



Published in final edited form as:

*Psychon Bull Rev.* 2009 October ; 16(5): 869–874. doi:10.3758/PBR.16.5.869.

## Looking for Meaning: Eye Movements Are Sensitive to Overlapping Semantic Features, Not Association

Eiling Yee, Eve Overton, and Sharon L. Thompson-Schill

Department of Psychology, University of Pennsylvania

### Abstract

Theories of semantic memory differ in the extent to which relationships among concepts are captured via associative or via semantic relatedness. We examine the contributions of these two factors using a “visual world” paradigm in which participants select the named object from a four-picture display. We control for semantic relatedness while manipulating associative strength by using the visual world paradigm’s analogue to presenting asymmetrically associated pairs in either their “forward” or “backward” associative direction (e.g., *ham-eggs* vs. *eggs-ham*). Semantically related objects were preferentially fixated regardless of the direction of presentation (and the effect size was unchanged by presentation direction). However, when pairs were associated but not semantically related (e.g., *iceberg-lettuce*), associated objects were not preferentially fixated in either direction. These findings lend support to theories in which semantic memory is organized according to semantic relatedness (e.g., distributed models) and suggest that association by itself has little effect on this organization.

### Introduction

Models of semantic memory differ in how relationships among concepts are captured. Associative relations play an important role in spreading activation models (e.g., Anderson, 1983; Collins & Loftus, 1975), whereas in distributed models, semantic feature overlap is critical (e.g., Masson, 1995). However, both types of model assume that when a concept is activated, related concepts also become partially active. The paradigm most commonly applied to testing these models is semantic priming. Two recent reviews of this literature (Lucas, 2000; Hutchison, 2003) concluded that automatic semantic priming does not *require* association. While this conclusion is consistent with distributed models, both reviews also found some role for association, suggesting it too plays a role in the organization of semantic memory.

Using semantic priming to test associative vs. semantic relations can be problematic, however. Because the prime and target are typically paired, participants may notice relationships and engage in strategies to perform the task (Neely, Keefe & Ross, 1989), limiting the conclusions that can be drawn about the organization of semantic memory. It is therefore necessary to take great methodological care to ensure that semantic priming is not due to controlled processing – especially when testing associative relations, as they tend to be more obvious during semantic priming.

Evaluating the role of association is further complicated by the fact that most associated concepts are semantically related. In the priming literature, the strength of “association”

between a pair of concepts describes the probability that one will call the other to mind, and although this variable reflects a variety of relationships (e.g., antonyms, synonyms, words that co-occur in language, category coordinates, etc.), most of these are semantic. Further, concepts that are strongly semantically related are usually strongly associatively related as well, making it difficult to test either associative or semantic relatedness in isolation. One way to circumvent this problem is to take advantage of a key difference between associative and semantic relationships: as typically defined, association can be asymmetrical but semantic relations cannot. For example, *ham* and *eggs* are asymmetrically associated because while the cue *ham* frequently elicits the response *eggs*, the cue *eggs* does not often elicit *ham*. At the same time, because *ham* and *eggs* share the same number of semantic features regardless of which concept is first activated, their semantic relationship is symmetrical.

In an exploration of the effects of semantic and associative relatedness on automatic semantic priming, Thompson-Schill, Kurtz & Gabrieli (1998) (henceforth TKG) exploited the asymmetry of association strength. They used asymmetrically associated pairs that were either semantically related (e.g., *sleet-snow*) or semantically unrelated (e.g., *lip-stick*), and presented them in both the forward (e.g., *lip-stick*) and backward (e.g., *stick-lip*) directions. In this way, they varied association strength within a pair while keeping the semantic relationship constant<sup>1</sup>. Associative theories predict that backward priming should be weaker than forward priming because activation should spread only in the direction of association. They also predict that associated pairs should exhibit priming in the absence of semantic relatedness. Distributed theories, on the other hand (unless supplemented with associative links, e.g., Plaut, 1995), predict equal priming in both directions, and no priming for pairs that are semantically unrelated, regardless of their association.

TKG's findings were consistent with distributed theories. Priming was equal in the forward and backward directions for semantically related pairs. Furthermore, they observed no priming for semantically unrelated pairs (in either direction). These findings suggest that association alone plays little, if any role in semantic memory. This conclusion contrasts with demonstrations that semantically related pairs that are unassociated produce less priming than those that are semantically related *and* associated (see Lucas 2000, for a review). Two possible reasons for this discrepancy are that controlled processing played a role in the prior studies, or that pairs that are semantically related *and* associated are generally more strongly semantically related than those that are not associated (see McRae & Boisvert, 1998). Another possibility has been raised, however: according to the University of South Florida (USF) free association norms (Nelson, McEvoy, & Schreiber, 1998), TKG's semantically unrelated pairs were not as strongly associated as their semantically related pairs (Hutchison, 2003). It therefore remains possible that the lack of priming in the semantically unrelated condition was due to weaker association.

In the current study we test the predictions of associative vs. distributed models of semantic memory using stimuli that are better controlled than TKG's. We also avoid some of the problems with semantic priming by using a new paradigm. A new paradigm is desirable because standard semantic priming tasks (lexical decision and naming) make minimal demands on semantic processing, meaning that weaker relationships may not be detected. Furthermore, the manipulation typically used to prevent controlled processing – a short SOA – limits the paradigm's ability to detect relationships that may become active slowly; perhaps non-semantic associative relationships have failed to emerge under automatic

---

<sup>1</sup>TKG also took several steps to ensure that processing was automatic rather than controlled, including using a short SOA and a naming task.

conditions (Shelton & Martin, 1992; TKG) because associative relations emerge more slowly than semantic ones.

In a standard “visual world” eyetracking paradigm, participants are presented with a multi-picture display and asked to click on one of the objects (the target). If one of the objects is semantically related to the target word, participants are more likely to fixate on this related object than on unrelated objects (e.g., Huettig & Altmann, 2005; Yee & Sedivy, 2006). For example, when instructed to click on a candle, one is more likely to fixate on a picture of a lightbulb than on unrelated objects. This “relatedness effect” appears to be a result of a match between the semantic attributes of the related object and the conceptual representation of the activated target word (cf. Altmann & Kamide, 2007).

The visual world paradigm can provide useful insight into semantic vs. associative relationships because the related object is one of three distracter objects and is never singled out. Hence participants are unlikely to notice the relationship between objects in the display and so to engage in controlled processing. Furthermore, mapping a spoken word onto a picture places demands on semantic processing that are not required in word pronunciation and lexical decision. Finally, because eye movements are sampled continuously over a long window, we are not limited to probing at a single SOA and hence may be able to detect effects that occur at different times.

## Methods

### Participants

Thirty male and female undergraduates from the University of Pennsylvania received course credit for participating. All were native speakers of English and had normal or corrected-to-normal vision and no reported hearing deficits.

### Apparatus

An EyeLink II head-mounted eye tracker monitored participants’ eye movements. Stimuli were presented with PsyScript (Bates & Oliveiro, 2003).

### Materials

**Stimulus selection and norming**—Using the USF free association norms, we selected stimulus pairs that were more strongly associated in one direction than the other. To assess semantic relatedness, a separate set of 30 participants rated pairs of clip-art images (from Rossion & Pourtois, 2004, or a commercial collection) from 1 to 7 on “how similar are these things”. Pairs rated lower than 3 were assigned to the “non-semantic” condition (mean=1.6, e.g., *iceberg-lettuce*); pairs rated higher than 3 were assigned to the “semantic” condition (mean=4.2, e.g., *ham-eggs*). Mean forward and backward association strength was equivalent across the two conditions (non-semantic forward=.15, backward=.00; semantic forward=.15, backward=.01). Target words were recorded by a female speaker (E.Y.) in a quiet room.

Two lists, each 179 trials, were created. In the non-semantic condition (22 trials), the target<sup>2</sup> was associatively but not semantically related to one of the other three objects in the display (Figure 1, left panel). Related pairs appeared in both their “forward” and “backward” associative directions: the target in the forward direction (e.g., *iceberg*) was the word that commonly elicited the associate. The target in the backward direction (e.g., *lettuce*)

<sup>2</sup>In the current study, the *target* refers to the one and only word uttered, whereas in the semantic priming literature, the *target* follows the prime and is used to gauge the prime’s activation, and in the word association task, the *target* is the response to the *cue*.

infrequently elicited the associate. The other two objects in the display were semantically, associatively, and phonologically unrelated to the related pair. One unrelated object's name was frequency-matched with that of the target, and the other with that of the related object. The design of the semantic condition (24 trials) was analogous, except that the pairs were semantically as well as associatively related (Figure 1, right panel). These pairs also appeared in their "forward" (e.g., target: *ham*) and "backward" (e.g., target: *eggs*) associative directions.

All pairs (non-semantic and semantic) appeared on each list, but half were presented in the forward direction, and half in the backward direction (counterbalanced between lists such that no participant saw any object more than once). Each list also included 16 "semantic distracter" trials in which two of the objects in the display were related to each other but not to the target. Therefore, the presence of a related pair could not be used to predict the target. There were 117 filler trials in which no objects in the display were related.

To determine whether visual similarity was comparable across conditions, an independent group of 24 participants provided visual similarity ratings (from 1 to 7) for all related and control pairs (i.e., the target and the word-frequency-matched unrelated object in the same display). In the non-semantic condition, visual similarity ratings did not differ between target-related and target-control pairs (forward=1.5, forward control=1.2, backward=1.4, backward control=1.4). In the semantic condition, however, ratings were higher for target-related pairs than for target-control pairs (forward=2.1, forward control=1.4, backward=2.4, backward control=1.4). (Since semantically related objects often have similar shapes, this difference is expected.) We therefore used these visual similarity ratings as covariates in analyzing the eyetracking data.

## Procedure

**Picture exposure**—To ensure that participants knew what the pictures were supposed to represent, they viewed each picture and its label immediately before the eyetracking experiment.

**Eyetracking**—Participants were seated approximately 18 inches from a touch-sensitive monitor. They were presented with a 3×3 array (each cell was 2×2 in) with four pictures, one in each corner (Figure 1). Two seconds after the display appeared, a sound file named an object in the display. After the participant selected a picture by touching it, the trial ended and the screen went blank. There were two practice trials.

Eye movements were recorded starting from when the array appeared and ending when the participant touched the screen. Only fixations initiated after the onset of the target word were included in our analyses. We defined four regions, each corresponding to a corner cell in the array. We defined a fixation as starting with the beginning of the saccade that moved into a region and ending with the beginning of the saccade that exited that region.

## Results & Discussion

Across all participants, seven trials (0.5%) were excluded because the wrong picture was selected. Nine percent of trials did not provide any data because there were no eye movements after target word onset (in most, the participant was fixating the target picture at word onset). For the remaining data, we computed the proportions (across trials) of fixations on each picture type (target, related, control) over time in 100-ms bins (see Figures 2 and 3). For each bin of each condition, fixation proportions more than 2.5 standard deviations from the mean were replaced with the mean for that bin of that condition.

For data analysis, we defined a trial as starting at 200 ms after target onset (approximately 180 ms is needed to initiate a saccade to a target in a task like this [Altmann & Kamide, 2004]), and ending at 1100 ms after target onset, the average time that the probability of fixating on the target reached asymptote. We binned data into nine bins (from 200-1100 ms after target onset) and submitted these bins to ANOVA, with condition (related or unrelated) and time as repeated measures<sup>3</sup>. In the item analyses, picture-based visual similarity ratings were included as a covariate.

In the non-semantic condition (Figure 2), there were no effects of relatedness in either the forward,  $F_1(1, 29)=1.1, p=.31, F_2(1, 20)=0.8, p=.38$ , or backward direction,  $F_1(1, 29)=1.4, p=.25, F_2(1, 20)=0.3, p=.61$ . In neither direction was there an interaction of relatedness with time. In the semantic condition (Figure 3), in contrast, there was a significant effect of relatedness in both the forward,  $F_1(1, 29)=9.7, p<.01, F_2(1, 22)=4.9, p=.04$ , and backward directions,  $F_1(1, 29)=15.7, p<.001, F_2(1, 22)=4.4, p=.05$ . Because picture-based visual similarity was co-varied out in the items analysis, this effect cannot be attributed to visual similarity. In both directions there was an interaction of relatedness with time that was significant by participants (forward,  $F_1(8, 29)=3.4, p=.01$ ; backward,  $F_1(8, 29)=2.8, p=.01$ ), but not by items (forward  $F_2(8, 22)=1.9, p=.12$ ; backward  $F_2(8, 22)=0.8, p=.54$ )<sup>4</sup>. The interaction is due to the fact that the relatedness effect did not emerge until after the first few time bins.

To address the critical question of whether direction of association influenced the size of the relatedness effect, we computed the average relatedness effect (fixations on related – control averaged across the entire trial) for each condition. In neither condition was there a difference between the forward and backward directions,  $F_1(1, 29)=2.3, p=.14, F_2(1, 20)=1.1, p=.31$ , and  $F_1(1, 29)=.43, p=.52, F_2(1, 22)=0.4, p=.53$ , for the non-semantic and semantic conditions, respectively. To test whether direction of presentation might have an influence for the most strongly associated pairs, we analyzed the 10 most strongly associated pairs in each condition (non-semantic forward=.25, backward=.00; semantic forward=.25, backward=.01). The effects were unchanged.

A post-test questionnaire indicated that most participants (22 of 30) did not notice any relationships between the objects in the displays. Removing the eight participants who reported noticing some relationships did not change the pattern of results. To further test for controlled processing, we examined responses in the first 25% of trials (as strategies would presumably require time to develop). The pattern in these trials was the same as that in the later trials.

Why did we observe no relatedness effect for the non-semantic pairs in their forward-associative direction? One possibility is that these relations are weakly (if at all) encoded in semantic memory, and instead are encoded lexically. If so, association effects in priming studies may occur because priming tasks typically emphasize lexical processing. The eyetracking task, in contrast, emphasizes semantic processing – a feature that makes it more suitable for probing semantic memory. Another possibility, however, is that pictorial depictions bias the focus of attention to some aspects or features of the concept at the expense of others that might be more relevant for a given association. Due to constraints on stimulus selection (that we return to below), both of these differences would be particularly relevant for the non-semantic condition.

<sup>3</sup>Main effects of time bin were obtained in all analyses. This is unsurprising because eye movements in early bins reflect the processing of only a small amount of acoustic input, while later eye movements converge on the target object. We therefore limit our discussion to main effects of condition and to condition by time interactions.

<sup>4</sup>It is likely that the interaction with time is not significant by items because of variability in the point at which different words can be identified.

To test whether, under conditions that increase demands on lexical processing, the pictures used in the eyetracking study would elicit non-semantic associative priming, 23 separate participants completed a picture naming task with these stimuli. After an exposure phase (see eyetracking procedure), participants were presented with a sequence of individual pictures to name. Pictures were preceded by either a related picture (e.g., *lettuce* preceded by *iceberg*, or vice versa, between subjects) or an unrelated picture (overall frequency and word length matched). Two lists were created so that no picture appeared twice. Response times were measured from the onset of the image to the onset of the name, recorded by a voice key. In the backwards direction, there was no priming in either the non-semantic or the semantic condition. In the forward direction, the priming effect in the non-semantic condition was significant  $t(21)=1.8, p=.04$ , and the effect in the semantic condition approached significance  $t(22)=1.7, p=.07$ . This priming suggests that the absence of non-semantic associative priming in the eyetracking study has more to do with the balance of demands on semantic vs. lexical processing than the use of pictures per se.

Our eyetracking findings, which revealed semantic but not associative relatedness effects, converge with those of TKG using an independent paradigm. Because we used the USF norms to match association across our non-semantic and semantic conditions, this work is not subject to the concern that the null result in the non-semantic condition could have been due to weaker association for these pairs. The current study also suggests that the absence of non-semantic, associative priming under automatic conditions in prior studies (Shelton & Martin, 1992; TKG) was not because associative relationships are slow to arise and therefore undetectable at short SOAs.

The effect that we observed for pairs that are related both semantically and associatively—i.e., a significant relatedness effect that, importantly, was unchanged by direction of presentation—is consistent with distributed models. Although it is possible to supplement distributed models with associative links (e.g., Plaut, 1995), the fundamental organizing principle of the distributed semantic architecture is that relatedness is represented via overlapping semantic features. Thus (in the absence of associative supplementation), concepts that are associated but semantically unrelated should not activate each other, but concepts that are associated and semantically related should activate each other equally, regardless of which concept is activated first.

On the other hand, the null result (in both directions) for associatively related but semantically unrelated pairs shows that association alone is not sufficient to produce a relatedness effect even in a sensitive paradigm like eyetracking, which allows us to monitor activation over a more extended period of time than other paradigms. Thus, our result suggests that when semantic processing is emphasized, any automatic conceptual activation of things that are associatively but not semantically related is weak at best. This finding casts doubt on theories of semantic memory that assume an important role for association because these theories predict that activation should spread to associated concepts regardless of whether they are semantically related.

It is important to recall, though, that associatively related concepts tend also to be semantically related. One exception is a subset of associated concepts that co-occur in language but are semantically unrelated. Our non-semantic condition was therefore drawn almost exclusively from these co-occurring pairs, and primarily comprised the individual components of compound nouns (e.g., *iceberg lettuce*, *ant hill*)<sup>5</sup>. These concepts are nevertheless associated according to norms, and hence, according to associative theories, should activate each other. Because we found no evidence of such activation under automatic conditions, our results suggest that these simple co-occurrence relations play little, if any role in the organization of semantic memory.



At first glance, the absence of a significant effect of co-occurrence may appear incompatible with a recent study showing that two corpus-based metrics of semantic similarity (LSA [Landauer & Dumais, 1997] and contextual similarity [McDonald, 2000]) can reliably predict relatedness effects in a similar eyetracking paradigm (Huettig, Quinlan, McDonald, & Altmann, 2006). However, because both metrics use distance in high-dimensional space to quantify similarity, they capture much more than simple co-occurrence. In fact, our finding that simple co-occurrence does not lead to a relatedness effect lends support to the possibility (raised by Huettig et al.) that LSA may be a weaker predictor of semantic relatedness effects than contextual similarity because LSA relies more heavily on textual co-occurrence.

Although useful for evaluating the role of association in semantic memory, it is important to keep in mind that the co-occurrence relationship represents an unusual case of an associative relationship that is not semantic. This leads to a point raised previously by McRae & Bosivert (1998) and others: given the heterogeneity of relationships that appear in association norms, associative relations should not be treated uniformly. A deeper understanding of the organization of semantic memory can be achieved by examining the individual semantic relationships that comprise associations.

## Conclusions

The current study indicates that association by itself does not modulate the size of a semantic relatedness effect. This finding is inconsistent with models of semantic memory that posit an important role for associative relations (e.g., spreading activation). The finding is consistent with distributed models of semantics because these models encode semantic relatedness via overlapping semantic features, and this overlap remains constant regardless of which concept is activated first.

## Acknowledgments

This research was supported by NIH Grant (R01MH70850) awarded to Sharon Thompson-Schill and by a Ruth L. Kirschstein NRSA Postdoctoral Fellowship awarded to Eiling Yee. We are grateful to Emily McDowell and Steven Frankland for assistance with data collection, and to Katherine White and Jesse Hochstadt helpful comments.

## References


- Altmann, GTM.; Kamide, Y. Now you see it, now you don't: Mediating the mapping between language and the visual world. In: Henderson, JM.; Ferreira, F., editors. *The interface of language, vision, and action: Eye movements and the visual world*. Psychology Press; New York: 2004. p. 347-386.
- Altmann GTM, Kamide Y. The real-time mediation of visual attention by language and world knowledge: Linking anticipatory (and other) eye movements to linguistic processing. *Journal of Memory and Language* 2007;57:502–518.10.1016/j.jml.2006.12.004
- Anderson JR. A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior* 1983;22:261–295.10.1016/S0022-5371(83)90201-3
- Bates TC, Oliveiro L. PsyScript: A Macintosh Application for Scripting Experiments. *Behavior Research Methods, Instruments, & Computers* 2003;35(4):565–576.

---


<sup>5</sup>Because our non-semantic associative pairs were primarily compound nouns, the absence of an effect in the forward condition contrasts with prior work showing a relatedness effect for objects semantically related to onset competitors of a heard word (Yee & Sedivy, 2006): the target word in the forward direction (e.g., *iceberg*) could be considered an onset competitor of the compound noun that comprises the target and competitor (e.g., *iceberg lettuce*), which is semantically related to the competitor (e.g., *lettuce*). We suspect that no such onset-mediated semantic effect was obtained in the current study because the compound words that comprised the “onset competitors” are low frequency and hence should be only weakly activated.

- Collins AM, Loftus EF. A spreading-activation theory of semantic processing. *Psychological Review* 1975;82:407–428.10.1037/0033-295X.82.6.407
- Huetting F, Altmann GTM. Word meaning and the control of eye fixation: semantic competitor effects and the visual world paradigm. *Cognition* 2005;96:B23–B32.10.1016/j.cognition.2004.10.003 [PubMed: 15833303]
- Huetting F, Quinlan PT, McDonald SA, Altmann GTM. Models of high-dimensional semantic space predict language-mediated eye movements in the visual world. *Acta Psychologica* 2006;121:65–80.10.1016/j.actpsy.2005.06.002 [PubMed: 16098943]
- Hutchison KA. Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review* 2003;10(4):785–813. [PubMed: 15000531]
- Landauer TK, Dumais ST. A solution to Plato's problem: the Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review* 1997;104:211–240.10.1037/0033-295X.104.2.211
- Lucus M. Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review* 2000;7(4):618–630. [PubMed: 11206202]
- Masson MEJ. A distributed model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1995;21(1):3–23.10.1037/0278-7393.21.1.3
- McDonald, SA. Environmental determinants of lexical processing effort. University of Edinburgh; Scotland: 2000. Unpublished doctoral dissertation
- McRae K, Boisvert S. Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory and Cognition* 1998;24:558–572.10.1037/0278-7393.24.3.558
- Nelson, DL.; McEvoy, CL.; Schreiber, TA. The University of South Florida word association, rhyme, and word fragment norms. 1998. From <http://www.usf.edu/Free-Association/>
- Neely JH, Keefe DE, Ross K. Semantic priming in the lexical decision task: Roles of prospective prime-generated expectancies and retrospective relation-checking. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1989;15:1003–1019.10.1037/0278-7393.15.6.1003
- Plaut, DC. Semantic and Associative Priming in a Distributed Attractor Network. *Proceedings of the 17th Annual Conference of the Cognitive Science Society; Pittsburgh, PA.* 1995. p. 37-42.
- Rossion B, Pourtois G. Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception* 2004;33:217–236.10.1068/p5117 [PubMed: 15109163]
- Shelton JR, Martin RC. How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory and Cognition* 1992;18:1191–1210.10.1037/0278-7393.18.6.1191
- Thompson-Schill SL, Kurtz KJ, Gabrieli DE. Effects of semantic and associative relatedness on automatic priming. *Journal of Memory and Language* 1998;38:440–458.10.1006/jmla.1997.2559
- Yee E, Sedivy JC. Eye movements to pictures reveal transient semantic activation during spoken word recognition. *Journal of Experimental Psychology, Learning, Memory, and Cognition* 2006;32(1): 1–14.10.1037/0278-7393.32.1.1






**Figure 1.** Sample displays from the non-semantic (left panel) and semantic (right panel) conditions. In the non-semantic condition the participant hears a target word (e.g., *iceberg*) which is associatively but not semantically related to one of the other objects in the display (e.g., the lettuce). In the semantic condition the target word (e.g., *ham*) is associatively and semantically related to one of the other objects in the display (e.g., the eggs). The other two objects are unrelated to the target or the related object. Each display is used in both the forward and the backward direction (between participants).



**Figure 2.** Mean proportions of trials across time containing a fixation on the target, related object and control object in the forward (left panel) and backward (right panel) directions of the non-semantic condition. Error bars represent standard error.



**Figure 3.** Mean proportions of trials across time containing a fixation on the target, related object and control object in the forward (left panel) and backward (right panel) directions of the semantic condition. Error bars represent standard error.