Semantics in Action: Nameability Predicts Verb Extension

Eleanor Flanigan
eflanigan@wisc.edu
Psychology, University of Wisconsin-Madison

Aja Altenhof
altenhof@wisc.edu
Linguistics, University of Wisconsin-Madison

Gary Lupyan
lupyan@wisc.edu
Psychology, University of Wisconsin-Madison

Abstract

What determines whether a novel referent is judged as similar enough to extend an existing term, or different enough to coin a new one? We investigated how nameability and event structure influence verb extension. Participants learned novel words for unlexicalized actions varying in nameability, then saw variants differing in end state, process, or object and chose whether to extend the learned word or select an alternate. Nameability was the strongest predictor: more nameable actions saw more extension. For event structure, results replicated prior findings—end-state changes resisted extension most, followed by process changes, with object changes permitting the most extension. Extension was the dominant response overall, consistent with efficiency pressures favoring reuse of existing vocabulary. These findings suggest that lexical decisions reflect a tradeoff between the accessibility of existing representations and the degree of semantic difference introduced by variation.

Keywords: psychology, linguistics, language acquisition, language comprehension, representation

Introduction

When people learn a language, they acquire a set of words that can express a wide range of meanings. Yet many meanings lack conventional labels (Lehrer, 1974; Lyons, 1977), requiring speakers to produce a descriptive phrase ('the awkward dance when two people keep stepping in the same direction'), extend the meaning of an existing word (*surf* the web), or coin a new term (*googling*). What determines which strategy a speaker chooses?

Each approach has costs. Descriptive reference is flexible, but inefficient: speakers communicate more effectively when their language lexicalizes a target category than when they must compose a description (Rissman et al., 2023). Coinage avoids this inefficiency but may risk miscommunication: a novel word could be clear to the speaker, but opaque to listeners who haven't encountered it. Faced with competing pressures for efficiency and accuracy (Gibson et al., 2019; Zaslavsky et al., 2018), semantic extension, when a word's original meaning broadens to include new meanings, offers a middle path, exploiting perceived similarity between established and novel referents to leverage familiar vocabulary. Perhaps for this reason, extension is the dominant strategy. Coinage is comparatively rare (Keyes, 2021).

If speakers prefer to assimilate new referents into existing categories, what determines whether a novel referent is judged as similar enough to extend an existing category term, or different enough to require a new one? Relational attributes—how things function, what they cause, and what roles they play—are central to category structure (Gentner, 2005; Markman & Stilwell, 2001; Rehder & Hastie, 2001). Actions, which encode richer relational information than objects (Gentner, 1982), may therefore provide a stronger test of how relational structure influences lexical choice.

The relational attributes of actions can be characterized as varying along three dimensions: the object involved, the end-state achieved, and the process it takes to achieve it (Behrend, 1990; Hovav & Levin, 2010; Talmy, 1985, 2000). Previous literature suggests that both the status of an action's end-state and process of completion may be semantically privileged. Accordingly, linguistic theories of verb classification distinguish between *manner* verbs, which encode process (e.g., wipe, scrub, slap), from *result* verbs, which encode end-state (e.g., clean, break, open) (Hovav & Levin, 2010). Comparatively, object changes are less salient. The same action can be completed with a wide range of objects—one can cut paper, fabric, or hair—reflecting that verbs lexicalize relations between entities rather than the entities themselves (Gentner, 1982; Gleitman, 1990). Given that these elements are encoded in verbs' meaning, the likelihood of extension may vary if they are changed.

Evidence from verb learning studies suggests that this is the case. Behrend (1990) taught children and adults three types of novel events: action verbs (encoding manner), result verbs (encoding end-state), and instrument verbs (encoding the object used). Afterwards, participants were asked to extend these verbs to new events where one component had changed. They were least willing to extend when the result had changed and result bias increased with age. Manner changes had moderate effects; instrument changes had the weakest effect. This pattern suggests that from early in development, the outcome of an event is treated as more salient to a verb's meaning than the process of achieving it.

Beyond the structure of the event itself, extension may depend on how well the original action is represented in speakers' existing lexical knowledge. In production, extension is driven primarily by the accessibility of lexical forms: more frequent, more accessible forms are extended to novel meanings, even when less accessible alternatives might be a better semantic fit (Harmon & Kapatsinski, 2017). However, when speakers are prompted to describe novel referents that frequent forms don't efficiently capture, less frequent or novel forms may be used to denote the unlexicalized meaning.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Some unlexicalized actions are easy to describe. For example, *sealing an envelope* evokes a clear, convergent image. Other actions, like *making a hole in a pile* are more challenging. We index this difference as *nameability*, or how readily speakers converge on a description for a referent. Some actions have more natural, consistent boundaries than others (Zacks & Tversky, 2001), and verbs present a unique challenge to language acquisition because the meanings they encode are more variable, context-dependent, and less consistently bounded than object concepts (Gentner, 1982; Gillette et al., 1999). Description convergence for a given event reflects the degree to which the underlying conceptual structure is consistent and stable across language users.

This conceptual coherence has consequences for learning and use. More nameable concepts are easier to learn (Zettersten & Lupyan, 2020), and the consistency of an action's structure influences how readily a novel verb is extended to new instances (Kersten & Smith, 2002). Thus, highly nameable actions may have more accessible representations, prompting higher rates of extension. If an action is difficult to name, no existing term may provide a strong basis for extension. Instead, speakers may reject an extension in favor of a novel label.

The present study asks which features of a novel action influence whether speakers extend an existing verb or reject an extension in favor of a new word. We operationalized two types of variance: nameability, how easily speakers can describe a new action, and dimensions of event structure. Participants learned novel verbs for nine news actions (e.g., joicate means "to seal an envelope"). Each base action was manipulated on three dimensions: end state (what the action achieves), process (how an action is performed), or object (what undergoes the action), creating six potential targets for extension. Our extension targets included both one and two-dimension manipulations. For example, "joicate," where the base action was sealing an envelope, was manipulated on the process dimension so the shown action in the target was tapping an envelope closed, keeping the end-state and object the same. For a two-dimension manipulation, "pulact," where the base action was tearing a piece of paper out of a notebook, was changed by the process dimension (tearing vs. cutting with scissors) and object dimension (notebook vs. textbook), while the end-state was preserved. After learning, participants were shown videos of these variants and asked, in a forced choice task, to either extend the learned word or select a new one.

We predicted that harder-to-name actions will see less extension with weaker representations and no existing term to generalize from. Aligned with earlier findings (Behrend, 1990; Forbes & Farrar, 1993), we expect that participants would be least likely to extend for end-state changes, followed by process changes, and finally object changes.

Method

Participants

58 participants (46 female, 13 male) were recruited from the University of XXX psychology participant pool, and received course credit for their time.

Stimuli

We created videos depicting nine base actions that do not have single word names in English. For example, "sealing an envelope." Each action was paired with a novel word. There were three versions of each base action to provide variability (e.g, a person sealing a white envelope and another person sealing a blue envelope). These base actions were systematically varied along three dimensions, object, process, and end state, to create seven versions (Figure 1) that varied in one or two dimensions (a version where all dimensions varied was included as well, as an attention check).

Single-dimension variations (one feature changed from the original):

- *Different object*: The action was performed with a different patient or object.
- *Different process*: The same end state was achieved with the same object, but the manner of completion differed (e.g., different tool use or movement).
- *Different end state*: The same object and manner were used, but the goal was not achieved.

Two-dimension variations (two features changed from the original):

- *Same object*: Process and end state differed; only the object remained the same.
- *Same process*: Object and end state differed; only the process remained the same.
- *Same end state*: Object and process differed; only the end state remained the same.

Nameability To assess the nameability of our stimuli, we recruited a separate sample of 48 participants from an undergraduate participant pool. These participants viewed all video stimuli from the main experiment—both the base actions and its variations—and asked to "describe what happens in the video" in a free response box. We first computed the semantic similarity of participants' descriptions for each using a pretrained sentence transformer model (*all-MiniLM-L6-v2*). For each video, we took the pairwise cosine similarities across all participant descriptions and averaged these to produce a nameability score. Conventionality scores were computed separately for each base action and variations. Higher scores indicate greater agreement among participants (Figure 2).

Traimin	Same Object		Different Object	
	Different End State	Same End State	Different End State	Same End State
Same Process	*Different State*	*Base*	*Same Process*	*Different Object*
	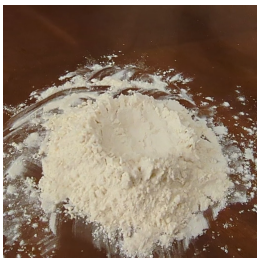			
	Hands are used to make and indent in the pile of flour. No hole is made.	Hands are used to make a hole in a pile of flour.	Hands are used to spread a pile of cards. No hole is made.	Hands are used to make a hole in a pile of cards.
Different Process	*Same Object*	*Different Process*	*All Different*	*Same State*
				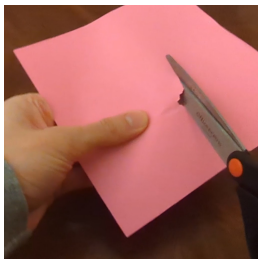
	A spoon is used to make an indent in a pile of flour. No hole is made.	A spoon is used to make a hole in a pile of flour.	A piece of paper is folded in half. No hole is made.	Scissors are used to cut a hole in a piece of paper.

Figure 1: Stimulus set for one of the base actions, "to make a hole in a pile of something" (*traimin*). For this base action, the process uses someone's hands, the object is flour, and the end state is a hole in the object. End-state variations did not result in hole. Process variations used different tools to make the hole (e.g., scissors and a spoon). Object variations used paper and cards instead of flour.

Procedure

Training Participants completed 141 training trials to learn the names of the different actions (3 videos per action). On each trial, participants saw a video of the action and were given a choice between two novel words. After each trial, participants received accuracy feedback.

Extension Following training, participants completed 62 extension trials. On each trial, participants viewed a video depicting either the base action or one of its variations. Participants then had to choose between two words: the learned word for the base action or a novel word not seen during training. Selecting the learned word would indicate that participants extended its meaning to encompass a new variation of the action. Selecting the novel word would indicate that participants viewed the variation as outside the scope of the learned word, effectively coining a new one. All six variations appeared for each of the nine words (54 trials), plus one presentation of base action as an attention check (9 trials).

Results

Participants successfully learned the words ($M = 0.85, SD = 0.09$), reaching 95% accuracy ($SD = 0.28$) by the final 25% of trials. Three participants were excluded for failure to learn (mean accuracy $< 60\%$). On test trials where the action was unchanged from training, participants used the learned word 97% ($SD = 0.17$) of the time, further confirming learning.

We analyzed the likelihood that a (just-learned) word was extended as a function of which dimensions were changed from the action that the word was originally associated with. Removing trials testing the base actions, we fit a generalized linear mixed-effects model, with main effects for the presence of a process change, a state change, an object change and nameability. We included random intercepts for participants and words to account for individual differences in overall extension rates and baseline differences between words. To compare effect sizes, we standardized all predictors (Figure 3).

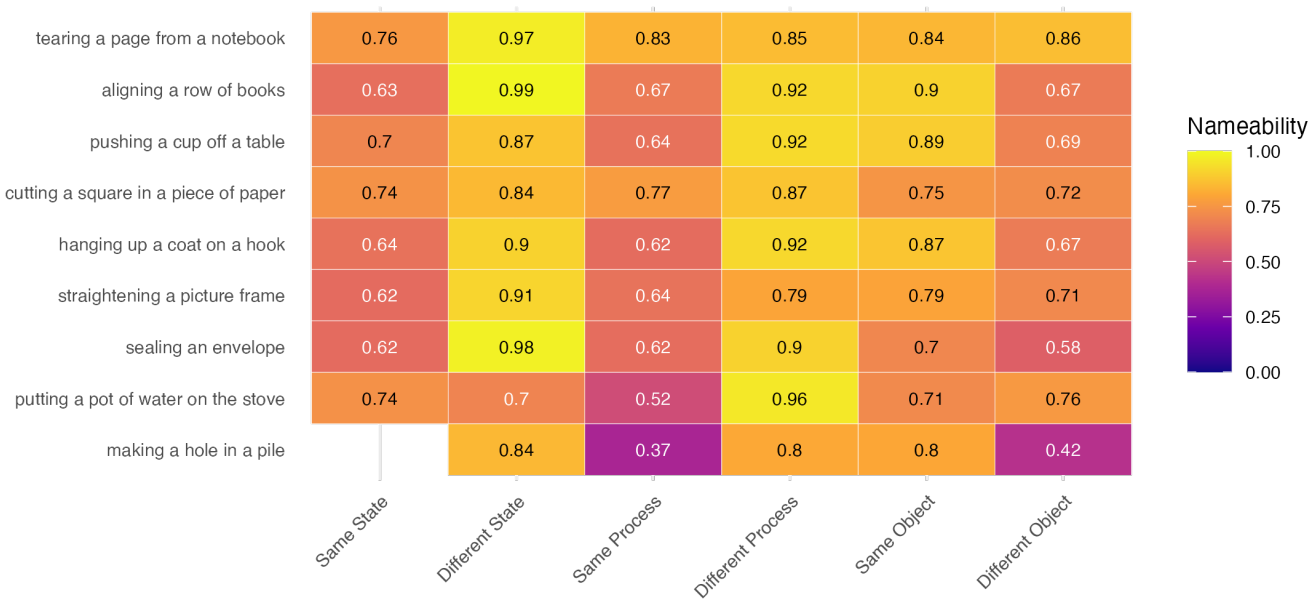

Figure 2: Nameability of base action variations across different dimension changes. Rows represent base actions ordered by mean nameability; columns indicate which feature was manipulated.

The overall rate of extension was high (M = 0.83, SD = 0.37), though it was modulated by the number and type of dimensions changed (Table 1). Though some variability in extension can be attributed individual differences in participant behavior ($SD = 1.45, ICC = 0.36$), there was explained by word ($SD = 0.62, ICC = 0.07$). Future work could address how individual representations of meaning interact with extension.

Unsurprisingly, as more dimensions were changed, participants were less likely to extend: both state and process changes reduced the likelihood of extension. State changes has the largest effect, strongly reducing the likelihood of extension($b = -0.77$, SE = 0.07, $p < 0.001$). This suggests that end-states are most central to a verb's meaning—when the outcome differs, it's less likely to be considered the same action.

Table 1: Extension rates by dimension changes and nameability

Condition	Extension Rate	Nameability
Different Process	92.4%	0.88
Different Object	89.7%	0.67
Same State	82.6%	0.68
Different State	80.7%	0.89
Same Process	71.6%	0.63
Same Object	68.9%	0.81

Extension rate is the proportion of trials where participants used the learned novel word for the variation of the action. Nameability is mean semantic similarity of participant descriptions.

Participants were also less likely to extend process changes ($b = -0.28$, SE = 0.07, $p < 0.001$), though less than state changes. Process may be relevant to an action's meaning, but more tolerant of variation than end-state.

There was no main effect of object changes on extension ($b = 0.08$, SE = 0.11, p = 0.47), consistent with speaker knowledge that verbs regularly take different objects (e.g., "She mailed the letter" and "She mailed the package"). Objects may be more peripheral to verb meaning.

More nameable actions were more likely to be extended ($b = 0.57$, SE = 0.10, $p < 0.001$; Figure 4). More nameable actions may reflect consolidated concepts that extend fluently, but create conflict when extension is rejected.

Discussion

We manipulated action nameability and event structure dimensions to investigate which conditions most influence the likelihood of semantic extension for previously unlexicalized actions. Nameability was a predictor of extension: more nameable actions were more likely to be extended. For structural dimensions, our results were consistent with previous literature (Behrend, 1990; Forbes & Farrar, 1993). Semantic extension was least likely for end-state changes, moderately likely for process changes, and most likely for object changes, suggesting that each dimension does not contribute equally to verb meaning.

The low extension rate for state changes suggests that the result of a given action is often most salient to verb meaning. If "to break" means "to cause something to become broken," then an event without that outcome may not count as the same action (Hovav & Levin, 2010). The primacy of state changes

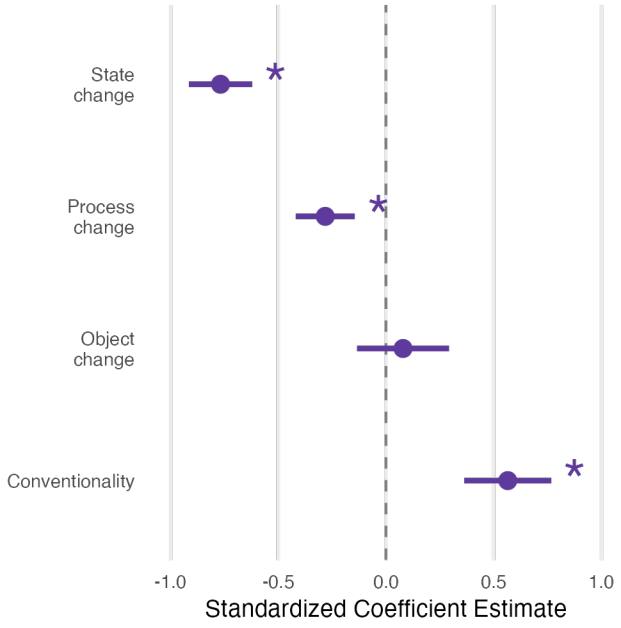

Figure 3: Effects of dimension changes (state, process, and object) on extension likelihood. Points represent standardized coefficient estimates from a generalized linear mixed-effects model with 95% confidence intervals. The dashed line indicates no effect. State changes strongly reduced extensions, followed by process changes. Object changes had no effect. Higher nameability increased extension.

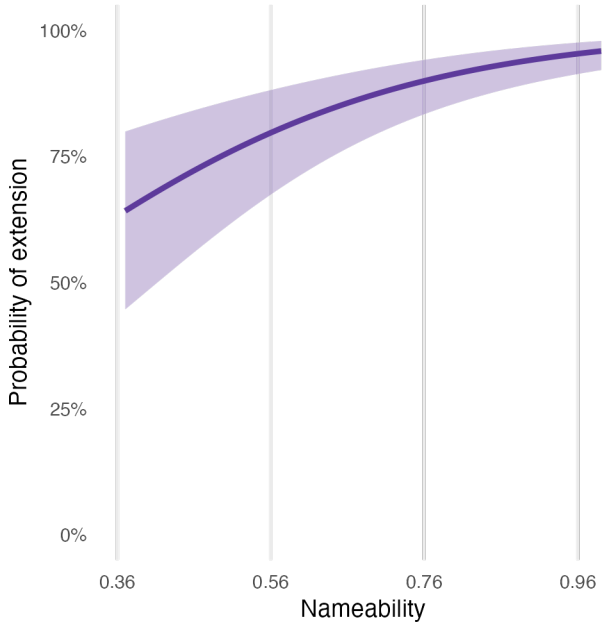

Figure 4: Probability of extension as a function of nameability. Line represents model-predicted probabilities; the shaded region indicates 95% confidence intervals. More nameable actions were more likely to be extended.

could also reflect sensitivity to the agent's goals: actions that do not reach their intended end state are construed as a fundamentally different events (Lakusta & Landau, 2005). This interpretation is consistent with broader evidence that goal structure plays a privileged role in event representation and learning (Zacks & Tversky, 2001).

Process changes showed higher rates of alternate word selection than end-state changes. While actions varied by process accommodate more variability than actions varied by end state, they still lose part of their relevant meaning. Since English prioritizes manner verbs (Talmy, 1985, 2000), changing the process by which the action is complete may feel like changing the verb itself, particularly for the native English speakers in our sample.

Object changes prompted the most extension, confirming that the object in use is often not a defining feature of verb meaning. Verbs are relational, and depend on connections between entities, not entities themselves (Gentner, 1982).

Beyond event structure, more nameable actions were more likely to be extended. This may reflect familiarity—frequently encountered actions yield more flexible representations that accommodate variation (Bybee, 2006). Additionally, more nameable actions likely occupy denser regions of semantic space with more close neighbors. When both an action and its variants are nameable, there is pressure to extend because similar actions already share lexical territory; existing forms preempt coinage when they are close enough to cover a new meaning (Clark, 1990). Selecting a new word for a nameable action may also require overriding an activated existing representation, since producing a word recruits inhibitory mechanisms to suppress co-activated lexical alternatives (Britt et al., 2016), making extension the more efficient choice. In contrast, low-nameability actions have no close semantic neighbors, so the alternate word fills an unlexicalized gap.

Extension was the dominant response overall: participants consistently chose to extend the learned word to the varied action, rather than adopting a new word. If extension were driven purely by novelty or task demands, we would not expect the systematic differences across dimensions that we observed. Alternate word choice was not because participants failed to learn the original label; high learning accuracy indicates they did. Since verbs are typically seen in more varied contexts than our initial video presentations, having more narrow representations of each verb may have made extension less likely. The high rate of extension indicates that the efficiency of using an existing label outweighs the accuracy benefits of using a new one, at least when semantic difference is small. Thus, high extension rates suggest that most variations were considered sufficiently similar to the original verb. Faced with competing pressures for efficiency and accuracy, language users optimize existing form-meaning mappings (Gibson et al., 2019) , reusing words that maximize in-

formativeness while minimizing complexity and vocabulary (Xu et al., 2024).

Several limitations warrant consideration. First, participants learned our novel actions in a non-pragmatic and impoverished context. Natural verb learning occurs in richer environments, and because our verbs were newly lexicalized, they may not have been fully integrated into semantic space (Gaskell & Dumay, 2003). Because our learning task was restricted to video examples, representations may have been narrower than those formed through varied exposure (Bybee, 2006; Goldberg, 2006). However, given that more frequent and variable exposure promotes extension, the prevalence of extension may have been more robust if participants had more familiarity. Regardless, responses still reflected known constraints on event representation, showing consistent patterns of extension.

Second, generating truly unnameable actions is difficult, and there was limited variability in our nameability ratings. Though we purposely chose unlexicalized actions, the absence of an existing label does not preclude an existing representation, and participants may have assimilated our actions into familiar categories. Future work could use genuinely unfamiliar or abstract actions—perhaps from esoteric domains or designed as animations—to more robustly test what patterns hold when prior knowledge is truly minimal.

Finally, semantic extension is ultimately a population level phenomenon (Croft, 2000). Our paradigm tested individual participants making isolated judgments. Iterated learning paradigms, dyadic communication games, or longitudinal corpus analyses would be needed to assess whether the individual-level preferences documented here predict community-level patterns of lexical change.

Our nameability finding suggests that not only the accessibility of the novel verb, but also the nameability of the action that the verb refers to, determines the likelihood of extension. However, the effect of nameability could be explained by greater accesibility of nameable verbs. In forced-choice tasks, accessibility differences between alternatives are reduced, and extension decisions are driven primarily by semantic match (Harmon & Kapatsinski, 2017). This dissociation is important for interpreting nameability effects: if nameability operates purely through accessibility, its effects should be muted in forced choice; if nameability also reflects semantic coherence, it should predict extension even when accessibility is leveled. In our study, the actions that were most nameable could be more retrievable, and see higher rates of extension, simply because they are more accessible. Additionally, the forced-choice design of our experiment may have prompted participants to pick the learned word over the novel option, given that the learned word is familiar and had been the correct option in the learning task. Future work will examine patterns of extension in language production tasks, using less nameable actions, aiming to dissociate nameability from accessibility.

Another direction concerns individual differences in semantic representation. We found that a substantial portion of variance in extension behavior was attributable to participants. Speakers may differ in the richness of their conceptual representations, and how they weight features when judging similarity. Nameability itself may interact with these differences: when an action is highly nameable, individual variation in representation should matter less, because representation is more likely to be convergent. For low-nameability actions, individuals differences in conceptual structure may drive variance in extension patterns. Future work pairing productive extension paradigms with measures of semantic richness and feature weighting could test these predictions.

In summary, extension appears to be the default response, consistent with principles of cognitive economy that favor reusing existing lexical resources over expanding the vocabulary. However, the likelihood of extension depends on which dimension of an action is changed: variations that alter the end state or process of an action are more likely to prompt a new word than variations that merely change the object involved. Speakers implicitly weigh semantic dimensions according to their relevance to verb identity, with event structure playing the largest role. At the same time, the threshold for rejecting an extension is not fixed: it is modulated by nameability.

These results provide behavioral evidence for theoretical distinctions that have been proposed on linguistic and typological grounds. Regularities in verb semantics may reflect universal cognitive biases in what speakers choose to lexicalize and extend (Brochhagen et al., 2023; Fugikawa et al., 2023). More broadly, our results indicate that lexical decisions are indeed systematic: speakers extend existing words when possible, coin new ones when necessary, and calibrate this tradeoff based on both the event structure and the strength of the existing representations.

References

Behrend, D. A. (1990). The Development of Verb Concepts: Children's Use of Verbs to Label Familiar and Novel Events. *Child Development*, *61*(3), 681. https://doi.org/10.2307/1130953

Britt, A. E., Ferrara, C., & Mirman, D. (2016). Distinct Effects of Lexical and Semantic Competition during Picture Naming in Younger Adults, Older Adults, and People with Aphasia. *Frontiers in Psychology*, *7*. https://doi.org/10.3389/fpsyg.2016.00813

Brochhagen, T., Boleda, G., Gualdoni, E., & Xu, Y. (2023). From language development to language evolution: A unified view of human lexical creativity. *Science*, *381*(6656), 431–436. https://doi.org/10.1126/science.ade7981

Bybee, J. (2006). From usage to grammar: The mind's response to repetition. *Language*, *82*(4), 711–733.

Clark, E. V. (1990). On the pragmatics of contrast. *Journal of child language*, *17*(2), 417–431.

Croft, W. (2000). *Explaining language change: An evolutionary approach*. Pearson Education.

Forbes, J. N., & Farrar, M. J. (1993). Children's initial assumptions about the meaning of novel motion verbs: Biased and conservative? *Cognitive Development*, *8*(3), 273–290. https://doi.org/10.1016/S0885-2014(93)80002-B

Fugikawa, O., Hayman, O., Liu, R., Yu, L., Brochhagen, T., & Xu, Y. (2023). A computational analysis of crosslinguistic regularity in semantic change. *Frontiers in Communication*, *8*, 1136338. https://doi.org/10.3389/fcomm.2023.1136338

Gaskell, M., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, *89*(2), 105–132. https://doi.org/10.1016/S0010-0277(03)00070-2

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. *BBN report; no. 4854*.

Gentner, D. (2005). The development of relational category knowledge. In *Building object categories in developmental time* (pp. 263–294). Psychology Press.

Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How efficiency shapes human language. *Trends in cognitive sciences*, *23*(5), 389–407.

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, *73*(2), 135–176.

Gleitman, L. (1990). The structural sources of verb meanings. *Language acquisition*, *1*(1), 3–55.

Goldberg, A. E. (2006). *Constructions at work: The nature of generalization in language*. OUP Oxford.

Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44. https://doi.org/10.1016/j.cogpsych.2017.08.002

Hovav, M. R., & Levin, B. (2010). Reflections on manner/result complementarity. *Syntax, lexical semantics, and event structure*, 21–38.

Kersten, A. W., & Smith, L. B. (2002). Attention to Novel Objects during Verb Learning. *Child Development*, *73*(1), 93–109. https://doi.org/10.1111/1467-8624.00394

Keyes, R. (2021). *The hidden history of coined words*. Oxford University Press.

Lakusta, L., & Landau, B. (2005). Starting at the end: The importance of goals in spatial language. *Cognition*, *96*(1), 1–33.

Lehrer, A. (1974). Semantic fields and lexical structure.

Lyons, J. (1977). *Semantics* (Vol. 2). Cambridge university press.

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence*, *13*(4), 329–358.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, *130*(3), 323.

Rissman, L., Liu, Q., & Lupyan, G. (2023). Gaps in the Lexicon Restrict Communication. *Open Mind*, *7*, 412–434. https://doi.org/10.1162/opmi_a_00089

Talmy, L. (1985). Lexicalization patterns: Semantic structure in lexical forms. *Language typology and syntactic description*, *3*(99), 36–149.

Talmy, L. (2000). *Toward a cognitive semantics: Concept structuring systems* (Vol. 1). MIT press.

Xu, A., Kemp, C., Frermann, L., & Xu, Y. (2024). Word reuse and combination support efficient communication of emerging concepts. *Proceedings of the National Academy of Sciences*, *121*(46), e2406971121. https://doi.org/10.1073/pnas.2406971121

Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological bulletin*, *127*(1), 3.

Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, *115*(31), 7937–7942.

Zettersten, M., & Lupyan, G. (2020). Finding categories through words: More nameable features improve category learning. *Cognition*, *196*, 104135. https://doi.org/10.1016/j.cognition.2019.104135