academy of Sciences*, 115(10), 2323–2328.
- Rodd, J. M., Berriman, R., Landau, M., Lee, T., Ho, C., Gaskell, M. G., & Davis, M. H. (2012). Learning new meanings for old words: Effects of semantic relatedness. *Memory & Cognition*, 40, 1095–1108.
- Shayan, S., Ozturk, O., & Sicoli, M. A. (2011). The Thickness of Pitch: Crossmodal Metaphors in Farsi, Turkish, and Zapotec. *The Senses and Society*, 6(1), 96–105.
- Speer, R. (2022). *Rspeaker/wordfreq: V3. 0. Version v3. 0.2. Sept*.
- 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(2), 217–241.
- Thibodeau, P., & Durgin, F. H. (2008). Productive figurative communication: Conventional metaphors facilitate the comprehension of related novel metaphors. *Journal of memory and language*, 58(2), 521–540.
- Wagner, A. (2011). *The origins of evolutionary innovations: A theory of transformative change in living systems*. OUP Oxford.
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior research methods*, 45, 1191–1207.
- Winter, B., & Srinivasan, M. (2022). Why is Semantic Change Asymmetric? The Role of Concreteness and Word Frequency and Metaphor and Metonymy. *Metaphor and Symbol*, 37(1), 39–54.
- Xu, Y., Duong, K., Malt, B. C., Jiang, S., & Srinivasan, M. (2020). Conceptual relations predict colexification across languages. *Cognition*, 201, 104280.
- Xu, Y., Malt, B. C., & Srinivasan, M. (2017). Evolution of word meanings through metaphorical mapping: Systematicity over the past millennium. *Cognitive psychology*, 96, 41–53.
- Zhu, H., & Malt, B. C. (2014). Cross-linguistic evidence for cognitive foundations of polysemy. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 36(36).