

The Goal-Dependent Nature of Automatic Semantic Priming

Lin Khern A. Chia (lachia2@illinois.edu)

Jon A. Willits (jwillits@illinois.edu)

Department of Psychology, University of Illinois at Urbana-Champaign, IL 61820 USA

Abstract

Despite the fact that priming is one of the most studied phenomena in cognitive psychology, many questions remain about exactly when, why and under what task conditions we ought to observe priming in the lab, and what types of relationships between words or concepts reliably lead to priming. This project contrasted two priming experiments where the primary manipulation was the decision the subjects were making about words (as well as manipulating other factors, like relatedness proportion, known to affect priming). We found evidence that: 1) automatic priming for semantically related words does happen under some conditions, but 2) semantic priming, and whether it happens independent of association, is dependent on the task in which participants are engaged. These results provide evidence for the context sensitive nature of the activation of semantic memory.

Keywords: Semantic memory; Semantic Priming; Associative Priming; Goals; Explicit Awareness

Introduction

Priming, or the improvement in performance in a perceptual or cognitive task relative to some baseline, is one of the most studied effects in cognition (McNamara, 2005). Much of this interest is because of priming's potential for giving us a window into our representations and how we access them. For example, if the word or concept *dog* is responded to more quickly when it is preceded by the word or concept *cat* than when it is preceded by *shoe*, it suggests that our representation of *dog* and *cat* share some relation, association, or overlap that *dog* and *shoe* do not (Collins & Loftus, 1975).

By discovering systematicities about what kinds of words or concepts prime each other, cognitive scientists hope to unravel the nature and structure of how knowledge is represented. For example, McRae, de Sa, & Seidenberg (1997) found different patterns of priming for human-made artifact and natural kind words, leading to claims about differences in the nature of the representations of those words' meanings. Statistically significant priming was only found between natural kind words if those words had high correlated feature overlap. In other words, priming was observed between words like *canary* and *robin*, which share a set of intercorrelated features that co-occur across a broad

range of words, like "*has wings*", "*has feathers*", and "*can fly*". But no priming was observed between words like *raspberry* and *ruby*, despite the fact that they superficially share many features in common (like "*is red*", "*is small*" and "*is round*"). Unlike "*has wings*", "*has feathers*", and "*can fly*", these features are not correlated across a broad set of items. McRae et al found that, in contrast to natural kind words, priming occurred for human-made artifact words that had high feature overlap, regardless of whether those features were correlated or uncorrelated. Based on these results, McRae et al. argued that correlated features are in some way important to the representational structure of natural kind concepts but not artifact concepts.

Decades of research has investigated a wide range of relationships between words, and whether those relationships lead to priming, including: normative association strength, co-occurrence in language, synonymy, antonymy, perceptual similarity, feature overlap, shared category membership, shared script/schema membership, functional relations, and others (for review, see Hutchison, 2003; & McNamara, 2005). However, conclusions about what types of relations systematically lead to priming are made difficult by the fact that many factors unrelated to the target-prime relationship also influence the extent to which semantic priming occurs. One such factor found by Moss et al. (1995) is that words belonging to the same category prime when presented auditorily, but not as text.

A second moderating factor is the type of task subjects are asked to perform during presentation of a target, can influence priming results. Examples of tasks are naming the target word aloud, or deciding whether the target is a legal string in the English language. The latter is called a lexical decision task, which we will abbreviate LDT.

A third moderating factor is the time duration between the prime and target (the stimulus onset asynchrony, or SOA). Hutchison (2003) reviewed 36 experiments (shown in Figure 1 below) examining priming for words belonging to the same category (and which were not normatively associated). He found that in experiments where the task was lexical decision, priming almost always occurred regardless of SOA, whereas in naming studies, priming effects were much less consistent.

A final moderating factor is the relatedness proportion (RP) of words in the study. In a typical study this can range from as high as 50% of the items being related, to sometimes being as low as 5% (Hutchison, 2003). Like SOA, RP effects can dramatically alter whether words of certain types prime each other. RP and SOA are often seen as working in a similar mechanistic fashion, by altering the extent to which the subject is explicitly aware of the potential for a relationship between the prime and target.




Figure 1. Priming effect sizes for same category words (in ms) in 36 experiments, as a function of task (lexical decision vs. naming) and short (<300 ms) or long (>300 ms) SOA. Data taken from Hutchison (2003) Table 4.

When subjects have a long time to process the prime, and a high proportion of items are related, there is a high chance they are making active, explicit, or strategic predictions about the items (Neely, 1991). In contrast, when the speed is very quick, or a small proportion of items are related, the chance of this is much reduced, and priming effects are often attributed to unconscious or automatic effects like spreading activation (Collins & Loftus, 1975).

With so many factors moderating or eliminating priming effects, we do not yet understand priming well enough to use it as a tool for probing which words’ or concepts’ representations are related. In this paper, we hope to bring some clarity to these issues. We do so by controlling and contrasting task, SOA, and RP within the same experiment and using the same items. This will allow us to see if semantic priming can be consistently obtained and the factors that affect these priming effects. In doing so, we hope to answer three primary questions.

First, can we reliably obtain priming for items that are semantically related (in terms of being from the same category), in the absence of other types of relations? For example, one other factor moderating the studies shown in Figure 1 was the extent to which the words were normatively associated (e.g. the prime reliably elicited the target in a free association task). Of the studies shown that failed to find semantic priming, the overwhelming majority used same category words that were not normatively

associated, whereas the studies that did find priming used words that were both associated and from the same category. This had led some to argue that most priming is “associative” priming, and that purely semantic (e.g. category or feature-based) priming are rare, weak, or nonexistent. This explanation is somewhat dissatisfying, however, because the fact that two words are associated in a normative task does not tell us much about the nature or cause of that association.

One possible explanation for the lack of priming without association deals with the strength of the similarity of the items. In some of the studies testing same-category priming for unassociated words, the category-based relationship was rather weak. For example, Shelton & Martin (1992)’s experiments 2 and 3 found priming times of 2 ms and -23 ms (in a long SOA LDT task), but many of their “related” items were of questionable relatedness, including *dirt* and *cement*, *soup* and *juice*, *barn* and *home*, and *duck* and *cow*. Thus, in this experiment the lack of associativity was confounded with the lack of strong semantic similarity (in terms of feature overlap or any other definition). Other studies, such as McRae and Boisvert (1998) that used more strongly related words, found evidence for semantic priming in the absence of association. In order to address the question of whether semantic priming exists independent of association, in our experiments we choose items that are from the same set of eight categories (mammals, birds, fruits, vegetables, vehicles, clothing, weapons, and tools), and were as highly similar as possible, but varied the degree of association so that its effect could be investigated statistically as a covariate in our analyses.

Second, is semantic category-based priming consistent across different relatedness proportions and stimulus onset asynchronies? As noted, there has been inconsistency in whether semantic priming is found with short SOAs or low RPs, leading some to suggest that semantic priming (as opposed to associative priming) is only an explicit or strategic phenomenon that occurs when subjects might be aware of the fact that words in the study are related, and that therefore automatic unconscious semantic priming does not occur. But again, many of these studies have problems, ranging from small sample sizes to relatively dissimilar “semantically related” words, to not fully crossing RP and SOA. In the experiments described below we ran different sets of subjects in a 2x2x2 design crossing extremely short SOAs (50ms) and moderately short SOAs (250ms), two RP conditions (0.25 and 0.50 related), in addition to whether the words’ meanings are related or unrelated (priming would or would not be expected).

A **third** question being tested in this paper is whether automatic priming is dependent on the task-related goals of the subject. Contrary to the depiction of semantic priming as a static phenomenon by a large majority of the literature, Willits et al. (2015) found that what types of

verb-instrument relations led to priming could be manipulated by changing the task. In tasks that had a linguistic bias (such as naming words aloud), priming was observed for words that have strong linguistic co-occurrence relationships, but not for words that were semantically related that do not co-occur frequently. In contrast, in tasks that were heavily semantic (such as making a category decision about the target word), priming occurred for those words that shared a semantic relationship, regardless of their linguistic co-occurrence probability. Across the current two experiments, we manipulated the task in a similar fashion, observing whether semantic priming is independent of the tasks-specific goals. In Experiment 1, subjects' task was to decide whether the target was a concrete (vs. abstract) entity. Unlike other tasks often used in semantic priming (like naming and lexical decision), this task is one that involves activation of semantic information, and thus may make semantic priming more likely. In Experiment 2, the subjects performed a semantic categorization decision, deciding if words belonged to a particular category (selected from the same eight categories from which the stimulus words were drawn). Critically, sometimes the related pairs were aligned with the category decision being made (e.g. *eagle-hawk* for "is the second word a bird" vs. "is the second word a vehicle). Thus, the contrast between Experiments 1 and 2 allows us to investigate the extent to which semantic priming is consistent across tasks, and whether or not it matters that the kind of semantic relationship being primed is consistent with the subject's current goal.

Experiment 1:

Priming in a Concrete/Abstract Decision Task

In Experiment 1, subjects saw a sequence of 128 prime-target pairs. They were asked to judge whether the target was a concrete real object (like a *rock*, *bird*, or *cloud*) or an abstract concept (like *truth*, *beauty*, or *honesty*). Half the items were concrete, and half the items were abstract. Among the concrete items, either 50% or 25% of the items were semantically related. Subjects were randomly assigned to each RP condition, and to either a 50 ms or 250 ms SOA condition. This resulted in a 2x2x2 mixed design, with RP and SOA as between-subject factors, and prime-target relatedness as a within-subject factor.

Method

Subjects. There were 339 undergraduate students who participated in the experiment for course credit. All subjects were fluent speakers of English.

Procedure. Subjects were seated in front of a computer screen. Each of 128 trials consisted of the presentation of the following sequence of events. First, a fixation cross for 50 ms. Second, the prime word (for either 25 ms or 225 ms, depending on SOA condition). Third, a pattern mask ("&&&&&&") for 25 ms (with the duration of the prime

word plus the pattern mask constituting the SOA). Fourth, the target word, which stayed on the screen until a response. The inter trial interval was one second. Subjects were required to answer yes or no as to whether the target word was a concrete real object. The trials were randomly divided into eight blocks of 16 words, allowing the subjects a brief resting period between each block.

Materials. Each subject saw 128 noun-noun trials which consisted of 64 concrete-abstract pairs and 64 concrete-concrete pairs. The specific 64 concrete-concrete trials varied across subjects depending on their RP condition. The 128 words making up the 64 concrete-concrete pairs were chosen according the following parameters. First, 16 words from eight semantic categories (mammals, birds, fruits, vegetables, vehicles, clothing, weapons, and tools) were chosen resulting in 64 pairs that were from the same category, maximized semantic similarity, while varying normative association strength.

The experiment's 64 related pairs were then arranged into counterbalanced lists that re-paired 50% or 75% of the targets with unrelated primes (depending on the RP condition). These lists also ensured that each prime and target occurred only once in each list, and that each word occurred as a related prime and target, and as an unrelated prime and target across different lists. For example, the RP=.50 condition had four lists, so that *dog* could occur as a related prime and target (*dog-cat*, and *cat-dog*) and an unrelated prime and target (*dog-shoe* and *shoe-dog*) across the four lists. Each subject saw only one list.

The 128 words making up the concrete-abstract trials were chosen by selecting 64 abstract words and then pairing each one with an unrelated concrete prime word (chosen equally distributed from the same eight categories). These same 64 concrete-abstract pairs were added to each of the lists described above. Note that this means that the RP conditions could be considered .25 and .125 rather than .5 and .25, depending on whether you are considering the relatedness of all trials, or of just the concrete-concrete *yes* trials which constituted our analyses.

Results and Discussion

As per standard convention in priming experiments, we first inspected accuracy scores to check to make sure there were no speed accuracy tradeoffs. Then we analysed the reaction times in the *yes* trials, after removing outlier trials that were shorter than 400 ms or greater than three standard deviations of the mean leaving 20,560 trials (out of 21,234 total *yes* trials) left for analyses. The resulting mean reaction times for related and unrelated trials in our four RP-by-SOA conditions are shown in Figure 2.

Next, we used relatedness, SOA, and RP as fixed factors predicting RT in a mixed-effects regression model, with subject and target word as random factors (Bates, Maechler, Bolker & Walker, 2015). The results of this model are shown in Table 1. We found significant main

effects of relatedness which did not interact with either SOA or RP. Thus we found evidence for a priming effect independent of SOA or RP, and no evidence that our




Figure 2. Mean reaction times (and standard deviations computed across subjects) in an abstract/concrete judgement task for related and unrelated trials as a function of relatedness proportion and stimulus onset asynchrony.

Table 1. Fixed effects of mixed effect model analyzing reaction time on *yes* trials in Experiment 1. As per convention, *t* scores of greater than 2 are typically considered statistically “significant” (Baayen, 2008).

Fixed Effect	<i>b</i>	<i>t</i>
Relatedness	-15.6	-5.39*
SOA	-26.7	-2.34*
RP	-1.20	-0.10
Relatedness x SOA	-2.92	-0.51
Relatedness x RP	5.74	0.99
SOA x RP	7.58	0.33
Relatedness x SOA x RP	-21.6	-1.87

Table 2. Fixed effects predicting residual variance in RT after removing variance in RT predictable by normative association strength in Experiment 1.

Fixed Effect	<i>b</i>	<i>t</i>
Relatedness	-8.22	-1.48
SOA	-21.9	-1.73
RP	-2.76	-0.22
Relatedness x SOA	3.97	0.36
Relatedness x RP	12.4	1.11
SOA x RP	-1.06	-0.04
Relatedness x SOA x RP	-38.0	-1.71

subjects were generating strategic expectations about prime-target relationships, even in SOA/RP conditions that encouraged such expectations.

We also fit a second model to the RT data after removing the variance in RT that could be predicted by association strength. This removal of variance was done by excluding the 15,175 trials that: 1) involved normatively associated prime-target pairs, and 2) included targets shared

by the normatively associated prime-target pairs so as to ensure equal treatment, leaving 5385 trials for analysis. The effect of relatedness disappeared after removing variance in RT predictable by normative association strength. Thus, in an abstractness judgement task, although we found evidence for an RP and SOA-independent priming effect, we did not find evidence for priming due to “semantic” relatedness (i.e. high similarity items belonging to the same category), when the effect of normative association was removed. This is true even though our items were picked to maximize strength of the relationship between the related prime and target items.

Experiment 2: Priming in a Category Decision Task

Experiment 2 was designed to investigate the extent to which the priming results found in Experiment 1 were dependent on the task in which the subject was engaged. In Experiment 2, subjects performed a category decision task, deciding if the target word belonged to a specific category (one of the same eight from which the stimuli were drawn).

Method

Subjects. There were 253 undergraduate students who participated in the experiment for course credit. All subjects were fluent speakers of English.

Materials. Items here were identical to those of experiment 1 but for one difference: the *no* trials were, like the *yes* trials, concrete-concrete pairs drawn from the same eight categories. These *no* trials were chosen such that their relatedness proportion matched that of the subject’s condition (either 0.25 or 0.50). Thus, each block consisted of related pairs that aligned with the category decision relevant for that block, unrelated pairs with either prime or targets (but not both) aligned with the category, as well as related trials that were misaligned with the category. As an example, consider Table 3, with a sample showing two counterbalanced lists of eight items.

Table 3. Sample items demonstrating related and unrelated pairings in Experiment 2 when task was to decide “Is the second word a mammal?” and RP=0.50. For this sample of words, other counterbalancing lists would have been created allowing all words to serve as both primes and targets in both related and unrelated trials, across different subjects.

Prime	Target	Condition	Correct Response
dog	cat	Related	Yes
rat	mouse	Related	Yes
eagle	deer	Unrelated	Yes
hammer	cow	Unrelated	Yes
sword	knife	Related	No
subway	train	Related	No
moose	shirt	Unrelated	No
zebra	blueberry	Unrelated	No

Procedure. The procedure in Experiment 2 was identical to that of Experiment 1 except for the nature of the yes-no decision, now a category decision. The 128 trials were divided into eight blocks, such that the specific category about which the subject was evaluating the word changed every 16 trials. These eight categories were the same from which the items were drawn. The order of the eight blocks was randomized across subjects.

Results and Discussion

Data in Experiment 2 were analyzed the same way as we analyzed the data in Experiment 1. 14,977 trials (out of 17,092 total *yes* trials) were left after our trimming process. The mean reaction times in each condition are shown in Figure 3.




Figure 3. Mean reaction times (and standard deviations computed across subjects) in a category judgement task.

Again, we used relatedness, SOA, and RP as fixed factors predicting RT in a mixed-effects regression model, with subject and target word as random factors. The results of this model are shown in Table 3. We found a significant main effect of relatedness, and also a significant interaction of relatedness with RP (but not SOA).

Table 4. Fixed effects of mixed effect model analyzing reaction time on *yes* trials in Experiment 2.

Fixed Effect	b	t
Relatedness	-22.1	-4.95
SOA	-42.6	-2.95
RP	-3.15	-0.22
Relatedness x SOA	4.76	0.69
Relatedness x RP	-19.3	-2.77
SOA x RP	67.8	2.35
Relatedness x SOA x RP	-4.35	-0.31

As per Experiment 1, we also fit a second model to the RT data after removing the variance in RT that could be predicted by association strength. This removal of variance was done by excluding the 10,563 trials that 1) involved normatively associated prime-target pairs, and 2) included targets shared by the normatively associated prime-target

pairs so as to ensure equal treatment, leaving 4414 trials for analysis.

The results of Experiment 2 turned out interestingly different than those of Experiment 1. Here, the main effect

Table 5. Fixed effects predicting residual variance in RT after removing variance in RT predictable by normative association strength in Experiment 2.

Fixed Effect	b	t
Relatedness	-33.4	-3.05
SOA	-40.2	-2.40
RP	9.36	0.56
Relatedness x SOA	-2.23	-0.17
Relatedness x RP	-2.10	-0.16
SOA x RP	80.8	2.41
Relatedness x SOA x RP	0.63	0.02

of relatedness survived the removal of normatively associated word pairs. Thus, the priming observed in Experiment 2 was at least partly due to semantic relatedness, independent of association. This stands in sharp contrast to the way that the priming observed in Experiment 1 can be attributable to effects of normative association.

Why the difference between Experiments 1 and 2? In comparison to an abstract/concrete judgement decision, semantic information (in particular, semantic similarity, overlapping semantic features, or the category to which a word belongs) is clearly more relevant when the task in which the subject is engaged is a category judgment decision. The results of Experiment 2, and their contrast with Experiment 1, strongly suggested that the manifestation of semantic priming depends on the goals or task in which a person is engaged. If their goals beg heavy use of knowledge about semantic features, empirical phenomena of cognitive access like priming should be organized semantically.

There was a significant Relatedness x RP interaction found in the model that included normatively associated pairs. However, this interaction disappeared in the model that excluded normatively associated pairs. This is a curious finding. Alone, these results might suggest that RP effects are selective and only relevant to associative priming. Unfortunately, this conclusion is untenable. Experiment 1, where priming was associative in nature, showed no RP effect at all. We are still left with some uncertainty about the exact role that RP plays in priming.

General Discussion

“Priming is an improvement in performance in a perceptual or cognitive task, relative to an appropriate baseline, produced by context or prior experience.” (McNamara, 2005). Priming is typically called semantic when the improvement in performance is brought about by prior experience with semantically related concepts. Due to the fleeting nature of semantic priming, some have expressed doubts that it reflects the true organization of

concepts in our mind. If semantic priming only manifested itself in conditions that encouraged strategic processing, no researcher would be able to use it as evidence that semantic memory is organized semantically. Given further evidence that semantic priming is predictable by word association norms, it might even be reasonable to say that semantic memory is instead organized associatively. However, this is not the conclusion warranted by our data. In Experiment 2, we found evidence for automatic priming (i.e. priming even with very short SOAs and low RPs) for words with no associative relationship.

Despite these findings, it should be stressed that automatic semantic priming is nonetheless a fleeting phenomenon. Consistent with other work about the task and/or goal dependent nature of semantic activation (Willits et al., 2015), we found that priming does not occur independent of task, or with words of low relatedness (Shelton & Martin, 1992). Our results indicate that semantic priming should only manifest reliably when subjects' goals involve heavy usage of semantic information. Subjects' task in Experiment 2 was heavily reliant on semantic information, where they were required to make decisions about the words' membership to categories that were directly aligned to the related vs. unrelated contrasts in stimuli. Experiment 1 on the other hand, required a lighter use of semantic information where nuanced distinctions between and matches of sets of features were not needed. Instead, all that was needed was whether or not the word referred to something that is tangible.

Given the ubiquity of SOA effects in the priming literature, it might be remarkable that our experiments showed no effects of SOA on priming at all. It is worth noting, though, that one of our limitations lie in our SOA manipulations: they were relatively small (50 ms vs. 250 ms) compared to some previous work (which has investigated SOAs as long as 1000 ms). While many have argued that 250 ms is where strategic effects begin to appear, it lies too close to the borderline for us to be certain of any conclusions drawn about pure-SOA effects. Future work could extend these studies to using much longer SOAs, resolving this uncertainty about strategic effects.

Other future directions include investigating the true nature of associative priming, a phenomenon that, because it is defined by a word norm task, is unsatisfying as a mechanistic explanation for priming. An alternative would be to ground associative priming in something more tangible as a mechanistic explanation. Willits et al. (2015), for example, used language co-occurrence statistics to fruitfully predict priming results. Corpora analyses therefore offers a step away from defining the phenomenon by a word norm task. Finally, given our promising results of the existence of semantic priming under at least some circumstances, it is natural that another next steps would be a computational model that is able to predict priming at the level of individual words and/or subjects. Such a model

would be a major step towards a truly mechanistic and unambiguous account of the source of lexical priming effects.

References

- Baayen, R. H. (2008) *Analyzing Linguistic Data. A Practical Introduction to Statistics Using R*. Cambridge, UK: Cambridge University Press.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67, 1–48.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.
- Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review*, 10, 785-813.
- McNamara, T. P. (2005). *Semantic Priming: Perspectives from Memory and Word Recognition*. New York: Psychology Press.
- McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 558-572.
- McRae, K., De Sa, V. R., & Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126, 99-130.
- Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 863-883.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and the, ones. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-337). Hillsdale, NJ: Erlbaum.
- Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming?. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1191-1210.
- Willits, J. A., Amato, M. S., & MacDonald, M. C. (2015). Language knowledge and event knowledge in language use. *Cognitive Psychology*, 78, 1-27.