Exposure to Co-occurrence Regularities in Language Drives Semantic Integration of New Words

Olivera Savic*, Layla Unger*, and Vladimir M. Sloutsky

Department of Psychology, Ohio State University

Word Count: 10,924

Author Note

This work was supported by National Institutes of Health Grants R01HD078545 and P01HD080679 to Vladimir Sloutsky.

*Both authors are the first author on the manuscript, and the present order is purely alphabetical.

All data and materials have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/dt84u/?view_only=b9de59ea94c840c2b232058cfb0a0841.

This study was not preregistered.

Please address correspondence to:

Olivera Savic, Ohio State University, Department of Psychology, 1835 Neil Avenue, PS 268, Columbus, Ohio 43210. Phone: (614) 822-1156. Email: savic.7@osu.edu

Abstract

Human word learning is remarkable: We not only learn thousands of words, but also form organized semantic networks in which words are interconnected according to meaningful links, such as those between *apple*, *juicy*, and *pear*. These links play key roles in our abilities to use language. How do words become integrated into our semantic networks? Here, we investigated whether humans integrate new words by harnessing simple statistical regularities of word use in language, including: (1) Direct co-occurrence (e.g., *eat-apple*) and (2) Shared co-occurrence (e.g., *apple* and *pear* both co-occur with *eat*). In four reported experiments ($N=139$), semantic priming (Experiments 1-3) and eye-tracking (Experiment 4) paradigms revealed that new words became linked to familiar words following exposure to sentences in which they either directly co-occurred, or shared co-occurrence. This finding highlights a potentially key role for co-occurrence in building organized word knowledge that is fundamental to our unique fluency with language.

Total: 150

Keywords: semantic priming; distributional semantics; semantic integration; syntagmatic, paradigmatic

Exposure to Co-occurrence Regularities in Language Drives Semantic Integration of New Words

People learn tens of thousands of words over the course of their lives (e.g., Nation & Waring, 1997). Critically, these words become linked to each other based on their meanings. For example, as we learn the word *melon*, we also learn that it is linked to words such as *apple*, *eat* and *juicy*. Therefore, our vocabularies are not merely large collections of words, but instead function as organized *networks* of words that are connected based on their meanings. Such *semantic* organization plays a key role in our ability to use language to express and comprehend ideas. For example, because *ripe*, *juicy*, and *melon* are semantically linked to *eat*, we can understand that “*I’d like a ripe, juicy melon*” expresses a desire to *eat* a *melon*, even though eating was never mentioned (e.g., DeLong, Urbach, & Kutas, 2005).

Although it is hardly controversial that word knowledge is semantically organized, there is much less clarity with respect to how semantic networks emerge and expand. The current research seeks to shed new light on how our semantic networks expand by integrating new words.

Semantic Links that Organize Word Knowledge

To ground our investigation into the semantic integration of words, we first review two critical types of semantic links that organize our vocabularies: associative and taxonomic links. Associative links (also often referred to as *syntagmatic* or *thematic*) connect words that are commonly used together in language, such as *eat* and *strawberry*¹. On the other hand, taxonomic (or *paradigmatic*) links connect words similar in meaning, such as *melon* and *strawberry*. Both types of semantic links play key roles in our everyday use of language. For example, we can

¹ Note that we refer here to associations between words *in language*, and not to associations between words measured from responses on free association tasks. Responses on free association tasks are an *outcome* of semantic links (e.g., McRae, Khalkhali, & Hare, 2012). Thus, these responses do not tell us how these links are learned. By the same token, the extensive semantic priming research that has investigated the role of “association strength” measured from free association norms (for review, see Hutchison, 2003) cannot shed light on whether *exposure to co-occurrence regularities in language* contribute to learning semantic links between words.

understand that the sentence “*I’d like a ripe, juicy strawberry*” expresses a desire to *eat a strawberry*, even though eating was never mentioned. Such understanding is possible because associative links, such as those connecting *ripe*, *juicy* and *strawberry* to *eat*, are activated during language comprehension (Arnon & Snider, 2010; Kutas & Federmeier, 2000; McDonald & Shillcock, 2003). Similarly, we can use associative and taxonomic links to infer meanings of novel words. For example, we can infer that *dax* is an animal when it appears with associated words such as in “*We saw a furry dax in the zoo*”, or taxonomically related words such as in “*We saw a lion, a zebra, a giraffe and a dax*” (Sloutsky, Yim, Yao, & Dennis, 2017). Moreover, taxonomic links support one of the critical aspects of human intelligent behavior – the generalization of knowledge (e.g., Heit, 2000; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990). For instance, even if we have only ever heard about *strawberries* being *juicy*, taxonomic links between *strawberry* and *melon* allow us to understand and produce utterances about *juicy melons*.

Critically, associative and taxonomic links contribute to language comprehension and production because they support the *rapid* activation of semantically linked words *without the need for deliberative processing*. This is important because language comprehension and production unfold in real time, with words being processed at the timescale of hundreds of milliseconds. We refer to this property of semantically linked words to automatically activate each other as *priming*. The importance of priming is attested by evidence that the activation of semantically related words occurs *online* during language comprehension (DeLong et al., 2005; Kamide, Altmann, & Haywood, 2003; Kutas & Federmeier, 2000), and can thus aid comprehension by activating words that are likely to occur as an utterance unfolds. Further evidence for the importance of priming comes from several findings from research with developing language learners, in which semantic priming is correlated with language comprehension or production skills (Bonnotte & Casalis, 2010; Nation & Snowling, 1999; Nobre & Salles, 2016). Thus, a critical aspect of learning new words is the formation of semantic links with familiar words that can support language comprehension and production through priming. We refer to this

formation of semantic links between new and familiar words that can support priming as *semantic integration*.

Routes for Semantic Integration

As described above, a vital aspect of word learning is the integration of new words into our existing semantic networks. Prior research has investigated some routes through which new words may become semantically integrated, such as when new words are repeatedly paired with their definitions (Clay, Bowers, Davis, & Hanley, 2007; Dagenbach, Horst, & Carr, 1990; Tamminen & Gaskell, 2013) or their real-world referents (Breitenstein et al., 2007; Dobel et al., 2010). While some words may be semantically integrated through these routes, new words are often learned without access to definitions or observable referents.

An alternative route for semantic integration that requires neither definitions nor observable referents is exposure to input that provides high constraints on the meanings of novel words. For example, hearing the sentence “*It was a windy day, so Peter went to the park to fly a dax*” (Borovsky, Elman, & Kutas, 2012; Mestres-Missé, Rodriguez-Fornells, & Münte, 2006) can rapidly foster a link between *dax* and *kite*, because semantic links between *windy*, *park*, and *fly* and the word *kite* generate a strong expectation of *kite* when *dax* appears (DeLong et al., 2005; Kutas & Federmeier, 2000). While this is a potentially important route, it is limited in scope because: (1) Highly constraining input may be uncommon in everyday language (Huettig & Mani, 2016), and (2) It cannot operate when the learner does not already possess semantically organized word knowledge that can generate strong expectations, such as early in development (Mani & Huettig, 2012).

Due to these limitations, previously investigated routes for semantic integration cannot explain the ease with which semantic integration often unfolds - during mere exposure to everyday language, such as while reading, watching television, or having conversations (Nagy, Herman, & Anderson, 1985). The focus of the current research is to examine how new words may become integrated into semantic networks from simple exposure to everyday language.

How Co-Occurrence Regularities May Drive Semantic Integration

The present research investigates whether new words can become integrated into semantic networks from exposure to a source of information that is ubiquitous in everyday language: Regularities with which words co-occur (Ervin, 1961; Harris, 1954). This investigation is motivated by extensive computational evidence that language is rich in co-occurrence regularities that can in principle foster semantic links between words (Asr, Willits, & Jones, 2016; Griffiths, Steyvers, & Tenenbaum, 2007; Hofmann et al., 2018; Huebner & Willits, 2018; Jones & Mewhort, 2007; Jones, Willits, & Dennis, 2015; Landauer & Dumais, 1997; Lund & Burgess, 1996; Miller & Charles, 1991; Rohde, Gonnerman, & Plaut, 2004; Spence & Owens, 1990). Critically, semantic priming phenomena can be predicted by both measurements of co-occurrence regularities in language (Roelke et al., 2018; Willits, Amato, & MacDonald, 2015), and models that form representations of words based on co-occurrence regularities in language (Hare, Jones, Thomson, Kelly, & McRae, 2009; Jones & Mewhort, 2007; Lund & Burgess, 1996). These findings highlight the possibility that co-occurrence regularities can foster semantic integration.

As illustrated in Figure 1, language is rich in two types of regularities of word use that can foster semantic integration: *direct* co-occurrence and *shared* co-occurrence. Direct co-occurrence is the regularity with which words occur together in language, either adjacently or separated by intervening words. For example, when we learn the word *melon*, it can become linked to the words *eat* and *juicy*, because these words reliably directly co-occur in input in sentences such as “*I’d like to eat a nice, juicy melon*”. Hence, direct co-occurrence can foster *associative* semantic links, such as *eat-melon* or *juicy-melon*. Critically, direct co-occurrence regularities can also contribute to semantic links between words that do not themselves directly co-occur, because words similar in meaning tend to share similar patterns of direct co-occurrence with other words. For example, *strawberry* and *melon* may both directly co-occur with words like *eat*, *juicy* or *sweet*. Thus, taxonomic links can be formed based on these *shared* patterns of direct co-occurrence.

Figure 1. Illustration of direct and shared co-occurrence regularities in language input (panel A) and how they can foster associative and taxonomic links between words (panel B).

Co-occurrence is a potentially simple and powerful route for semantic integration because it is ubiquitous in language, and can integrate new words into semantic networks even when input otherwise provides no information about what new words mean. For example, exposure to sentences like “*Sally saw a foobly mipp*” and “*I took a picture of a foobly strawberry*” could drive both *foobly* and *mipp* to become semantically linked with *strawberry*. Hence, following exposure to these sentences, a person could answer questions such as “*Would you like a mipp?*” with some certainty, even if they have never seen or tasted *mipp*, and have only heard it described as *foobly*. Therefore, even *without* observing or being taught what *foobly* or *mipp* refer to, these words could nonetheless *inherit* semantic content from their direct and shared co-occurrence with *strawberry*.

However, for co-occurrence regularities to support semantic integration of newly learned words, it is not sufficient that these regularities are ubiquitous in language. In addition, human learners need to be sensitive to these regularities and able to use them to form semantic links that can be rapidly activated during language processing. In the next section, we evaluate prior evidence supporting this possibility.

Can Exposure to Word Co-Occurrence Foster Priming?

Although no prior studies have *directly* tested the unique contributions of co-occurrence to the semantic integration of new words in human learners, several bodies of evidence suggest that such contributions are plausible.

One body of evidence comes from statistical learning research. This work has primarily focused on sensitivity to direct co-occurrence, and has demonstrated that even infants can spontaneously and rapidly form links between directly co-occurring items across a variety of domains, including speech sounds (Pelucchi, Hay, & Saffran, 2009; Romberg & Saffran, 2010; Saffran, Aslin, & Newport, 1996), tones (Saffran, Johnson, Aslin, & Newport, 1999), and shapes (Bulf, Johnson, & Valenza, 2011; Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002) (see also Coane & Balota, 2010; McKoon & Ratcliff, 1979 for evidence that direct co-occurrence between familiar words can foster priming between them given explicit study or extensive real-world exposure).

Evidence that humans form links based on both direct *and* shared patterns of co-occurrence primarily comes from research in the domains of language and memory. Importantly, participants in these studies were asked to deliberately reason about links between words or objects based on direct and shared co-occurrence (Johns, Dye, & Jones, 2016; Jones & Love, 2007; Matlen, Fisher, & Godwin, 2015; Ouyang, Boroditsky, & Frank, 2017; Zeithamova, Dominick, & Preston, 2012; Zeithamova & Preston, 2010). For example, in Jones and Love (2007), participants explicitly judged words that shared co-occurrence with the same word as more similar than words that co-occurred with different words. While this evidence suggests that humans can link inputs based not only on direct but also on shared co-occurrence, performance in tasks that require or allow participants to reason about these links does not provide evidence for *semantic integration*. As we explained above, such evidence is critical because a key property of semantic links is that they are activated rapidly, without the need for reasoning or other deliberative processes. Research that has evaluated the possibility that shared co-occurrence fosters links that can be automatically activated is comparatively scarce. However, a handful of relevant findings come from research on syntactic learning (Gerken, Wilson, & Lewis, 2005; Lany & Saffran, 2010). For example, to learn grammatical categories of *noun* and *article*, one can capitalize on the fact that members of same syntactic categories tend to share patterns of co-occurrence (e.g., nouns

typically share each other's co-occurrence with articles). Although much of this research has assessed learning from deliberative processing (e.g., Frigo & McDonald, 1998; Mintz, 2002), a small number of studies have assessed learning from spontaneous behavior such as looking in infants (Gerken et al., 2005; Lany & Saffran, 2010). The results of these studies suggest that shared co-occurrence can support learning of syntactic categories *when they are correlated with other cues*. For example, links between words that have the same number of syllables are learned more readily when these words also share co-occurrence. However, while cues such as syllable number or morphological markers of grammatical class (e.g., English verbs typically end in '-ed', '-s', or '-ing') can support the learning of *grammatical* categories such as article or verb, they are not useful for learning *semantic* links. For example, these cues are not available to support learning taxonomic links between words for fruits such as "strawberry" and "melon". Thus, it remains an open question whether co-occurrence regularities *alone* can drive the integration of new words into semantic networks.

Present Study

The present research investigated whether exposure to language input containing direct and shared co-occurrence regularities can foster new semantic links between novel and familiar words. Participants were presented with sentences in which familiar words directly co-occurred and shared co-occurrence with novel words (Figure 2). For example, in sentences such as, "*I'd like a foobly apple*" and "*She saw a foobly mipp once*", the novel word *foobly* directly co-occurs with both a familiar word *apple*, and a novel word *mipp*. Moreover, although *apple* and *mipp* never appear in the same sentence, they share co-occurrence with *foobly*. Immediately following exposure to sentences, we assessed semantic integration by testing whether the familiar words came to be *primed* by novel words with which they either directly co-occurred or shared co-occurrence in training sentences. We measured priming from a reaction time paradigm (Experiments 1-3) and from spontaneous looking behavior in an eye tracking paradigm (Experiment 4).

There are four important aspects of the current design. First, this research examined whether co-occurrence regularities *alone* can foster semantic links, without additional, supportive cues (cf. Gerken et al., 2005; Lany & Saffran, 2010). Second, because a key hallmark of semantic links is that they are activated without the need for deliberative processing (McNamara, 2005; Neely, 2012), the current research assesses whether familiar words come to be spontaneously activated or *primed* by novel words with which they directly co-occur or share co-occurrence (cf. Jones & Love, 2007; McNeill, 1963). Third, the research explicitly examines sensitivity to both direct and shared co-occurrence in order to test whether this route for semantic integration can support both types of semantic links that are vital in semantic networks: associative and taxonomic links. Finally, the exposure to co-occurrence regularities was designed based on the consideration that much of prior research into other routes for semantic integration has required extensive training over days and weeks to observe evidence that newly learned words prime familiar words (e.g., Breitenstein et al., 2007; Clay et al., 2007; Dobel et al., 2010; Mestres-Missé et al., 2006; Tamminen & Gaskell, 2013). To provide an evaluation of the co-occurrence route to semantic integration, we instead opted for an approach modeled on statistical learning studies, which typically investigate whether brief exposure to input containing a constrained set of regularities can foster learning. Specifically, we investigated whether brief exposure to a constrained set of co-occurring novel and familiar words can foster semantic integration. This choice was additionally motivated by the goal to provide a paradigm that can be used to evaluate this route to semantic integration during development, when much of word learning takes place.

Experiment 1

Method

Participants

Participants were 45 undergraduate students from a large Midwestern university located in a metropolitan area, who received course credit for their participation. This study was approved by the *Blind for Review* University Institutional Review Board (protocol entitled *Blind for Review*), and

all participants provided informed consent. The target sample size was based on both prior studies that used priming measures to assess the effects of other, non-co-occurrence-based routes for semantic integration (Coutanche & Thompson-Schill, 2014; Tamminen & Gaskell, 2013), and prior studies that used other, non-priming measures to assess the effects of co-occurrence regularities on links between either novel or familiar words (Jones & Love, 2007; McNeill, 1963). Four participants were excluded from main analyses: Two due to failure to reach the minimum of 75% accuracy on control questions designed to check whether participants were actually reading sentences during self-paced reading, and two due to failure to reach a criterion of 80% overall accuracy in the priming task.

Training Stimuli

Training stimuli and data from a norming study used to select these stimuli are available on OSF (https://osf.io/dt84u/?view_only=b9de59ea94c840c2b232058cfb0a0841).

See Figure 2 for an illustration of the training stimuli. The training stimuli were two triads of words (1: *foobly-apple-mipp*; 2: *dodish-horse-geck*), each consisting of a pseudoadjective (*foobly*; *dodish*), familiar noun (*apple*; *horse*), and a pseudonoun (*mipp*; *geck*). Within each triad, the pseudoadjective preceded (a) the familiar noun in half the sentences, or (b) the novel pseudonoun in other sentences. As noted above, the choice to use these constrained sets of stimuli was motivated by goals to: (1) Evaluate whether even brief exposure to co-occurrence regularities can foster semantic integration (in comparison to the days and weeks of repeated exposure studied in other routes for semantic integration, e.g., Breitenstein et al., 2007; Clay et al., 2007; Dobel et al., 2010; Mestres-Missé et al., 2006; Tamminen & Gaskell, 2013), and (2) Provide a paradigm that can be used to investigate the co-occurrence route to semantic integration during development, when much word learning takes place.

Figure 2. Left: Examples of training sentences. Right: Illustration of the two triads. Direct co-occurrences are denoted by solid blue lines, and shared co-occurrences are denoted by dashed red lines. Word font and capitalization are for illustrative purposes; all text was in the same font and case in the experiment.

The pseudowords used in Training were tested in a norming study ($N = 20$) to ensure that they did not evoke any meaning (Chuang et al., 2020) related to the selected familiar words (i.e. *apple* and *horse*). The norming study used: (1) A free association task to assess whether participants spontaneously produced fruit or mammal words in response to pseudowords, and (2) A labeling task to assess whether participants preferentially labeled fruits or mammals with some of the pseudowords over the others. Pseudowords were not associated with fruits or mammals according to either measure (see Supplemental Materials). Therefore, we used the same assignment of pseudowords to triads across all experiments and across all participants².

Each directly co-occurring word pair from these triads (*foobly-apple*, *foobly-mipp*, *dodish-horse*, and *dodish-geck*) was embedded in 10 unique sentence frames, for a total of 40 Training sentences. These sentences were designed to allow us to test whether exposure to the direct and

² Note that although the reported norming study does address the issue of the potential effects of pre-existing preferences in linking pseudowords and familiar words, a design in which assignment of pseudowords and familiar words would be fully counterbalanced across subjects (i.e., for some participants *foobly* would be linked to word *apple*, for others to word *horse*) would be preferred and is recommended for future studies.

shared co-occurrence regularities in these sentences fosters the formation of semantic links between each familiar noun (*apple*, *horse*) and (1) The pseudoadjective with which it directly co-occurred in sentences (*foobly-apple*, *dodish-horse*) and (2) The pseudonoun with which it never directly co-occurred, but instead shared patterns of co-occurrence across different sentences (*mipp-apple*, *geck-horse*).

To ensure that any new semantic links could be attributed only to exposure to co-occurrence regularities during Training, Training sentences did not convey any other cues to pseudoword meaning (e.g., “I’d like a foobly mipp”).

Testing Stimuli

To test semantic integration, we examined whether the novel words primed the familiar words with which they directly co-occurred or shared co-occurrence (see Table 1 for full set of Prime-Target word pairs). We therefore constructed two types of Related Prime-Target word pairs for which priming was expected: (1) Direct, in which a pseudoadjective Prime was paired with the familiar noun with which it directly co-occurred in Training sentences (*foobly-apple*; *dodish-horse*), and (2) Shared, in which a pseudonoun Prime was paired with the familiar word with which it shared co-occurrence in Training sentences (*mipp-apple*; *geck-horse*).

Table 1

Full set of Related and Unrelated Prime-Target word pairs.

Relatedness	Type	Triad 1		Triad 2	
		Prime	Target	Prime	Target
<i>Related</i>	Direct	foobly	apple	dodish	horse
	Shared	mipp	apple	geck	horse
<i>Unrelated</i>	Direct	dodish	apple	foobly	horse
	Shared	geck	apple	mipp	horse

In addition to Related Prime-Target pairs, we also constructed two types of Unrelated word pairs for which priming was *not* expected. Unrelated word pairs violated the co-occurrence regularities in Training sentences. Specifically, Target words from one triad were paired with words from the other triad (Table 1). Unrelated pairs either violated direct co-occurrence regularities, by pairing a pseudoadjective from one triad and the familiar noun from a different triad (*foobly-horse*; *dodish-apple*), or violated shared co-occurrence regularities, by pairing a pseudonoun from one triad and the familiar noun from a different triad (*mipp-horse*; *geck-apple*). Because cross-pairings in Unrelated pairs violated Training sentence regularities, the pseudoword Primes should not prime the familiar word Targets, and may even instead produce interference. In addition to Related and Unrelated word pairs, we also generated filler pairs in which Targets were paired with novel pseudowords (*nuppical*; *bof*) that did not appear in Training. Therefore, the set of stimuli used in the test phase consisted of six pseudowords (*foobly*; *dodish*; *nuppical*; *mipp*; *geck*; *boff*), two familiar words (*apple*; *horse*) and two pictures depicting the familiar words from the Triads, one of an apple and one of a horse.

Procedure

The experiment had two phases: Training and Test. Verbatim instructions for all tasks across experiments are available in the Supplemental materials.

Training. Training consisted of three blocks. In each block, participants read the 40 Training sentences in a random order at their own pace. In addition, to identify whether participants actually read sentences rather than clicked through them, each block contained attention check questions and a free association task.

First, attention check questions followed three randomly selected Training sentences, and prompted participants to type the novel words from the sentence they had just read. Second, free association trials were presented at the end of each block of training. In free association trials, participants were presented with one of the pseudowords from the training sentences as a prompt and asked to respond by typing the first (pseudo) word that came to their mind upon reading the

prompt. During the free association task, each of the pseudowords (*foobly*, *dodish*, *mipp*, *geck*) was presented 3 times in a randomized order. Note that the free association task was originally included following McNeill (1963, 1966), to provide a comparable assessment of the formation of direct and shared co-occurrence-based links. However, the free association task effectively assesses only participants' dominant responses and not all the links between words they have formed (as such, free association responses in this and subsequent experiments were overwhelmingly based on direct co-occurrence). Therefore, our main test of semantic integration used the semantic priming task described in the next section. The full set of three Training blocks took approximately 10-15 minutes.

Test. For the test phase, participants performed a *priming* task (see Figure 3 for timing of events in trials). At the start of each trial, participants saw a fixation cross (500ms), followed by two images, one on either side of the screen: A horse, and an apple. The two-image display was presented alone for 500ms. After 500ms, the Prime and then the Target word also appeared as text on the top of the screen (see Table 1). The Prime word appeared for 300ms, and the subsequently presented Target word remained on the screen until participant responded. The participants' task was to read both words and identify whether the second word (Target word) labeled the image on the right or the image on the left. The position of the pictures of the apple and the horse, and the position of the Related and Unrelated Target, were counterbalanced across trials. To get a precise measure of reaction time, participants responded by clicking the right or left mouse button (note that participants merely clicked the left or right mouse button: they did *not* move the mouse to and click on the labeled image).

Participants first completed 8 practice trials of this task, in which the two words consisted of filler word pairs (i.e., a new pseudoword that did not appear during Training, followed by apple or horse). Participants then completed the Test task itself, which contained a total of 96 trials: 8 repetitions of the full set of Prime-Target pairs (64 trials), and 32 filler trials.

Figure 3. Timing of events in the priming task (Experiments 1-2).

Participants were given an unlimited time to make their responses. However, to encourage participants to respond quickly rather than deliberatively, participants were instructed to respond quickly at the beginning of the task, and additionally encouraged to respond faster on trials in which their response time was greater than 800ms.

Results

All analyses were performed in the R environment for statistical computing (R Development Core Team, 2008). Analyses using mixed effects models were conducted using the lme4 (Bates, Maechler, Bolker, & Walker, 2015) and lmerTest (Kuznetsova, Brockhoff, & Christensen, 2017) packages. Data and analysis scripts for all experiments reported here are available at https://osf.io/dt84u/?view_only=b9de59ea94c840c2b232058cfb0a0841.

Preliminary analyses: Attention to Sentences and Pseudoword Forms

To assess whether participants attended to the Training sentences and learned the pseudoword forms, we analyzed participants' responses on the attention check questions and the free association task.

Performance on attention check questions was high ($M = .94$, $SD = .08$), which confirmed that participants read the sentences. Performance of two participants was below .25 accuracy, so their data were excluded from the further analyses.

In the free association task, participants were asked to respond to the prompt word with one of the training triad words. They responded as instructed on an average 96% of the free association trials presented at the end of training. In addition, they tended to respond with training words that had directly co-occurred with the prompt word. While 81% of participants' responses were based on direct co-occurrence, only 2% were based on shared co-occurrence regularities³.

The free association data confirms that participants read sentences and learned the pseudowords. In addition, it shows that participants were sensitive to direct co-occurrence regularities in the input. As we noted previously, this task only probes participants' most *dominant* association with a word. Thus, any sensitivity to shared co-occurrence regularities could have been masked by a dominant tendency to respond based on direct co-occurrence (cf. McNeill, 1963). Therefore, we did not use this task as the main test of whether semantic links were formed. For this purpose, we instead used the priming task.

Main Analyses: Priming

The purpose of the main analyses was to investigate whether the novel pseudowords were semantically integrated with familiar words on the basis of co-occurrence regularities in Training sentences. We accomplished this by measuring whether the novel pseudowords affected the speed of processing familiar words in a semantic priming task. Specifically, we tested whether participants more rapidly identified a familiar noun (Target: *apple*, *horse*) when it was preceded by a novel pseudoword (Prime) in the Related (Direct and Shared) versus the Unrelated (Direct and Shared) condition. Following the logic of extensive semantic priming research (e.g., McRae & Boisvert, 1998), if participants linked pseudowords with familiar words based on direct and

³ Please note that here and in all subsequent experiments the proportion of responses congruent with direct and shared co-occurrence regularities was corrected for guessing. This was needed in order to more accurately reflect true learning and differentiate it from high proportions of congruent responses that could spuriously result from simple guessing given that the number of possible responses was restricted to six words. Complete data, coding schema and steps in analyses of attention check questions and free association data are available at https://osf.io/dt84u/?view_only=b9de59ea94c840c2b232058cfb0a0841.

shared co-occurrence, pseudowords should prime the familiar words from the same triad. Specifically, novel pseudowords should allow participants to respond more quickly to Targets from the same triad (Related condition) than to Targets from the opposite triad (Unrelated condition). Prior to analyzing reaction times, we removed data from both incorrect trials, and trials with extremely short (<200ms) and extremely long response latencies (>1500ms). This resulted in a removal of 5.6 % of all trials. Summary statistics are reported in Table 2.

Table 2

Mean reaction times and standard deviations across the conditions in Experiments 1-3

Type	Relatedness	Experiment 1 (N = 41)	Experiment 2 (N = 42)	Experiment 3 (N = 36)
		M (SD)	M (SD)	M (SD)
Direct	Related	534 (138)	523 (126)	524 (156)
	Unrelated	547 (129)	534 (134)	540 (145)
Shared	Related	535 (129)	516 (129)	524 (143)
	Unrelated	547 (132)	530 (135)	547 (156)
Filler		534 (130)	514 (126)	548 (153)
Priming effect (ms)		12.5	13.0	19.8

We analyzed reaction times by fitting them to linear mixed effects models with fixed effects of Prime Type (levels: Direct and Shared), Relatedness (levels: Related and Unrelated), and their interaction. The random-effects structure was based on the log likelihood ratio test (Wagenmakers & Farrell, 2004). Specifically, following Wagenmakers and Farrell (2004), we compared models with the same fixed-effects structure but varying complexity in their random-effects structure, and settled on the simplest among the candidate models that provided best fit of the data. The best fitting random effects structure, as indicated by log-likelihood ratio test, included only a random

intercept for participant and random intercept for stimuli (i.e. Triad)⁴. This model revealed no significant effect of Prime Type, neither as a main effect nor in interaction with Relatedness ($F_s < 1.0$, $p_s > .10$). Critically, the model revealed a significant effect of Relatedness, $F(1, 2443.4) = 5.85$, $p = .016$, with participants responding faster in Related than in Unrelated conditions (Figure 4, left panel). In other words, participants responded faster to familiar words (Targets) when they were preceded by novel pseudowords with which they directly co-occurred or shared co-occurrence in training (Related Prime), than when they were preceded by novel pseudowords that directly co-occurred or shared co-occurrence with a different familiar word (Unrelated Prime).

Figure 4. Mean reaction times for direct and shared co-occurrence primes across three experiments. Error bars indicate the standard errors of the means.

Follow-Up Analyses: Facilitation and Inhibition

⁴ Note that in this and the following experiments, we additionally tested whether effects of Relatedness and Type varied across the two triads. We found no significant interactions between the triad and other fixed effects, which confirmed that the Relatedness effect did not depend on a specific stimulus set. Given that comparing differences between the triads is not otherwise of interest in the study, we included triad as a random effect in the reported analyses.

The analyses presented so far focused on the contrast between the related and unrelated primes on word processing. However, the experimental design also includes filler trials which may offer additional insights on the mechanism of the reported priming effect. Specifically, on filler trials the primes were novel words not presented at training. These words were thus completely neutral – they could neither speed up nor slow down word processing. Therefore, the analyses of filler trials can clarify whether the reported priming effects, the difference between the related and unrelated trials, results from the facilitation effects, inhibition effects, or both.

To analyze participants' performance on filler trials and compare it to the performance on related and unrelated trials, we ran a mixed effects model that included a fixed effect of Relatedness (three levels: Related, Unrelated, and Filler) and a random intercept for participants. The analysis revealed that in Experiment 1 participants responded slower on unrelated than on filler trials, $B = 13.21$, $SE = 4.92$, $t = 2.68$, $p < .01$, but were equally fast to respond on related and filler trials, $B = 1.46$, $SE = 4.92$, $t = 0.30$, $p > .10$. Thus, the observed priming effects in Experiment 1 could have stemmed from inhibition - unrelated primes slowed participants' responses relative to filler and related primes (see Figure 5, left panel). Follow-up analyses confirmed that this effect did not depend on prime type, i.e. both direct and shared co-occurrence primes produced inhibition effects.

Note that the co-occurrence training in the current study could teach the participants not only which words go together (positive information), but also which words do not go together (negative information). While the difference between the related and unrelated conditions could be driven by either of these types of information, the reported inhibition effect potentially suggests that the automatic semantic priming in Experiment 1 was primarily driven by the negative information. This interpretation is motivated by the recent work on the role of the negative information in word learning (Johns, Mewhort, & Jones, 2019; Ramscar, Sun, Hendrix, & Baayen, 2017; Roembke, Wasserman, & McMurray, 2016).

Figure 5. Mean reaction times for Related, Unrelated and Filler trials across three experiments. Error bars represent standard errors of mean.

Discussion

In this experiment, we observed that novel pseudowords came to prime familiar words following brief exposure to sentences in which they directly co-occurred, or shared co-occurrence with other words. This priming effect transpired with a brief 300ms presentation of novel pseudowords (as in studies of semantic priming between familiar words, Hutchison et al., 2013; McNamara, 2005; Neely, 2012) in responses that occurred on average under 550ms following the onset of familiar words. Thus, this finding captures a rapid process that does not require slow deliberative reasoning about relations between words (cf. Jones & Love, 2007). Importantly, the formation of novel semantic links was supported solely by the regularities with which novel and familiar words occurred in sentences participants read during the training (cf. Gerken et al., 2005; Lany & Saffran, 2010). Therefore, our findings support the proposal that co-occurrence regularities can drive the semantic integration of new words into semantic networks.

It is worth noting that the pattern of similar priming effects for both direct and shared co-occurrence was not captured by responses on the free association task. Although the free association task was initially included to provide a converging measure of the formation of links

between novel and familiar words modeled on McNeill (1963), free association responses were overwhelmingly dominated by direct co-occurrence, and therefore missed evidence for the formation of links based on shared co-occurrence that the priming task captured. This discrepancy is important for three reasons. First, it highlights the importance of the priming task as a sensitive measure of the links that participants formed. Second, it suggests that the links that the priming task detected likely came from exposure to co-occurrence regularities in the training, rather than that free association itself provided some form of training. Specifically, if the free association task played a role in forming the links assessed in the priming task, it could only serve to rehearse and strengthen links based on direct co-occurrence, as they dominated responses in free association task. In contrast, free association could not have rehearsed links based on shared co-occurrence, because responses based on shared co-occurrence were extremely rare. Finally, it is possible that participants interpreted the training not as a reading task as instructed, but as a form of a memory task due to the attention check and free association questions. While this aspect of the study can be taken as a potential limitation, it is important to note that this type of participants' strategy could only support better memory for the individual words and potentially the direct co-occurrence links, but is highly unlikely that it could support the formation of shared co-occurrence links. The reported discrepancy between the performance in the free association task, which is slow and explicit, and the priming task, which is fast and targets automatic processing, supports this interpretation.

In Experiment 2, we further tested the robustness of the priming effects observed in Experiments 1. Specifically, in Experiment 1 the formation of links between the novel pseudowords and the familiar noun in the same triad may have been facilitated by two characteristics of the pseudoadjectives. Specifically, the pseudoadjectives had suffixes (i.e., -ly and -ish) that are diagnostic of adjectives in English. Moreover, these suffixes rendered the pseudoadjectives in the two triads very distinctive from each other. This distinctiveness could have helped participants form links between a given pseudoadjective and the specific familiar

noun and pseudonoun from the same triad with which it directly co-occurred. Therefore, in Experiment 2, we modified the pseudoadjectives to have the same suffix – i.e., -ing – that is not uniquely diagnostic of adjectives in English.

Experiment 2

Participants

Out of 45 undergraduates recruited from the same population as Experiment 1, three were excluded due to not reaching the accuracy criterion (75%) for control questions during the training phase.

Stimuli and Procedure

The stimuli and procedure used in this experiment were identical to Experiment 1, with the exception that the pseudoadjectives *foobly* and *dodish* were replaced with *foobing* and *doding*.

Results

Preliminary analyses: Attention to Sentences and Pseudoword Forms

Participants responded accurately on attention check questions on an average of 95% of all trials ($SD = .11$). Three participants performed very low on these questions (below 25% accurate responses). Therefore, their data was excluded from further analyses. Similar to Experiment 1, 98% of responses in the free association task were training triad words, out of which 79% was based on direct co-occurrence and 2% on shared co-occurrence regularities.

Main Analyses: Priming

Main analyses followed the logic and steps described for Experiment 1. Data from incorrect trials, and trials with extremely short (<200 ms) and extremely long response latencies (>1500 ms) was removed. This resulted in a removal of 3.9 % of all trials.

The linear mixed model with Prime Type, Relatedness, and their interaction as fixed effects and random intercepts for participants and stimuli revealed significant effects of Relatedness, $F(1, 2540.3) = 7.57$, $p = .006$ and no significant effect of Prime Type or the interaction ($p > .10$). Thus,

Experiment 2 replicated the main effect of Relatedness found in Experiment 1, as participants were faster on Related than Unrelated trials (Figure 4, panel c).

Follow-Up Analyses: Facilitation and Inhibition

Following the same logic as in Experiment 1, participants' performance on filler trials was compared to the performance on related and unrelated trials. We ran a mixed effects model that included a fixed effect of Relatedness (three levels: Related, Unrelated, and Filler) and a random intercept for participants. The analyses revealed that in Experiment 2, as in Experiment 1, priming effects consisted of inhibition. Participants were equally fast on related and filler trials, $B = 5.37$, $SE = 4.74$, $t = 1.13$, $p > .10$, but they were slower on unrelated than filler trials, $B = 18.41$, $SE = 4.73$, $t = 3.90$, $p < .001$, (see Figure 5, central panel). Again, the effect was the same for both direct and shared co-occurrence primes.

Discussion

Experiment 2 replicated the evidence provided by Experiment 1 that exposure to direct and shared co-occurrence fosters the semantic integration of novel words with familiar words. The main contribution of Experiment 2 is to support the robustness of this effect. Specifically, Experiment 2 replicated the evidence in the preceding studies for semantic integration even when morphological markers did not distinguish between pseudoadjectives with which familiar and pseudonouns shared co-occurrence.

As stated in the introduction, the main goal of the current study was to examine whether new co-occurrence-based links can support the *rapid* activation of linked words *without the need for deliberative processing*. Therefore, following prior studies of automatic semantic priming, Experiments 1 and 2 used a short prime duration. However, some prior evidence suggests that the time course of semantic priming effects may vary across the two kinds of semantic links (i.e., associative and taxonomic links) that may be fostered by direct and shared co-occurrence (Hutchison, 2003; McNamara, 2005; Roelke et al., 2018). Therefore, Experiment 3 serves to test

whether links formed based on direct and shared co-occurrence both remain robust given a longer prime duration of 1500ms.

Experiment 3

Participants

Participants were 45 undergraduate students. Nine participants were excluded from main analyses: One due to failure to reach the minimum of 75% accuracy on control questions, one due to failure to complete the experiment within the 30-minute test session, and seven due to failure to reach a criterion of 80% overall accuracy in the priming task.

Results

All analyses reported here followed the logic described for Experiments 1-2.

Preliminary analyses: Attention to Sentences and Pseudoword Forms

Performance on attention check questions was high ($M = .96$, $SD = .07$), which confirmed that participants read the sentences. Performance of only one participant was below .25 accuracy, so their data were excluded from the further analyses.

In the free association task, participants responded as instructed by responding to the prompt with one of the training words on an average 98% of the trials. In addition, they tended to respond with training words that had directly co-occurred with the prompt word. While 86% of the responses were based on direct co-occurrence, only 6% were based on shared co-occurrence regularities. This data (1) confirms that participants read sentences and learned the pseudowords and (2) shows that participants were sensitive to direct co-occurrence regularities in the input.

Main Analyses: Priming

Prior to analyzing reaction times, we removed data from both incorrect trials, and trials with extremely short (<200ms) and extremely long response latencies (>1500ms). This resulted in a removal of 7.3 % of all trials.

The linear mixed model with Prime Type, Relatedness, and their interaction as fixed effects and random intercept for participants revealed no significant effect of Prime Type, neither as a

main effect nor in interaction with Relatedness ($F_s < 1.0$, $p_s > .40$). Critically, the model revealed a significant effect of Relatedness, $F(1, 2085.9) = 10.06$, $p = .002$, with participants responding faster in Related than in Unrelated conditions (Figure 4, panel C). In other words, participants responded faster to familiar words (Targets) when they were preceded by novel pseudowords with which they directly co-occurred or shared co-occurrence in training (Related Prime), than when they were preceded by novel pseudowords that directly co-occurred or shared co-occurrence with a different familiar word (Unrelated Prime).

Follow-Up Analyses: Facilitation and Inhibition

A mixed effects model that included a fixed effect of Relatedness (three levels: Related, Unrelated, and Filler) and a random intercept for participants revealed that in Experiment 3, there was no difference in reaction times on filler and unrelated trials, $B = -5.10$, $SE = 6.11$, $t = -0.67$, $p > .10$, but participants responded faster on related than on filler trials, $B = -23.53$, $SE = 6.12$, $t = -3.89$, $p < .001$. Thus, while in Experiments 1 and 2 priming effect consisted of inhibition, in Experiment 3 they consisted of facilitation - related primes speeded responses relative to filler and unrelated primes. The difference in the mechanisms driving semantic priming may be related to the manipulation of prime duration. While in Experiments 1 and 2 we used a short prime duration (300ms), in Experiment 3 long prime duration (1,500ms) was used. Future research might shed light on the time course with which newly learned and semantically integrated words activate and compete with known words.

Discussion

Experiment 3 replicates and extends the key findings of Experiments 1 and 2. Specifically, it demonstrates that new words can prime familiar words with which they directly co-occur or share patterns of co-occurrence at both short and long prime durations.

Credibility of Priming Effects

Across three experiments, we demonstrated that exposure to co-occurrence regularities in language fosters semantic integration of novel words. However, similar to prior studies of other

routes for semantic integration (e.g., Dagenbach et al., 1990; Tamminen & Gaskell, 2013), the magnitude of the priming effects that provided this evidence was small (11-14ms). Therefore, to assess the credibility of these small priming effects, we conducted an additional Bayesian analysis of priming effects across the two experiments.

To assess the credibility of priming effects, we fit the priming data from the Experiments 1-3 using Bayesian multilevel linear models. We analyzed the data both from all three experiments together, and from just Experiments 1 and 2 that used a short, 300ms prime duration. All analyses were conducted in the R environment using the *rstanarm* package to fit models (Goodrich, Gabry, Ali, & Brilleman, 2020) and the *bayestestr* package to calculate credible intervals (Makowski, Ben-Shachar, & Lüdtke, 2019). Note that in contrast with confidence intervals calculated for frequentist analyses, credible intervals directly capture the probability that the magnitude of an effect is within the interval. These models were designed to be comparable to the frequentist mixed effects models used to analyze the results of Experiments 1 and 2. In these models, reaction time was predicted by Prime Type, Relatedness, and their interaction, with estimated intercepts for participants and experiments. To determine whether each factor produced a priming effect, we calculated 89% Highest Density credible intervals, and inferred a priming effect when the interval did not include 0. We used the 89% credible interval following some Bayesian statisticians (Makowski et al., 2019; McElreath, 2020) to highlight the distinction with confidence intervals.

Figure 6. Posterior distributions of priming effects for Relatedness, Prime Type, and their interaction for: All prime durations (Experiments 1-3) and 300ms prime duration experiments only (Experiments 1-2). The 89% credible interval for each distribution is shown in blue.

Posterior distributions for priming effects are depicted in Figure 6. Across experiments, only the effect of Relatedness had an 89% probability of falling within a range that did not include 0 (All prime durations: 6.54ms - 20.08ms, 300ms prime duration only: 4.36ms - 19.67ms).

Discussion

The results of the Bayesian analyses corroborate the evidence from frequentist analyses that exposure to direct and shared co-occurrence regularities caused novel pseudowords to prime familiar words, even for short (300ms) prime durations.

Experiment 4 further investigated the robustness and generalizability of co-occurrence-driven semantic integration by investigating whether integration of pseudowords is evident when assessed using an entirely different paradigm. Specifically, Experiment 4 was an eye tracking study which used a version of the extensively used visual world paradigm (Altmann & Kamide, 1999; Huettig, Rommers, & Meyer, 2011; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy,

1995) to investigate whether pseudowords would influence spontaneous looking at pictures of the familiar nouns with which they directly co-occurred or shared co-occurrence.

Importantly, following the design of the recently published eye-tracking study (Unger, Savic, & Sloutsky, 2020), participants in Experiment 4 only made behavioral responses when they heard familiar words, and were not required to provide any response on the trials on which they heard the pseudowords. Since participants could completely ignore the pseudoword trials and still excel in the task, any priming effects in Experiment 4 would reflect the spontaneous activation of familiar words in the mere presence of newly learned pseudowords with which they directly co-occurred or shared co-occurrence in the preceding training.

Experiment 4

Method

Participants

Participants were 24 undergraduate students from a large Midwestern university located in a metropolitan area, who received either course credit or \$10 for their participation. The sample size was based on the sample size needed to detect semantic links between familiar words in a previously published version of this paradigm (Unger et al., 2020). Post-hoc power analyses confirmed that this sample size was sufficient to detect the effects observed (see Supplemental Materials).

Training Stimuli

Training stimuli in Experiment 3 were similar to those used in prior experiments, with the following exceptions: (1) we created a separate set of 40 sentences for each block, and (2) sentences were recorded by a female native English speaker.

Testing Stimuli

Testing stimuli included four pictures of apples, four pictures of horses, and the six words from the triads recorded by the same speaker as the Training sentences. Word recordings were edited in Praat to be 850ms in duration.

Apparatus

This experiment used an EyeLink Portable Duo eye tracking system that measures eye gaze by computing the pupil-corneal reflection at a sampling rate of 500Hz. We additionally constructed a non-functional “button box” with yellow and blue buttons for use in the cover task that participants completed during the experiment (see Procedure).

Procedure

The experiment had two phases: Training and Test.

Training. The Training phase in Experiment 3 was similar to Training in the preceding experiments, with the exception that participants listened to recordings of the Training sentences rather than reading them. In addition, due to the auditory presentation of sentences, the attention check questions were removed. These questions were used in the self-paced reading paradigm in Experiments 1 and 2 to catch whether participants proceeded through sentences without reading them, but were not necessary in Experiment 3 because participants could not skip through the sentence recordings without hearing them in full.

Test. The Test phase used an eye tracking paradigm. During each trial, participants saw a picture of an apple and a picture of a horse on diagonally opposite corners of the computer screen, one on a yellow square, and one on a blue square (see Figure 6). Pictures were presented alone for 500ms, after which participants heard a word. The pictures remained on the screen while gaze was recorded for 3,000ms following the onset of the word. There were 16 cover task trials and 32 experimental trials. For the cover task, participants were instructed that they would sometimes hear the word “apple” or “horse”, and that their task was to click a button on a button box corresponding to the color of the square on which the labeled picture appeared. On experimental trials, participants had no task. On these trials, participants heard a Prime word, i.e. one of the four triad pseudowords, and made no response. As in Experiments 1-2, the experimental trials belonged to one of two Prime Type conditions: Direct, in which the Prime was a pseudoword that directly co-occurred with *apple* or *horse* during Training, and Shared, in which the Prime was a

pseudoword that shared co-occurrence with *apple* or *horse* during Training (see Table 1). Thus, on each experimental trial, one picture was the Related picture (i.e., the picture of the familiar noun that directly co-occurred or shared co-occurrence with the Prime), and the other was the Unrelated picture. Note that the current design allows for differences in looking to be attributed solely to the formation of new co-occurrence-based links, given that the visual input is identical across the Relatedness conditions. The position of the Related and Unrelated pictures, the position of the apple and the horse, and the color of the square on which each picture appeared were counterbalanced across trials.

The pictures of apples and horses were each 250 x 225 pixels (subtending approximately 4.5 x 4° of visual angle). Areas of Interest for the Related and Unrelated pictures on experimental trials were defined as 425 x 400 pixel rectangles.

Figure 6. Timing of events in the eye tracking paradigm in Experiment 3. Dashed lines depict Areas of Interest.

Results

Preliminary analyses: Attention to Sentences and Pseudoword Forms

In the free association task, participants responded with one of the training words on all trials, which confirmed that they listened to training sentences and learned novel pseudowords. In

addition, 45% of their responses were based on direct co-occurrence, and only 4% were based on shared co-occurrence regularities from training sentences.

Main Analyses: Primed Gaze

As in the preceding experiments, the purpose of the main analyses was to investigate whether the novel pseudowords were semantically integrated with familiar words on the basis of co-occurrence regularities in Training sentences. Specifically, we measured whether hearing a novel pseudoword prompted more looking at the picture of the familiar noun with which it directly co-occurred or shared co-occurrence (i.e., the Related picture), versus the picture of the familiar noun from the contrasting triad (i.e., the Unrelated picture). Thus, similar to the preceding experiments, we investigated the effects of Relatedness (Related: picture from same triad vs. Unrelated: picture from opposite triad) for two Prime Types (Direct and Shared).

To conduct this analysis, we first divided the 3,000ms of gaze recorded following Prime onset into six 500ms time bins, then calculated the total dwell time to the Related and Unrelated picture during each bin in each experimental trial. We then analyzed whether dwell time was greater for the Related versus Unrelated picture in each 500ms time bin during Direct and Shared Prime trials (Figure 7). We opted for this analysis approach instead of more fine-grained analyses of the time course of looking such as growth curve analysis (Mirman, Dixon, & Magnuson, 2008) due to concerns about the statistical robustness of existing approaches for such analyses (Huang & Snedeker, 2020; McMurray, 2020). For example, fine grained time course analyses typically treat dwell time as a continuous variable that unfolds over short timescales without correlated errors, whereas dwell time is in fact the outcome of a series of relatively long, discrete events – i.e., fixations separated by saccades – such that looking during one sample is often part of the same fixation as looking in a preceding sample.

Figure 7. Average looking time to Related and Unrelated pictures in each 500-ms time bin over the 3,000ms following Direct Prime onset (left) and Shared Prime onset (right). Error bars depict standard errors of the mean.

We analyzed the effects of Relatedness and Prime Type in each bin using the same Bayesian multilevel linear models used to analyze the credibility of priming effects across Experiments 1-2. Specifically, each model predicted dwell time in a time bin based on fixed effects of Prime Type (levels: Direct and Shared), Relatedness (levels: Related and Unrelated), and their interaction, as well as a random intercept for participant. Main effects and interactions were inferred when their Highest Density credible intervals did not include 0. We followed up any interactions by fitting a model in which dwell time was predicted by Relatedness separately for each of the two Prime Type conditions.

These analyses revealed a pattern similar to that found in Experiments 1-2, in which pseudowords prompted more spontaneous looking at Related versus Unrelated pictures. In the first time bin (0-499ms) only the Direct Prime Type influenced looking behavior: Analyses revealed no main effect of Relatedness or Prime Type, but did reveal an interaction between these factors (Direct Co-Occur 89% HDI = [1.16, 37.28], Shared Co-Occur 89% HDI = [-33.86, 2.77]). Critically,

all subsequent time bins (500ms onward) consistently revealed a main effect of Relatedness that indicated greater looking at the Related picture for both Prime Types. In the final bin (2,500-3,000ms) only, an interaction with Prime Type indicated that this tendency was stronger in the Direct Prime Type condition (Direct Prime 89% HDI = [66.54, 107.18], Shared Prime 89% HDI = [33.85, 73.78]).

The presented analyses reveal that participants spontaneously looked more at pictures of familiar words upon hearing pseudoword primes with which they directly co-occurred or shared co-occurrence. Effects of both direct co-occurrence and shared co-occurrence emerged relatively rapidly.

Discussion

The results of Experiment 3 show that co-occurrence-driven semantic integration is evident in spontaneous looking behavior. Although participants did not perform any task on trials in which semantic integration was assessed, pseudowords nevertheless prompted spontaneous looking at pictures of the familiar nouns with which they directly co-occurred or shared co-occurrence. The timing with which pseudowords prompted looking is similar to both: (1) the timing with which familiar words prompted looking at pictures of semantically related familiar items (e.g., looking at an apple upon hearing “tree” or “grapes”) in a variant of the present paradigm (Experiment 3 in Unger et al., 2020), and (2) the timing of looking at referents of newly learned words whose meanings were learned through different routes (Trueswell, Medina, Hafri, & Gleitman, 2013). These findings present further evidence that co-occurrence regularities drive semantic integration

General Discussion

As we learn new words, they become organized in our mental lexicon according to meaningful semantic links. The present experiments highlight a potentially powerful human ability to build these interconnected semantic networks via sensitivity to simple co-occurrence regularities of word use in language.

Across three experiments, we observed that new words became semantically integrated with familiar words not only when they reliably directly co-occurred, but also when they never directly co-occurred, and instead only shared patterns of direct co-occurrence with other words. These links formed rapidly within a 10-15-minute training session and were strong enough to be detected from both a simple semantic priming task and spontaneous looking behavior. Critically, we detected the formation of novel semantic links even without instructing participants to reason about the links between the words during either the training or the assessment of semantic links (cf. Jones & Love, 2007; McNeill, 1963, 1966). These findings suggest that sensitivity to co-occurrence statistics can support the expansion of semantic networks to include new words. This sensitivity represents a potentially powerful and broadly useful route for semantic integration because it does not require explicitly learning definitions (Clay et al., 2007; Dagenbach et al., 1990; Tamminen & Gaskell, 2013), observing referents of new words (Breitenstein et al., 2007; Dobel et al., 2010), or encountering new words in contexts that provide strong cues to their meanings (Borovsky et al., 2012; Mestres-Missé et al., 2006).

Human Sensitivity to Co-Occurrence Statistics

Previous computational work has demonstrated that the two types of co-occurrence statistics that we manipulated in our experiments (i.e. direct and shared co-occurrence) are ubiquitous in everyday language, and capture much of the semantic links that connect human vocabularies (Asr et al., 2016; Frermann & Lapata, 2015; Hofmann et al., 2018; Jones & Mewhort, 2007; Jones et al., 2015; Landauer & Dumais, 1997; Miller & Charles, 1991; Rohde et al., 2004; Spence & Owens, 1990). Our work contributes to this literature by providing evidence that not only computational models but also humans can rely on these kinds of regularities to form semantic links between words.

In addition, the present findings expand our understanding of the cognitive domains that are shaped by statistical regularities in the environment. First, our observation that direct co-occurrence contributes to the semantic integration of new words adds to extensive prior evidence

that direct co-occurrence contributes to multiple domains, including speech sounds (Saffran et al., 1996), visual objects (Fiser & Aslin, 2002), and semantic reasoning (Fisher, Matlen, & Godwin, 2011; Matlen et al., 2015). Importantly, the contribution of shared co-occurrence adds to a smaller body of prior evidence for shared co-occurrence sensitivity in other domains, such as the formation of episodic memories (Hall, Mitchell, Graham, & Lavis, 2003; Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013; Zeithamova et al., 2012). Together, the present findings build upon this prior evidence by demonstrating the importance of statistical regularities of word use in building the semantic organization of word knowledge.

Exposure to Co-Occurrence Statistics

In contrast with prior studies of other routes to semantic integration that have provided extensive training over days and weeks (e.g., Breitenstein et al., 2007; Clay et al., 2007; Dobel et al., 2010; Mestres-Missé et al., 2006; Tamminen & Gaskell, 2013), we used an approach common in statistical learning studies in which we investigated whether even brief exposure to input containing a constrained set of regularities led to learning. Thus, as in other studies of sensitivity to statistical regularities, the exposure we provided was a condensed form of the exposure that people typically experience spread out over time and mingled with other input in real-world settings outside the lab. Our study thus demonstrates that short, condensed exposure to co-occurrence regularities can promote semantic integration.

Future research could investigate how co-occurrence-driven semantic integration might unfold over the course of exposure over time. Specifically, future research might investigate how rapidly words start to become semantically integrated based on co-occurrence, and whether co-occurrence-driven semantic integration remains (or is even more) robust when exposure is spaced over time. Such future research might further investigate how the timescale over which exposure occurs affects the nature of semantic integration. Specifically, the semantic integration observed in the present study following short, condensed exposure might have relied on the formation of episodic memory traces that only link the specific words experienced in the

experiment. From this perspective, novel words might only become integrated into an individual's broader semantic network with repeated exposure to consistent co-occurrence regularities over time (Coane & Balota, 2010; McClelland, McNaughton, & O'Reilly, 1995). Alternatively, there is evidence that new experiences that link to an existing network of semantic knowledge can become rapidly incorporated into it (Kumaran, Hassabis, & McClelland, 2016; Tse et al., 2007; van Kesteren et al., 2013; Van Kesteren, Fernández, Norris, & Hermans, 2010). Thus, new words might be rapidly incorporated into broader semantic networks when they co-occur with familiar words such as "apple" that are themselves interconnected in existing semantic networks. In contrast, new items may be more challenging to integrate when they cannot be incorporated into an existing semantic network (Frigo & McDonald, 1998; Ouyang et al., 2017; Yim, Savic, Unger, Sloutsky, & Dennis, 2019). Future research might investigate these possibilities by manipulating the time course of exposure to co-occurrence regularities and the availability of existing semantic networks while measuring the degree to which novel words become stably integrated over time.

Distinction from Other Drivers of Semantic Integration

It is particularly noteworthy that in the current study participants semantically integrated novel words based exclusively on exposure to co-occurrence regularities in language (Figure 2). The current study therefore demonstrates a route for semantic integration that is distinct from other routes for semantic integration studied in prior research.

First, given that it only requires *exposure* to everyday language, semantic integration based on co-occurrence regularities is clearly distinct from the routes that involve explicit learning, such as when new words are learned from studying definitions (Clay et al., 2007; Dagenbach et al., 1990; Tamminen & Gaskell, 2013) or from repeated exposure to new words paired with their referents (Breitenstein et al., 2007; Dobel et al., 2010).

Importantly, the co-occurrence route is also distinct from semantic integration based on exposure to language input that strongly constrains the meaning of novel words (Borovsky et al., 2012; Mestres-Missé et al., 2006). For example, after hearing the sentence "It was a windy day,

so Peter went to the park to fly a *dax*”, the novel word *dax* can become linked to a familiar word *kite*. However, for this link to be formed, the context words (e.g., *windy*, *park*, and *fly*) need to activate strong expectations of the word *kite* when *dax* is encountered. Therefore, this form of semantic integration requires both (1) highly constraining input and (2) the activation of strong expectations from pre-existing semantic links. In contrast, the present experiments deliberately focused on semantic integration that (1) relies on simple regularities that are abundant in everyday language and (2) does not require pre-existing semantic links to activate expectations. The present findings thus highlight a unique contribution of a simple but powerful mechanism of semantic integration that relies on exposure to co-occurrence regularities in language.

Potential Contributions to Semantic Development

Semantic networks of words begin to form early in development, just as early word learning is taking place (Bergelson & Aslin, 2017; Willits, Wojcik, Seidenberg, & Saffran, 2013). The nature of the learning processes that drive the formation of these networks has been a subject of research for decades, in which the focus has been to explain how we learn similarity in meaning (often referred to as *taxonomic or paradigmatic relatedness*). Similarity in meaning has typically been treated as distinct from *associative* links (also termed *syntagmatic, thematic, or schematic*) that are likely derived from co-occurrence between words in language or entities in the environment.

The present experiments may offer a key insight into how semantic networks of words emerge and develop. Specifically, sensitivity to the regularities with which words co-occur may give rise not only to associative links (through direct co-occurrence), but also to similarity in meaning (through shared co-occurrence). This possibility is supported by three lines of evidence. First, recent modeling studies have shown that infant and child language input, like adult language input, is rich in direct and shared co-occurrence regularities that capture semantic links between words (Asr et al., 2016; Frermann & Lapata, 2015; Huebner & Willits, 2018). Second, abilities to link inputs (such as words) based on direct co-occurrence emerge early in development (e.g.,

Wojcik & Saffran, 2015), and may be gradually supplemented by abilities to form links based on shared co-occurrence with development (e.g., Bauer & Larkina, 2017; Schlichting, Guarino, Schapiro, Turk-Browne, & Preston, 2017). Finally, recent evidence (Unger et al., 2020) suggests that developmental changes in semantic organization are well captured by mechanistic account (Sloutsky et al., 2017) in which such changes are driven by the formation of direct and shared co-occurrence-based semantic links. The present results further support this route for semantic development by demonstrating that exposure to co-occurrence regularities can contribute to building semantic networks of words.

Questions for Future Research: Mechanisms

One key avenue for future investigation is to uncover the mechanisms responsible for co-occurrence supported semantic integration of new words. One possibility is that associative learning may be sufficient for the formation of direct and shared co-occurrence-based links. Specifically, words that directly co-occur in language may become linked via associative learning based on contiguity (e.g., Jenkins, 1963; Pearce & Bouton, 2001). Consequently, words that share patterns of direct co-occurrence may become linked *indirectly*, via activation of direct co-occurrence-based links. For example, upon forming associative links based on observed direct co-occurrence such as *foobly-apple* and *foobly-mipp*, activation of *mipp* can activate *foobly*, which in turn activates *apple*. This possibility has been investigated in the work on mediated associations (Hall et al., 2003; Jenkins, 1963), where the shared element of associated pairs (e.g. *foobly*) is seen as a mediator which supports the spread of activation in a chain (e.g., *mipp* – [*foobly*] – *apple*). While this argument has received support in multiple studies and thus should not be entirely dismissed, its power to account for more complex co-occurrence regularities has been challenged (Jenkins, 1963).

One alternative explanation to simple associative learning mechanism is that co-occurrence based semantic integration may rely on forming *predictions*. Specifically, direct co-occurrence of words may prompt the formation of predictions about what words may be encountered next in a

sentence. For example, in our experiment, after reading *foobly* a participant may predict that the next word will be either *apple* or *mipp*. Consequently, word *foobly* will sometimes prompt a prediction of the wrong word, such as predicting *apple* when *mipp* occurs instead. This process would cause the words that share each other's patterns of direct co-occurrence (such as *apple* and *mipp*) to become co-activated during exposure to regularities in input. This co-activation alone, or in conjunction with an error signal prompted by encountering a word that was not predicted (e.g., Ervin, 1961; see also Zeithamova et al., 2012), may then facilitate forming shared co-occurrence-based links.

While the prediction-based process may not be specific to co-occurrence based semantic integration, it is likely that language provides the context in which prediction plays a more important role than in learning co-occurrence of completely novel elements in a novel context or in isolation. This possibility provides cues to why studies testing seemingly similar co-occurrence structures failed to find priming effects for shared co-occurrence regularities (e.g., Yim et al., 2019). Arbitrating between these and other possibilities may merit a systematic comparison and test of the predictions of different mechanisms for learning direct, shared, and more complex co-occurrence regularities that have been proposed in a variety of fields including statistical learning, memory, distributional semantics and behaviorist research (e.g., Hall et al., 2003; Jenkins, 1963; Jones & Mewhort, 2007; Kumaran & McClelland, 2012; Landauer & Dumais, 1997; Mikolov, Chen, Corrado, & Dean, 2013; Schapiro, Turk-Browne, Botvinick, & Norman, 2017; Zeithamova et al., 2012).

Conclusion

The present experiments provided novel evidence of a simple but powerful mechanism for integrating new words into our mental lexicons. Specifically, learners in our experiments formed links between novel and familiar words from mere exposure to sentences in which these words co-occurred or shared each other's patterns of co-occurrence with another word. This finding

indicates that sensitivity to distributional co-occurrence regularities is a potentially key contributor to building semantic networks fundamental to our fluency with language.

Context

The present study builds upon extensive computational evidence that suggests that statistical regularities of word use in language capture much of the semantic links in human semantic networks. Here we investigated whether human learners can exploit these regularities to form semantic links between newly learned and familiar words. The findings that humans did indeed link new with familiar words based on simple statistical regularities when no additional cues to word meaning were available suggest that this sensitivity represents a simple but powerful route for building semantic networks of words. Moreover, the present findings expand our understanding of the cognitive domains that are shaped by statistical regularities in the environment. This study is a part of a larger project investigating the role of statistical regularities in language in both semantic development, and the development of other abilities that depend on semantic organization, such as language comprehension.

References

- Altmann, G. T., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*, 73(3), 247-264.
- Arnon, I., & Snider, N. (2010). More than words: Frequency effects for multi-word phrases. *Journal of Memory and Language*, 62(1), 67-82.
- Asr, F. T., Willits, J. A., & Jones, M. N. (2016). Comparing predictive and co-occurrence based models of lexical semantics trained on child-directed speech. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 38.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67, 1-48.
- Bauer, P. J., & Larkina, M. (2017). Realizing relevance: The influence of domain-specific information on generation of new knowledge through integration in 4-to 8-year-old children. *Child development*, 88(1), 247-262.
- Bergelson, E., & Aslin, R. N. (2017). Nature and origins of the lexicon in 6-mo-olds. *Proceedings of the National Academy of Sciences*, 114(49), 12916-12921.
- Bonnotte, I., & Casalis, S. (2010). Semantic priming in French children with varying comprehension skills. *European Journal of Developmental Psychology*, 7(3), 309-328.
- Borovsky, A., Elman, J. L., & Kutas, M. (2012). Once is enough: N400 indexes semantic integration of novel word meanings from a single exposure in context. *Language Learning and Development*, 8(3), 278-302.
- Breitenstein, C., Zwitserlood, P., de Vries, M. H., Feldhues, C., Knecht, S., & Dobel, C. (2007). Five days versus a lifetime: Intense associative vocabulary training generates lexically integrated words. *Restorative Neurology and Neuroscience*, 25(5, 6), 493-500.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, 121(1), 127-132.

- Chuang, Y.-Y., Vollmer, M. L., Shafaei-Bajestan, E., Gahl, S., Hendrix, P., & Baayen, R. H. (2020). The processing of pseudoword form and meaning in production and comprehension: A computational modeling approach using linear discriminative learning. *Behavior Research Methods*, 1-32.
- Clay, F., Bowers, J. S., Davis, C. J., & Hanley, D. A. (2007). Teaching adults new words: the role of practice and consolidation. *Journal of experimental psychology: Learning, Memory, and Cognition*, 33(5), 970-976.
- Coane, J. H., & Balota, D. A. (2010). Face (and nose) priming for book. *Experimental Psychology*.
- Coutanche, M. N., & Thompson-Schill, S. L. (2014). Fast mapping rapidly integrates information into existing memory networks. *Journal of Experimental Psychology: General*, 143(6), 2296-2303.
- Dagenbach, D., Horst, S., & Carr, T. H. (1990). Adding new information to semantic memory: How much learning is enough to produce automatic priming? *Journal of experimental psychology: Learning, Memory, and Cognition*, 16(4), 581-591.
- DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature neuroscience*, 8(8), 1117.
- Dobel, C., Junghöfer, M., Breitenstein, C., Klauke, B., Knecht, S., Pantev, C., & Zwitserlood, P. (2010). New names for known things: on the association of novel word forms with existing semantic information. *Journal of Cognitive Neuroscience*, 22(6), 1251-1261.
- Ervin, S. M. (1961). Changes with age in the verbal determinants of word-association. *The American journal of psychology*, 74(3), 361-372.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99(24), 15822-15826.

- Fisher, A. V., Matlen, B. J., & Godwin, K. E. (2011). Semantic similarity of labels and inductive generalization: Taking a second look. *Cognition*, 118(3), 432-438.
- Frermann, L., & Lapata, M. (2015). Incremental Bayesian Category Learning From Natural Language. *Cognitive science*, 40, 1333–1381.
- Frigo, L., & McDonald, J. L. (1998). Properties of phonological markers that affect the acquisition of gender-like subclasses. *Journal of Memory and Language*, 39(2), 218-245.
- Gerken, L., Wilson, R., & Lewis, W. (2005). Infants can use distributional cues to form syntactic categories. *Journal of Child Language*, 32(2), 249-268.
- Goodrich, B., Gabry, J., Ali, I., & Brilleman, S. (2020). rstanarm: Bayesian applied regression modeling via Stan (Version R package version 2.21.1). Retrieved from <http://mc-stan.org/>
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological review*, 114(2), 211-244.
- Hall, G., Mitchell, C., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: evidence for associative mediation. *Journal of Experimental Psychology: General*, 132(2), 266-276.
- Hare, M., Jones, M. N., Thomson, C., Kelly, S., & McRae, K. (2009). Activating event knowledge. *Cognition*, 111(2), 151-167.
- Harris, Z. S. (1954). Distributional structure. *Word*, 10(2-3), 146-162.
- Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, 7(4), 569-592.
- Hofmann, M. J., Biemann, C., Westbury, C., Murusidze, M., Conrad, M., & Jacobs, A. M. (2018). Simple Co-Occurrence Statistics Reproducibly Predict Association Ratings. *Cognitive science*, 42(7), 2287-2312.
- Huang, Y., & Snedeker, J. (2020). Evidence from the visual world paradigm raises questions about unaccusativity and growth curve analyses. *Cognition*, 200, 104251.

- Huebner, P. A., & Willits, J. A. (2018). Structured semantic knowledge can emerge automatically from predicting word sequences in child-directed speech. *Frontiers in Psychology*, 9.
- Huetting, F., & Mani, N. (2016). Is prediction necessary to understand language? Probably not. *Language, Cognition and Neuroscience*, 31(1), 19-31.
- Huetting, F., Rommers, J., & Meyer, A. S. (2011). Using the visual world paradigm to study language processing: A review and critical evaluation. *Acta Psychologica*, 137(2), 151-171.
- Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review*, 10(4), 785-813.
- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse, C.-S., . . . Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45(4), 1099-1114.
- Jenkins, J. J. (1963). *Mediated associations: Paradigms and situations*. Paper presented at the Conference on Verbal Learning and Verbal Behavior, 2nd, Jun, 1961, Ardsley-on-Hudson, NY, US.
- Johns, B. T., Dye, M., & Jones, M. N. (2016). The influence of contextual diversity on word learning. *Psychonomic Bulletin & Review*, 23(4), 1214-1220.
- Johns, B. T., Mewhort, D. J., & Jones, M. N. (2019). The role of negative information in distributional semantic learning. *Cognitive science*, 43(5), e12730.
- Jones, M., & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive psychology*, 55(3), 196-231.
- Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, 114(1), 1-37.
- Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. In J. Busemeyer & J. Townsend (Eds.), *Oxford Handbook of Mathematical and Computational Psychology* (pp. 232-254). New York, NY: Oxford University Press.

- Kamide, Y., Altmann, G. T., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, 49(1), 133-156.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35-B42.
- Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7), 512-534.
- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: a model of the hippocampal system. *Psychological review*, 119(3), 573-616.
- Kutas, M., & Federmeier, K. D. (2000). Electrophysiology reveals semantic memory use in language comprehension. *Trends in cognitive sciences*, 4(12), 463-470.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*, 82(13).
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2), 211-240.
- Lany, J., & Saffran, J. R. (2010). From statistics to meaning: Infants' acquisition of lexical categories. *Psychological Science*, 21(2), 284-291.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28(2), 203-208.
- Makowski, D., Ben-Shachar, M., & Lüdtke, D. (2019). bayestestR: Describing Effects and their Uncertainty, Existence and Significance within the Bayesian Framework. *Journal of Open Source Software*, 4(0), 1541.

- Mani, N., & Huettig, F. (2012). Prediction during language processing is a piece of cake—But only for skilled producers. *Journal of Experimental Psychology: Human Perception and Performance*, 38(4), 843.
- Matlen, B. J., Fisher, A. V., & Godwin, K. E. (2015). The influence of label co-occurrence and semantic similarity on children's inductive generalization. *Frontiers in Psychology*, 6, 1146.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3), 419.
- McDonald, S. A., & Shillcock, R. C. (2003). Eye movements reveal the on-line computation of lexical probabilities during reading. *Psychological Science*, 14(6), 648-652.
- McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan* (2nd ed.). Boca Raton, FL: CRC Press.
- McKoon, G., & Ratcliff, R. (1979). Priming in episodic and semantic memory. *Journal of verbal learning and verbal behavior*, 18(4), 463-480.
- McMurray, B. (2020). I'm not sure that curve means what you think it means: Toward a [more] realistic understanding of the role of eye-movement generation in the Visual World Paradigm. Preprint retrieved from <https://psyarxiv.com/pb2c6/>.
- McNamara, T. P. (2005). *Semantic priming: Perspectives from memory and word recognition*. New York: Psychology Press.
- McNeill, D. (1963). The origin of associations within the same grammatical class. *Journal of verbal learning and verbal behavior*, 2(3), 250-262.
- McNeill, D. (1966). A study of word association. *Journal of Memory and Language*, 5(6), 548-557.
- McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of experimental psychology: Learning, Memory, and Cognition*, 24(3), 558.

- McRae, K., Khalkhali, S., & Hare, M. (2012). Semantic and associative relations in adolescents and young adults: Examining a tenuous dichotomy. In V. F. Reyna, S. B. Chapman, M. R. Dougherty, & J. Confrey (Eds.), *The adolescent brain: Learning, reasoning, and decision making* (pp. 39-66): American Psychological Association.
- Mestres-Missé, A., Rodriguez-Fornells, A., & Münte, T. F. (2006). Watching the brain during meaning acquisition. *Cerebral Cortex*, 17(8), 1858-1866.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Miller, G. A., & Charles, W. G. (1991). Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6(1), 1-28.
- Mintz, T. H. (2002). Category induction from distributional cues in an artificial language. *Memory & Cognition*, 30(5), 678-686.
- Mirman, D., Dixon, J. A., & Magnuson, J. S. (2008). Statistical and computational models of the visual world paradigm: Growth curves and individual differences. *Journal of Memory and Language*, 59(4), 475-494.
- Nagy, W. E., Herman, P. A., & Anderson, R. C. (1985). Learning words from context. *Reading research quarterly*, 20, 233-253.
- Nation, K., & Snowling, M. J. (1999). Developmental differences in sensitivity to semantic relations among good and poor comprehenders: Evidence from semantic priming. *Cognition*, 70(1), B1-B13.
- Nation, P., & Waring, R. (1997). Vocabulary size, text coverage, and word lists. In N. Schmitt & M. McCarthy (Eds.), *Vocabulary: Description, Acquisition, and Pedagogy* (pp. 6-19). Cambridge, UK: Cambridge University Press.
- Neely, J. H. (2012). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Bresner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 272-344). Hillsdale, NJ: Lawrence Erlbaum.

- Nobre, A. d. P., & Salles, J. F. d. (2016). Lexical-semantic processing and reading: Relations between semantic priming, visual word recognition and reading comprehension. *Educational Psychology, 36*(4), 753-770.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological review, 97*(2), 185-200.
- Ouyang, L., Boroditsky, L., & Frank, M. C. (2017). Semantic coherence facilitates distributional learning. *Cognitive science, 41*, 855-884.
- Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual review of Psychology, 52*(1), 111-139. doi:10.1146/annurev.psych.52.1.111
- Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Statistical learning in a natural language by 8-month-old infants. *Child development, 80*(3), 674-685.
- Ramscar, M., Sun, C. C., Hendrix, P., & Baayen, H. (2017). The mismeasurement of mind: Life-span changes in paired-associate-learning scores reflect the “cost” of learning, not cognitive decline. *Psychological Science, 28*(8), 1171-1179.
- Roelke, A., Franke, N., Biemann, C., Radach, R., Jacobs, A. M., & Hofmann, M. J. (2018). A novel co-occurrence-based approach to predict pure associative and semantic priming. *Psychonomic Bulletin & Review, 25*(4), 1488-1493.
- Roembke, T. C., Wasserman, E. A., & McMurray, B. (2016). Learning in rich networks involves both positive and negative associations. *Journal of Experimental Psychology: General, 145*(8), 1062.
- Rohde, D. L., Gonnerman, L. M., & Plaut, D. C. (2004). An improved method for deriving word meaning from lexical co-occurrence. *Cognitive psychology, 7*, 573-605.
- Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*(6), 906-914.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*(5294), 1926-1928.

- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27-52.
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature neuroscience*, 16(4), 486-492.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160049.
- Schlichting, M. L., Guarino, K. F., Schapiro, A. C., Turk-Browne, N. B., & Preston, A. R. (2017). Hippocampal structure predicts statistical learning and associative inference abilities during development. *Journal of Cognitive Neuroscience*, 29(1), 37-51. Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5130612/pdf/nihms826993.pdf>
- Sloutsky, V. M., Yim, H., Yao, X., & Dennis, S. (2017). An associative account of the development of word learning. *Cognitive psychology*, 97, 1-30.
- Spence, D. P., & Owens, K. C. (1990). Lexical co-occurrence and association strength. *Journal of Psycholinguistic Research*, 19(5), 317-330.
- Tamminen, J., & Gaskell, M. G. (2013). Novel word integration in the mental lexicon: Evidence from unmasked and masked semantic priming. *The Quarterly Journal of Experimental Psychology*, 66(5), 1001-1025.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632-1634.
- Team, R. D. C. (2008). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing.

- Trueswell, J. C., Medina, T. N., Hafri, A., & Gleitman, L. R. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive psychology*, 66(1), 126-156.
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science*, 316(5821), 76-82.
- Unger, L., Savic, O., & Sloutsky, V. M. (2020). Statistical regularities shape semantic organization throughout development. *Cognition*, 198, 104190. doi:10.31234/osf.io/2ruyn
- van Kesteren, M. T., Beul, S. F., Takashima, A., Henson, R. N., Ruiter, D. J., & Fernández, G. (2013). Differential roles for medial prefrontal and medial temporal cortices in schema-dependent encoding: from congruent to incongruent. *Neuropsychologia*, 51(12), 2352-2359.
- Van Kesteren, M. T., Fernández, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences*, 107(16), 7550-7555.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11(1), 192-196.
- Willits, J. A., Amato, M. S., & MacDonald, M. C. (2015). Language knowledge and event knowledge in language use. *Cognitive psychology*, 78, 1-27.
- Willits, J. A., Wojcik, E. H., Seidenberg, M. S., & Saffran, J. R. (2013). Toddlers activate lexical semantic knowledge in the absence of visual referents: Evidence from auditory priming. *Infancy*, 18(6), 1053-1075.
- Wojcik, E. H., & Saffran, J. R. (2015). Toddlers encode similarities among novel words from meaningful sentences. *Cognition*, 138, 10-20.
- Yim, H., Savic, O., Unger, L., Sloutsky, V. M., & Dennis, S. (2019). *Can Paradigmatic Relations be Learned Implicitly?* Paper presented at the Annual Meeting of the Cognitive Science Society.

Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, 75(1), 168-179.

Zeithamova, D., & Preston, A. R. (2010). Flexible memories: differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *Journal of Neuroscience*, 30(44), 14676-14684.