Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Memory and Language 55 (2006) 89–101
Journal of Memory and Language
www.elsevier.com/locate/jml

Priming via relational similarity: A COPPER HORSE is faster when seen through a GLASS EYE

Zachary Estes [a,*], Lara L. Jones [b]

[a] *Department of Psychology, University of Warwick, Coventry CV4 7AL, UK*
[b] *University of Georgia, USA*

Received 23 June 2005; revision received 20 January 2006
Available online 15 March 2006

Abstract

Relation priming is a phenomenon in which comprehension of a word pair (e.g., COPPER HORSE) is facilitated by the prior presentation of another word pair (e.g., GLASS EYE) that instantiates the same conceptual relation (i.e., *composed of*). We investigated whether relation priming is contingent on lexical similarity. Study 1 revealed that relational similarity, but not lexical similarity, reliably predicted noun phrase comprehension across several previously published experiments. Study 2 demonstrated relation priming between lexically dissimilar phrases (e.g., STEEL SCISSORS → STRAW HAT). Thus, across both studies, lexical similarity failed to explain relation priming. Rather, comprehension of a target phrase was a function of its relational similarity to the prime phrase. Results are inconsistent with models in which conceptual relations are bound to the particular concepts that instantiate them, and suggest instead that conceptual relations are independent representational units that can be utilized by various and dissimilar concepts.

Keywords: Conceptual combination; Relation priming; Semantic similarity; Relational similarity; Lexical similarity; Thematic relations

Semantic priming is among the most well-documented phenomena in cognitive psychology (e.g., Meyer & Schvaneveldt, 1971; for review, see Hutchison, 2003; Lucas, 2000). Relation priming, on the contrary, is a topic of current dispute. At contention is whether comprehension of a word pair (e.g., COPPER HORSE) can be facilitated by the prior presentation of another word pair (e.g., GLASS EYE) that instantiates the same conceptual relation (i.e., Y *composed of* X). Although several researchers have demonstrated such relation priming (Estes, 2003; Gerrig & Murphy, 1992; Spellman, Holyoak, & Morrison, 2001), others attribute this effect to semantic priming (Gagné, Spalding, & Ji, 2005; see also Gagné, 2001). At the heart of this empirical dispute is a deeper theoretical debate concerning the nature of relational representation.

In this article, we consider the evidence of relation priming in the absence of semantic priming. We introduce *relational similarity* as a critical determinant of relation priming. By "relational similarity" we mean the extent to which the relation instantiated by one phrase is similar to the relation instantiated by another phrase. Relational similarity is contrasted here from "lexical similarity," by which we mean the semantic similarity of the individual lexical concepts. Thus, lexical similarity refers to words (e.g., COPPER and GLASS), while

* Corresponding author.
E-mail address: z.estes@warwick.ac.uk (Z. Estes).

0749-596X/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jml.2006.01.004

relational similarity refers to word pairs (e.g., COPPER HORSE and GLASS EYE). After introducing two models of relational representation, we review the extant investigations of relation priming. Then we re-analyze some of the key data to test whether relational similarity does indeed predict comprehension above and beyond any influence of lexical similarity. Finally, we report a novel experiment that demonstrates relation priming in the absence of lexical similarity.

Models of relational representation

Concepts may be related in many ways, such as causally (e.g., ROPE BURN), temporally (e.g., WINTER HOLIDAY), spatially (e.g., TABLE VASE), compositionally (e.g., GLASS EYE), and so forth. The issue of present interest is whether such conceptual relations are represented as independent units in the semantic network, or whether these relations are represented as part of the particular concepts that instantiate them. According to a model of *bound representation* (henceforth the "bound model"), a relation is represented as part of the meaning of whatever concepts entail that relation. For instance, the *composition* relation of COPPER HORSE is represented as part of the meaning of COPPER. Gagné illustrates this position, arguing that "relations are associated with the modifier's representation, rather than existing as independent structures" (2001, p. 247). Alternatively, according to a model of *independent representation* (henceforth the "independent model"), relations constitute representational structures in and of themselves (Estes, 2003; see also Spellman et al., 2001). So, for example, the *composition* relation is represented independent of any particular concept; it is not part of the representation of COPPER, or GOLD, or any given concept. It may be activated by those concepts, but it is not a part of their representation. That is, COPPER may activate the *composition* relation, but that relation may nevertheless constitute an independent representation in and of itself.

The bound and independent models differ dramatically in terms of representational demand. Consider the phrases BEAR PAW and FAN BLADE, both of which instantiate a *part/whole* relation. If relations are bound to their particular concepts of instantiation, then it follows that these concepts (e.g., BEAR, FAN) must represent the *part/whole* relation separately. In fact, every concept that consists of one or more parts must represent the *part/whole* relation separately. Similarly, just about every concept must include the *cause* relation in its representation, since just about any concept can be involved in a causal relation (e.g., LOTTERY RETIREMENT). Stated alternatively, for every concept in one's semantic network, its representation must include every relation that the concept could possibly instantiate (cf. Murphy, 2002, p. 463). So for example, consider just a brief list of relations instantiated by the concept BEAR: BEAR PAW (*part/whole*), BEAR SCARE (*causal*), BEAR SEASON (*temporal*), BEAR TOY (*possessive*), BEAR TRACKS (*from*), BEAR CAVE (*habitat*), BEAR CUB (*subtype*), BEAR FAMILY (*of*), BEAR STORY (*about*), BEAR PLAYGROUND (*for*), etc. Given the variety of different relations that may be instantiated, in conjunction with the number of concepts that may instantiate them, such redundancy would be extremely taxing in terms of representational demand. Moreover, the processing demands entailed by such redundancy would seem to be computationally intractable; as the number of representations increases, so does the time and effort required to search through those representations. In contrast, the independent model does not suffer from this problem of redundancy. Because relations are independent of any particular concept, each relation need be represented only once, and hence the representational demand is minimal.

Thus, theoretical considerations appear to favor the independent model over the bound model. Nonetheless, empirical tests are ultimately necessary to reject either model. We therefore turn to the relation priming paradigm, which can empirically contrast the bound and independent models of relational representation.

Relation priming

In the relation priming paradigm, a target phrase (e.g., COPPER HORSE) is preceded by a prime phrase that uses either the same relation (e.g., GLASS EYE; Y *composed of* X) or a different relation (e.g., GLASS CUT; Y *caused by* X). If the target is comprehended faster and/or more accurately following the same relation prime, then relation priming has occurred. The bound and independent models make opposing predictions with regard to relation priming. According to the bound model, relation priming should only occur if a concept is repeated from prime to target (e.g., GLASS EYE → GLASS HORSE), since relational representations are concept-bound (Gagné, 2001). In contrast, the independent model claims that the same relational representation (e.g., *composition*) is activated by any phrase that instantiates that relation, and hence relation priming may occur without lexical repetition (e.g., GLASS EYE → COPPER HORSE; Estes, 2003). So this issue of relation priming critically discriminates between these alternative models of relational representation.

There have been few direct investigations of relation priming, and the results are mixed. There is clear evidence that relations can be contextually primed, in terms of facilitating relational comprehension in general (Wisniewski & Love, 1998), or of specific relations that occur frequently in a list of stimuli (McKoon & Ratcliff, 1995). Gerrig and Murphy (1992, Experiment 4) embedded critical word pairs in context stories. Some stories contained a prime phrase that used the same conceptual relation as a target phrase. For instance, one context initially described a woman who carved a trumpet out of an olive (i.e., TRUMPET OLIVE), and subsequently included a relationally similar target phrase (i.e., KITTEN APPLE). Gerrig and Murphy found that the target phrase was comprehended more easily and more accurately when the context included the relationally similar prime phrase. This finding suggests that relation priming may occur without lexical repetition.

Spellman and colleagues (2001) presented word pairs in a lexical decision task (Experiment 2) and in a naming task (Experiment 3). Target word pairs (e.g., BEAR CAVE) were preceded by prime word pairs that used either the same relation (e.g., BIRD NEST) or a different relation (e.g., COPPER HORSE). When participants were instructed to focus on the relations between the words, significant relation priming was obtained. But when participants' attention was not specifically directed to relations, priming was not obtained. Thus, Spellman et al. provided corroborative evidence of relation priming, but this effect was task-sensitive.

The failure of Spellman and colleagues to find relation priming in some conditions could be attributable to the lexical nature of their tasks (lexical decision and naming). In a more semantic test of relation priming, Gagné (2001, Experiment 6) used a sensicality judgment task in which participants indicated whether the word pair made sense as a phrase. As in the study by Spellman and colleagues (2001), target word pairs followed prime word pairs that used either the same relation or a different relation. Gagné found no relation priming in this paradigm. However, in another experiment Gagné (2002, Experiment 2) used prime and target combinations that were more lexically similar, and in this case she did obtain relation priming. For example, comprehension of STUDENT VOTE was faster and more accurate following presentation of SCHOLAR ACCUSATION (Same Relation) than following SCHOLAR CAR (Different Relation).

One possible explanation of these discrepant results is that the stimuli in these various studies may have differed in some critical way. In particular, Estes (2003) noted that Gerrig and Murphy's prime and target stimuli used very similar relations (e.g., TRUMPET OLIVE → KITTEN APPLE), whereas Gagné's (2001) prime and target relations intuitively seemed less similar (e.g., BABY CRY → STUDENT VOTE). And when Gagné (2002) increased the similarity of the prime and target relations (e.g., SCHOLAR ACCUSATION → STUDENT VOTE), then relation priming was obtained. Thus, relation priming apparently occurs when the prime and target relations are highly similar (e.g., Gagné, 2002; Gerrig & Murphy, 1992), but not when the relations are only vaguely similar (e.g., Gagné, 2001). Estes (2003) therefore suggested that *relational similarity*, or the similarity between the prime and target relations, may explain the previously discrepant results. Estes conducted two further experiments using the sensicality judgment task devised by Gagné (2001), but with prime and target stimuli that instantiated more similar relations. Contrary to Gagné's (2001) null result, Estes (2003) did obtain a large relation priming effect (cf. Gagné, 2002). So, for instance, COPPER HORSE was understood faster, and was more likely to be understood, when preceded by GLASS EYE than by SIGN POST.

To test whether lexical priming could explain the relation priming effect, Estes (2003) presented the experimental stimuli in a lexical decision task (following the procedure of Gagné, 2001). Upon finding no difference between the Same Relation and Different Relation conditions, Estes concluded that lexical priming could not explain the observed effect. On the basis of a power analysis, however, Gagné et al. (2005) argued that the effect obtained by Estes (2003) may be attributable to lexical priming after all. In their Experiment 2, they showed that the Same Relation primes used by Estes (2003) were in fact more lexically similar to the target combinations than were the Different Relation primes. In two subsequent experiments, Gagné and colleagues controlled the lexical similarity between the prime and target combinations, and they again failed to find any evidence of relation priming. From this they reasserted their prior conclusion that relation priming does not exist, at least not in the absence of identical or highly similar modifiers.

Unfortunately, the study by Gagné et al. (2005) suffers from the same critical shortcoming as Gagné (2001): The prime and target relations are not similar enough to elicit relation priming (cf. Estes, 2003). For instance, Gagné et al. used EAR RING as a Same Relation prime for SUMMER MONEY, and BUG SPRAY was supposed to have the same relation as JUVENILE UNDERWEAR. Other examples include ICE RABBIT → VAPOR CLOUD, and LIP SOAP → TRAVEL MONTHS. Each of these stimulus pairs was presumed to exhibit the same relation between prime and target. But as these examples illustrate, for many of the stimulus pairs it is unclear what that common relation is meant to be. Just as one would not expect lexical priming when the lexical similarity between prime and target is very low, likewise, one should not expect relation priming when relational similarity is very low.

To summarize, then, it remains unclear whether relation priming is merely a conflation of lexical similarity (Gagné et al., 2005), or whether it is a genuine phenomenon that is mediated by relational similarity (Estes, 2003). This empirical question is theoretically critical, in that its answer will discriminate between bound and independent models of relational representation. Thus, we begin our investigation with regression analyses of datasets from several previously published experiments.

We then report a new experiment, under more controlled conditions, to directly investigate the occurrence or absence of relation priming.

Study 1: Regression analyses

The purpose of this study was to evaluate the relative contributions of the two presumed explanatory variables to the speed and likelihood of comprehending a target word pair. We conducted multiple regressions on the data of Estes (2003, Experiment 2) and Gagné et al. (2005, Experiments 3 and 4), with lexical similarity and relational similarity as predictor variables, and with response times and sensicality judgments as criterion variables. These two data sources (i.e., Estes, 2003; Gagné et al., 2005) were selected for analysis because they provided the most direct investigations of relation priming to date.

Method

Lexical similarity

Gagné et al. used latent semantic analysis (LSA; Landauer & Dumais, 1997) as a measure of lexical similarity between prime and target word pairs (see also Howard & Kahana, 2002; Steyvers, Shiffrin, & Nelson, 2004). Essentially, LSA uses a massive text corpus to extract the meanings of words from the contexts in which they occur. The underlying assumption is that if two words frequently occur in the same types of contexts, then those words presumably have similar meanings. LSA scores are bound from −1 (i.e., the words tend to occur in different contexts) to +1 (i.e., the words tend to occur in similar contexts).

Gagné et al. examined the lexical similarity scores for the stimuli used by Estes (2003, Experiment 2), as well as for their own stimuli (Gagné et al., 2005, Experiments 3 and 4). Estes (2003) had 24 target combinations, each with a Same Relation and a Different Relation prime, for a sum of 48 prime–target pairs. Gagné et al. had 60 target combinations, each with a Same Relation and a Different Relation prime, for a sum of 120 prime–target pairs. Thus, there were a total of 168 prime–target pairs (e.g., GLASS EYE → COPPER HORSE). For each pair, Gagné et al. obtained a lexical similarity score for the modifier concepts (i.e., GLASS and COPPER), and another score for the head concepts (i.e., EYE and HORSE). They found that, in the study by Estes (2003), the modifier similarity and the head similarity were both significantly higher in the Same Relation condition than in the Different Relation condition. Thus, Estes' manipulation of prime-type was confounded with lexical similarity. The stimuli used by Gagné et al., on the other hand, were more closely matched for both modifier similarity and head similarity. For the purposes of our regression analyses (reported below), we calculated a Lexical Similarity score for each prime–target pair by averaging its modifier similarity and head similarity.[1]

Relational similarity

Thirty-one undergraduates at the University of Georgia provided relational similarity ratings for all 168 prime–target pairs. Relational similarity was defined in the instructions as follows: "Relational similarity refers to how similar the inferred relations are between two combinations. For instance, because SOAP BASKET and CUCUMBER SALAD both have a CONTAINS relation, you might give them a high rating. Note that this is different from the more common form of similarity in which two concepts have the same features or belong to the same category. So although cucumbers and pickles are pretty similar, the relation for CUCUMBER SALAD may not be very similar to the relation for PICKLE JUICE, and therefore you might give them a lower rating."

Each trial consisted of a prime combination presented in isolation for 500 ms, followed by a target combination for another 500 ms, and finally a rating scale that ranged from 1 ("not at all relationally similar") to 7 ("very relationally similar"). All three stimuli on each trial were centered horizontally on screen, with the prime combination presented on the upper part of the screen, the target combination in the vertical center of the screen, and the rating scale in the lower part of the screen. All three stimuli remained on screen simultaneously until the participant responded. Each prime–target pair was presented to every participant in random order.

Results

Lexical Similarity and Relational Similarity are reported separately below, followed by the regression analyses of Estes (2003, Experiment 2) and Gagné et al. (2005, Experiments 3 and 4). One item used by Estes (2003) did not occur in the LSA corpus (i.e., "burrito"); that item was therefore excluded from all analyses.

Lexical similarity

Mean lexical similarity scores are presented in Table 1. Because Estes (2003) and Gagné et al. (2005) used different numbers of items, the studies were analyzed separately. For the stimuli used by Estes (2003), the lexical similarity of the prime and target combinations was significantly greater in the Same Relation condition than in the Different Relation condition [$t(22) = 5.06$, $p < .001$].

[1] When a multiplicative model of Lexical Similarity was used instead, the identical pattern of results was obtained.

Table 1
Lexical similarity and relational similarity, Study 1

	Relation-type			
	Same	Different	Same	Different
	Estes (2003)		Gagné et al. (2005)	
Lexical Similarity	.20 (.02)	.06 (.01)	.07 (.01)	.06 (.01)
Relational Similarity	5.16 (.14)	2.19 (.08)	2.71 (.12)	2.31 (.08)

Note. Lexical Similarity scores were calculated as the average of Modifier Similarity and Head Similarity, as determined by Latent Semantic Analysis, and are therefore bound between −1.00 and +1.00. Relational Similarity ratings were scaled from 1 (low) to 7 (high).

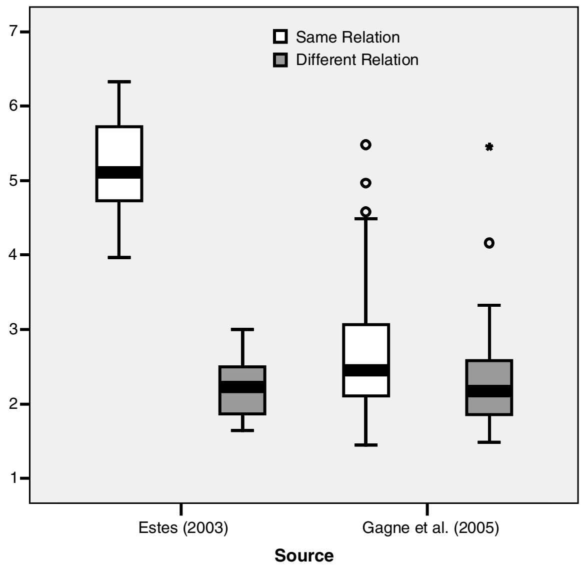

Fig. 1. Boxplots of relational similarity ratings.

This finding raises two important points. First, our averaged measure of Lexical Similarity corroborated Gagné et al.'s analysis, in which modifier similarity and head similarity were calculated separately. The present finding therefore supports the validity of our averaged measure. The second notable point is that the Same Relation and Different Relation conditions of Estes (2003) were conflated with lexical similarity. Thus, as correctly indicated by Gagné et al., the priming effect reported by Estes (2003) could be explained as lexical priming instead of relation priming.

For the stimuli used by Gagné et al., the Same Relation and the Different Relation primes marginally differed in lexical similarity to the target combination [$t(59) = 1.66$, $p = .10$]. Although this finding suggests that the stimuli of Gagné and colleagues were also somewhat conflated with lexical similarity, the magnitude of this difference was small (see Table 1), and was therefore unlikely to elicit differential lexical priming across conditions.

Relational similarity

Mean relational similarity ratings are presented in Table 1. As illustrated by the boxplots shown in Fig. 1, the difference between the Same Relation and Different Relation primes was greater for the stimuli of Estes (2003) than for the stimuli of Gagné and colleagues (2005). The Same Relation primes used by Estes were indeed high in relational similarity to their target combinations, and the Different Relation primes were in fact very low in relational similarity to their targets, thus yielding a significant difference in relational similarity [$t(22) = 18.16$, $p < .001$]. Although the Same Relation and Different Relation primes used by Gagné et al. also differed significantly in relational similarity [$t(59) = 4.07$, $p < .001$], both means were well below the midpoint of the scale, and hence relational similarity may be considered low for *both* relation-types in that study (see Fig. 1). This could explain why Estes (2003) found significant relation priming but Gagné et al. (2005) did not.

Also informative are the ranges of relational similarity from the two different studies. For the targets used by Estes (2003), relational similarity ranged from 3.97 to 6.32 for the Same Relation primes, and from 1.65 to 3.00 for the Different Relation primes. Because the ranges do not overlap, it is evident that the Same Relation prime was rated more relationally similar to the target than was the Different Relation prime for every prime–target pair. Thus, the stimulus set of Estes constitutes a strong manipulation of relational similarity. But for the targets used by Gagné et al. (2005), the range of relational similarity was virtually identical across prime conditions: Relational similarity ranged from 1.45 to 5.48 for the Same Relation primes, and from 1.48 to 5.45 for the Different Relation primes.

The high degree of overlap between Gagné's Same Relation and Different Relation primes suggests that, for some proportion of targets, the Same Relation prime may not have actually been as relationally similar as the Different Relation prime. To investigate this possibility further, we calculated for each target the difference in relational similarity between the Same Relation and Different Relation primes. By this analysis, positive scores indicate that the Same Relation prime was indeed rated more relationally similar to the target than was the Different Relation prime. For the stimuli used by Estes, the difference scores ranged from +1.32 to +4.29, thereby revealing that all primes were appropriately assigned to their respective conditions, and that indeed every target exhibited a substantial difference between prime-types. But for the stimuli used by Gagné et al.,

Table 2
Beta coefficients (β) for lexical similarity and relational similarity as predictors of response time (RT) and comprehension rate (CR), Study 1

	Estes (2003)		Gagné et al. (Experiment 3)		Gagné et al. (Experiment 4)	
	RT	CR	RT	CR	RT	CR
Lexical Similarity	+.18	−.02	+.06	−.05	+.10	−.22*
Relational Similarity	−.55**	+.62***	−.33***	+.27**	−.34***	+.35***

* $p < .05$.
** $p < .01$.
*** $p < .001$.

the difference scores ranged from −.84 to +2.45. In fact, 22 of the 60 targets yielded negative difference scores, and one target yielded a difference of zero. This finding indicates that 37% of the targets used by Gagné et al. had "Same Relation" primes that were actually *less* relationally similar than their Different Relation primes.

In sum, the stimuli used by Gagné and colleagues (2005) may be insufficient to elicit relation priming. The overall difference in relational similarity between the Same Relation and Different Relation primes was very small, both prime-types were well below the scalar midpoint, and indeed 37% of the "Same Relation" primes were in fact less relationally similar than the Different Relation primes. In contrast, the stimuli used by Estes (2003) exhibited a large difference in relational similarity, and every stimulus was correctly designated as either a Same Relation or a Different Relation prime. Hence, differential manipulations of relational similarity across the two studies may explain the differential finding of relation priming across those studies. This possibility is investigated further in the analyses reported below.

Analysis of Estes (2003, Experiment 2)

We conducted multiple regressions on the data of Estes (2003, Experiment 2) to determine whether the variance in response times and sensicality judgments was predicted by relational similarity, lexical similarity, or both.[2] Results are summarized in Table 2. In the response time analysis, Relational Similarity and Lexical Similarity collectively accounted for a significant amount of the variance, with a multiple R of .45, $R^2 = .20$, F (2, 43) = 5.43, $p < .01$. As illustrated in Fig. 2, Relational Similarity significantly predicted response times ($\beta = -.55$, $t = 2.99$, $p < .01$). However,

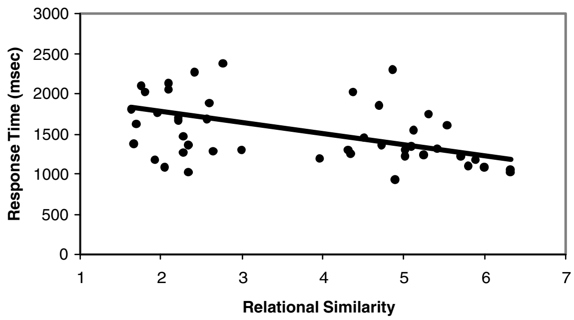

Fig. 2. Response time as a function of relational similarity, Estes (2003, Experiment 2).

contrary to the suggestion of Gagné et al., Lexical Similarity did not approach reliability as a predictor of comprehension time ($\beta = +.18$, $p = .33$). The analysis of sensicality judgments yielded an analogous pattern of results, with a highly significant multiple R of .61, $R^2 = .37$, $F(2,43) = 12.41$, $p < .001$. The variance again was explained by Relational Similarity ($\beta = +.62$, $t = 3.77$, $p < .001$) and not by Lexical Similarity ($\beta = -.02$, $p = .90$). Fig. 3 plots the likelihood of comprehension as a function of relational similarity. Thus, the claim of Gagné et al. that lexical similarity can explain the result of Estes (2003) received no support. Instead, comprehension was reliably predicted by the

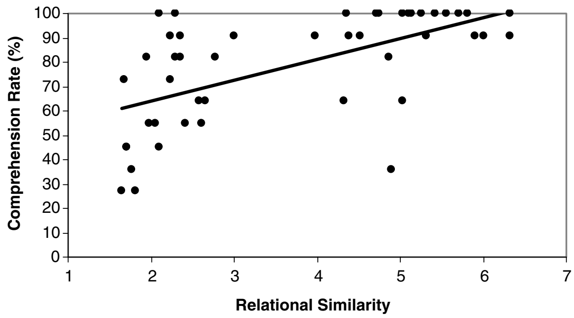

Fig. 3. Comprehension rate as a function of relational similarity, Estes (2003, Experiment 2).

[2] An alternative approach is to submit the data of Estes (2003) to an ANCOVA in which Lexical Similarity is treated as a covariate. This analysis yielded a significant main effect of prime-type on both response times [$F(1,20) = 6.09$, $p < .05$] and sensicality judgments [$F(1,20) = 14.83$, $p < .001$]. That is, even when the effect of Lexical Similarity is statistically controlled, the relation priming effect persists.

relational similarity between the prime and target combinations: As relational similarity increased, response times decreased (Fig. 2) and comprehension rates increased (Fig. 3).

Analysis of Gagné et al. (2005, Experiment 3)

Multiple regressions were also conducted on the data of Gagné et al. (Experiments 3 and 4). The two experiments were identical except for one methodological difference. Participants in Experiment 3 judged the sensicality of both primes and targets, whereas participants in Experiment 4 responded to the targets only. Gagné et al. suggested that the latter method might be more susceptible to strategic effects, which in turn might be more likely to induce relation priming (see Spellman et al., 2001). Despite this methodological difference, Gagné and colleagues failed to find relation priming in either study. Hence, it would be particularly striking if relational similarity were to explain a significant amount of the variance in either of those experiments. Analyses of the two experiments are reported separately, and results are summarized in Table 2.

In Experiment 3 of Gagné et al., Relational Similarity and Lexical Similarity collectively exerted a significant influence on response times, $R = .31$, $R^2 = .10$, $F(2, 117) = 6.31$, $p < .01$. In fact, Relational Similarity predicted response times ($\beta = -.33$, $t = 3.52$, $p < .001$), but Lexical Similarity did not ($\beta = +.06$, $p = .52$). Fig. 4 illustrates the significant relationship between response time and relational similarity. This relationship was corroborated by the sensicality judgments, as shown in Fig. 5. The multiple R of .26 was significant, $R^2 = .07$, $F(2,117) = 4.24$, $p < .05$. Sensicality judgments were reliably predicted by Relational Similarity ($\beta = +.27$, $t = 2.89$, $p < .01$), but not by Lexical Similarity ($\beta = -.05$, $p = .59$). Thus, an effect of relational similarity was apparent in the data of Gagné and colleagues, despite their failure to observe relation priming. Note also that this effect emerged from the method that Gagné and colleagues considered least likely to elicit relation priming, since strategic effects were minimized by disguising the prime–target structure of the experiment.

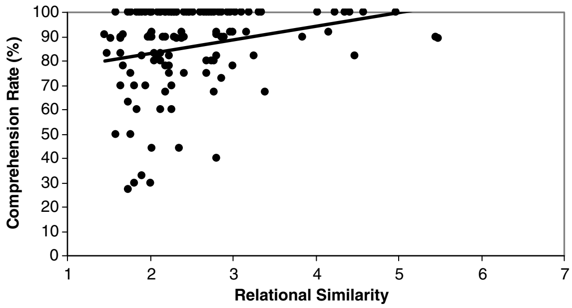

Fig. 5. Comprehension rate as a function of relational similarity, Gagné et al. (2005, Experiment 3).

Analysis of Gagné et al. (2005, Experiment 4)

In Experiment 4 of Gagné et al., Relational Similarity and Lexical Similarity together accounted for a significant amount of the variance in response times, with a multiple R of .32, $R^2 = .11$, $F(2, 117) = 6.88$, $p < .01$. More specifically, response times were predicted by Relational Similarity ($\beta = -.34$, $t = 3.71$, $p < .001$) but not by Lexical Similarity ($\beta = +.10$, $p = .29$). Fig. 6 presents response time as a function of relational similarity. The multiple regression was also significant in the analysis of sensicality judgments, with $R = .35$, $R^2 = .12$, $F(2,117) = 8.24$, $p < .001$. As in the previous analyses, Relational Similarity significantly predicted sensicality judgments ($\beta = +.35$, $t = 3.82$, $p < .001$). See Fig. 7 for an illustration of this relationship. The relation between Lexical Similarity and sensicality judgments was also significant ($\beta = -.22$, $t = 2.47$, $p < .05$), but surprisingly this relationship was negative. That is, as the lexical similarity between prime and target increased, the likelihood of comprehending the target decreased. As far as we are aware, no cognitive theory would predict this negative relationship. Moreover, no such relationship was observed in the preceding analyses of Estes (2003)

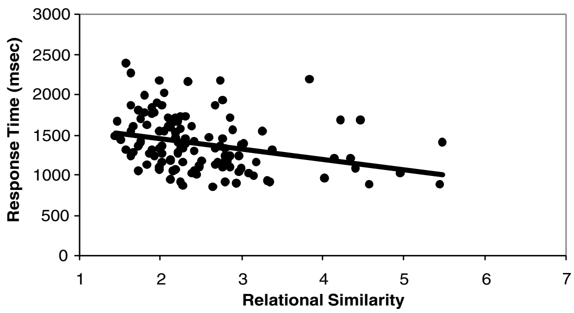

Fig. 4. Response time as a function of relational similarity, Gagné et al. (2005, Experiment 3).

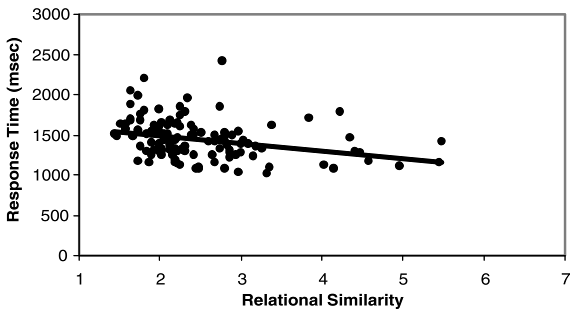

Fig. 6. Response time as a function of relational similarity, Gagné et al. (2005, Experiment 4).

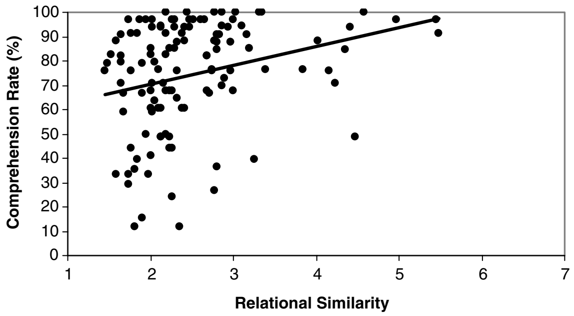

Fig. 7. Comprehension rate as a function of relational similarity, Gagné et al. (2005, Experiment 4).

and Gagné et al. (Experiment 3). In general, then, the speed and likelihood of comprehending a target phrase was a function of its relational similarity to the prime phrase.

Discussion

The impetus for the current analysis is evident in Table 1, and the critical results are summarized in Table 2. Essentially, Table 1 shows that the target combinations of Estes (2003) were more lexically similar *and* more relationally similar to the Same Relation primes than to the Different Relation primes. Thus, the priming effect reported by Estes could be attributable to either lexical priming or relational priming. To resolve this ambiguity, we conducted regression analyses. As shown in Table 2, these analyses clarify that in fact relational similarity predicted both response times and comprehension rates in the study by Estes (2003), whereas lexical similarity did not reliably predict either measure of comprehension. Moreover, despite the fact that Gagné et al. (2005, Experiments 3 and 4) failed to obtain relation priming, Table 2 reveals that relational similarity nonetheless predicted both response times and comprehension rates in both of those experiments as well. We suggest that the reason Gagné et al. failed to find relation priming is that their manipulation of relational similarity simply was not strong enough (see Fig. 1).[3]

Gagné (2001, 2002; Gagné et al., 2005) claims that the modifier concepts, in particular, must be similar (or identical) in order for relation priming to occur. One could argue, then, that the averaged measure of Lexical Similarity used in our regression analyses is insufficient insofar as it includes the presumably irrelevant factor of head concept similarity. We therefore analyzed the same datasets yet again, but with Modifier Similarity as a predictor variable (rather than the averaged similarity). This more restricted measure of lexical similarity yielded no better results. In the study by Estes (2003), Modifier Similarity did not reliably predict response times ($\beta = -.05$, $p = .77$) or sensicality judgments ($\beta = +.01$, $p = .95$). Nor did Modifier Similarity predict response times ($\beta = +.08$, $p = .37$) or sensicality judgments ($\beta < .01$, $p = .97$) in Experiment 3 of Gagné et al. (2005). In their Experiment 4, Modifier Similarity did predict both response times ($\beta = +.16$, $p = .07$) and sensicality judgments ($\beta = -.17$, $p = .05$), but again these results were in the opposite direction to what one would expect. As the prime and target modifiers increased in lexical similarity, comprehension of the target combination became slower and less likely. This result is counterintuitive, and counter to the prediction of Gagné and colleagues. Moreover, even if lexical similarity were to predict comprehension in a meaningful way, that would not diminish the contribution of relational similarity to comprehension.

The current analysis suggests that relation priming is a genuine phenomenon that cannot be explained as lexical priming. Across three experiments conducted in two different labs, the speed and likelihood of comprehending a target phrase was a function of its relational similarity to the prime phrase. We next sought to replicate this relation priming effect under improved experimental control.

Study 2: New evidence of relation priming

In Study 2, we controlled the lexical similarity of the Same Relation and Different Relation primes, while manipulating their relational similarity to the targets. This experiment essentially combined the control suggested by Gagné et al. (2005) with the strong manipulation of relational similarity demonstrated by Estes (2003).

Lexical similarity and relational similarity are naturally correlated—that is, similar concepts tend to instantiate the same relations.[4] If two concepts are highly similar (e.g., BICYCLE and SKATEBOARD), then they may instantiate many of the same relations (e.g., BICYCLE TRACK and SKATEBOARD RAMP). And conversely, if two concepts are highly dissimilar (e.g., BICYCLE and WIG),

[3] More specifically, the misallocation of stimuli to the Same Relation and Different Relation conditions by Gagné et al. (2005) may explain why their analysis of variance failed to detect the relationship between relational similarity (i.e., prime-type) and comprehension.

[4] To examine this relationship, we calculated the correlations between lexical similarity and relational similarity in the studies by Estes (2003) and Gagné et al. (2005). As suggested by the data in Table 1, there was a strong positive correlation in the study by Estes, $r(46) = +.67$, $p < .001$. There was also a significant positive correlation in the study by Gagné et al., $r(120) = +.30$, $p < .001$.

there will be few relations that both concepts instantiate. Thus, to date, relation priming has only been observed when the prime and target phrases were lexically similar. The relation priming found by Estes (2003) was obtained with stimuli that were relatively high in lexical similarity (see Table 1). Furthermore, Gagné found that comprehension of a target phrase (e.g., STUDENT VOTE) was facilitated by a prime with a lexically similar modifier (e.g., SCHOLAR ACCUSATION; Gagné, 2002), but not by a prime with a dissimilar modifier (e.g., BABY CRY; Gagné, 2001). Tagalakis and Keane (2005) also found that comprehension of a target combination was faster after a prime with a similar modifier (e.g., GAS TANK → COAL TANK) than after a prime with a dissimilar modifier (e.g., SAND TANK → COAL TANK).

Notice, however, that the independent model predicts that relation priming may occur even when the primes and targets are lexically dissimilar, so long as they instantiate the same relation. In Study 2, we therefore sought to provide a strict test of the independent model by using primes and targets that were lexically dissimilar from one another. COPPER HORSE may prime GLASS EYE just because COPPER and GLASS are relatively similar. The current question is whether relation priming will also obtain with stimuli that are lexically dissimilar: Does STEEL SCISSORS prime STRAW HAT? Such a result would provide strong evidence for the independent model.

Whereas the previous investigations of relation priming have included a multitude of relations (e.g., *cause*, *location*, *part-of*, etc.), all experimental targets in the current study used the same relation (i.e., *composed-of*). Our choice of this relation was largely practical— *composition* phrases may be easily generated according to a general rule (i.e., ⟨substance + ⟨object⟩;-ject⟩; e.g., COTTON SOCKS, CRYSTAL VASE, etc.). Furthermore, because most objects may be composed of several different substances, relational similarity could be manipulated independently of lexical similarity. The theoretical benefit of holding the relation constant is that it permits detailed conclusions about one particular relation. For example, Chaffin and Herrmann (1988; Chaffin, Herrmann, & Winston, 1988; Winston, Chaffin, & Herrmann, 1987) achieved a comprehensive evaluation of the *part/whole* relation, and more recently Fenker, Waldmann, and Holyoak (2005) investigated in detail the *causal* relation. In contrast, the failure of prior studies on relation priming (e.g., Estes, 2003; Gagné et al., 2005) to control or manipulate relation-type has precluded inference about any particular relation in detail. Of course, a potential cost of using a single relation-type is that participants may become attuned to the repeated occurrence of this relation. But such a general attenuation of this relation (cf. McKoon & Ratcliff, 1995) would apply across experimental conditions, and thus could not explain any difference that might obtain between them.

Gagné et al. (2005) suggested that the prime–target structure of Estes' (2003) experiment could induce strategic effects (as previously described). The implication is that the relation priming effect may be an artifact of the experimental procedure. Indeed, this implication is supported by the findings of Spellman and colleagues (2001), who obtained relation priming only when participants' attention was directed to the conceptual relations between stimuli. To minimize the possibility of such strategic effects, we used a continuous sensicality judgment task. Participants were required to indicate whether the prime combination was sensible, and then whether the target combination was sensible, so that primes and targets were procedurally indistinguishable. Thus, if a relation priming effect were to obtain, it would not be due to the prime–target structure of the experiment.

Method

Participants

Fifty undergraduates at the University of Georgia participated in the experiment proper, and an additional 101 participated in stimulus norming tasks. All received partial course credit for participation.

Materials

Target items consisted of 30 combinations instantiating a *composition* relation (e.g., FUR COAT). For each target, 30 Same Relation primes (e.g., CHAIN FENCE) and 30 Different Relation primes (e.g., CHAIN LOCK) were constructed, with both primes sharing the same modifier (Gagné et al., 2005). Primes and targets were pre-tested for both lexical similarity and relational similarity, as described below. See Appendix A for the complete set of experimental items.

Filler items included 60 additional sensible combinations and 120 nonsense combinations (e.g., CHOIR SEED) so that there would be an equal number of "yes" and "no" responses in the sensicality judgment task. These fillers were assigned as primes and targets such that the sensibility of the prime was not predictive of the sensibility of the target. Thus, there were 120 trials consisting of the following prime–target pairs: 30 sense–sense (15 Same Relation and 15 Different Relation), 30 sense–nonsense (e.g., PLUM WINE → FACT ROACH), 30 nonsense–sense (e.g., WORM SALON → COAL DUST), and 30 nonsense–nonsense (e.g., SWEATER VEGETABLE → MERCY DIRT).

Lexical similarity. In the preceding study, similarity was calculated via LSA (following Gagné et al., 2005), which is essentially a measure of lexical co-occurrence (see Landauer & Dumais, 1997). However, lexical co-occurrence likely reflects various other factors (e.g., McRae & Boisvert, 1998, p. 570) that are irrelevant to the present purposes (e.g., familiarity). Thus, in the present study we collected similarity ratings, which provide a more direct measure of similarity. Because both primes used the same modifier, lexical similarity of the modifiers was perfectly controlled (Gagné et al., 2005). Nonetheless, we collected similarity ratings of the 30 prime–target modifier pairs (e.g., CHAIN and FUR). Thirty-one undergraduates rated the similarity of the modifier pairs on a scale from 1 (not at all similar) to 7 (very similar). The lexical similarity of the modifiers was very low ($M = 1.71$, $SE = .10$). Thus, if relation priming were to obtain with these stimuli, it would not be attributable to the similarity of the modifier concepts.

We also controlled the lexical similarity of the head nouns via similarity ratings. Forty-two participants rated the similarity between all 60 prime and target head noun pairs (e.g., FENCE and COAT; LOCK and COAT) on the same 1 (not at all similar) to 7 (very similar) scale. Data were submitted to one analysis with participants random (t_p) and another with items random (t_i). Head noun similarity was virtually identical across the Same Relation ($M = 1.80$, $SE = .08$) and Different Relation ($M = 1.80$, $SE = .09$) conditions, $t_p(41) = .16$, $p = .88$ and $t_i(29) = .05$, $p = .96$. Any difference found between the Same Relation and the Different Relation conditions cannot be attributed to lexical priming.

Relational similarity. To validate our assignment of primes to the Same Relation and Different Relation conditions, we used a forced-choice task in which 28 participants selected which of the two primes (e.g., STEEL SCISSORS or STEEL FACTORY) was more relationally similar to the target (e.g., STRAW HAT). Relational similarity was defined as before, though a different example was provided in the instructions: "... a BASEBALL BRUISE is a bruise that was CAUSED by baseball. For this experiment, you will choose which of two combinations (e.g., WEDDING DEBT or TENNIS SKIRT) is more RELATIONALLY similar to the first combination (e.g., BASEBALL BRUISE). Relational similarity refers to how similar the inferred relations are between two combinations. Note that this is different from the more common form of similarity in which two concepts have the same features or belong to the same category. So although baseball and tennis are pretty similar, the relation for TENNIS SKIRT is not very similar to the relation for BASEBALL BRUISE. In contrast, baseball and weddings are not very similar at all, but WEDDING DEBT and BASEBALL BRUISE both have a causal relation. Therefore you might select WEDDING DEBT as the more similar combination."

Results confirmed that the Same Relation primes were more relationally similar to their targets than were the Different Relation primes. Overall, the Same Relation primes were selected more often than would be expected by chance ($M = .78$, $SE = .03$), $t_p(27) = 10.09$, $p < .001$ and $t_i(29) = 15.08$, $p < .001$. In fact, the Same Relation primes were the dominant choice (>50%) for each of the 30 targets. The manipulation of relational similarity was therefore successful and consistent. Furthermore, relational similarity was unrelated to the lexical similarity of the prime and target: the probability of choosing a given prime combination was not correlated with the lexical similarity of its modifier and head concepts, $r(60) = -.01$, $p = .94$. Thus, relational similarity and lexical similarity were indeed orthogonal.

Procedure

The procedure was modeled after that of Gagné et al. (2005, Experiment 3). Each trial began with a fixation cross presented in the center of the screen for 250 ms, followed immediately by presentation of the prime combination. The prime remained onscreen until the participant indicated by keypress whether the word pair made sense as a phrase. After a 1-s inter-stimulus interval, the fixation cross appeared again for 250 ms, followed by presentation of the target combination. The target remained onscreen until a sense/nonsense judgment was given. The inter-trial interval was 1 s, so that nothing in the presentation method indicated a pairing of the prime and target combinations. All stimuli were presented in 18-point red font in the center of a black background. The ⟨J⟩ key was used to indicate that the combination had a sensible interpretation and the ⟨F⟩ key to indicate that it did not. Participants were instructed to respond as quickly as possible without making errors. Ten practice trials preceded the experimental trials, which were presented in random order.

Results and discussion

One participant was excluded from all analyses on the basis of a low comprehension rate (<80%). For the primes, comprehension rates did not differ between the Same Relation and Different Relation conditions ($M = .89$, $SE = .03$), $t < 1$. For the targets, comprehension rates were virtually identical across the Same Relation and Different Relation conditions ($M = .96$, $SE = .01$), $t < 1$. Note that the comprehension rates for these stimuli were markedly higher than in previous studies of relation priming. For instance, Estes (2003) reported comprehension rates that ranged from 67 to 91% for relational phrases, and Gagné et al. (2005) had sensicality judgments in the range of 73 to 85%.

Target response times greater than 2.5 standard deviations from a participant's condition mean were excluded from analyses (1.3% of trials). Trials in which the target was judged to be nonsensical were also excluded from response time analyses (4.2% of trials). Critically, targets were comprehended significantly faster after a Same Relation prime ($M = 893$, $SE = 30$) than after a Different Relation prime ($M = 926$, $SE = 29$),

$t_p(48) = 2.12$, $p < .02$ and $t_i(29) = 1.78$, $p < .05$, one-tailed.[5] This result corroborates the analyses of Study 1.

Lexically dissimilar prime phrases (e.g., CHOCOLATE CAKE) reliably facilitated comprehension of relationally similar target phrases (e.g., GRAVEL ROAD). Although the magnitude of this effect (33 ms) was smaller than that found by Estes (2003), its occurrence nonetheless supports the conclusion that relation priming may obtain even in the absence of lexical priming. Because there was no lexical repetition across primes and targets, because the primes and targets were lexically dissimilar, and because the primes and targets were procedurally indistinguishable, a strategic search for relational similarity was unlikely. Thus, the present experiment is the most powerful demonstration of relation priming to date. The result provides strong support for the independent model of relational representation. If relations have concept-independent representations, then the same relational representation may be utilized by various unrelated concepts, thereby resulting in relation priming.

General discussion

Across three previously published experiments (Study 1) and one novel experiment (Study 2), the relational similarity between phrases exerted a consistent and significant effect on comprehension: As the relational similarity between prime and target phrases increased, the speed and likelihood of comprehending that target also increased. This effect of relational similarity was observed regardless of whether the prime and target phrases were lexically similar or dissimilar, and regardless of whether the prime–target nature of the experiment was apparent or disguised. The effect of relational similarity was even observed in two experiments that failed to obtain relation priming (i.e., Gagné et al., 2005).

The present results suggest that relational representations are concept-independent (Estes, 2003; see also Spellman et al., 2001). Reading a prime phrase (GLASS EYE) activates an independent relational representation (*composed of*), which may then be utilized by the target phrase (COPPER HORSE), thereby facilitating comprehension of that target. Moreover, a relationally similar prime (e.g., STEEL SCISSORS) can facilitate comprehension of a lexically dissimilar target (e.g., STRAW HAT). This latter result provides especially forceful evidence that relation priming is independent of lexical priming, and that relations are represented independent of the concepts that instantiate them.

Corroborative support for the independent model comes from the finding that conceptual relations exhibit typicality effects (Chaffin & Herrmann, 1984; Chaffin et al., 1988). That is, various instantiations of a given relation will differ in how typical they are of that relation. For example, Chaffin and Herrmann (1988) presented instances of the *part/whole* relation (e.g., toe-foot; bed-mattress), and had participants rate each pair on a scale from "a very good example of the part-whole relation" to "very dissimilar from the part-whole relation." They then had another group of participants make speeded judgments indicating whether one concept (e.g., toe) was a part of the other concept (e.g., foot). Chaffin and Herrmann found that typical instances of the relation were verified faster than atypical instances. This finding suggests that the *part/whole* relation has a general prototype or idealized representation against which the particular instantiations of that relation were compared. In contrast, if the *part/whole* relation were represented with each particular concept, then the typical and atypical instances should be verified equally fast. Thus, our conclusion of independent relational representations is not only consistent with other studies of relation priming (i.e., Estes, 2003; Gerrig & Murphy, 1992; Spellman et al., 2001), but also accounts for relational typicality effects (Chaffin & Herrmann, 1984, 1988; Chaffin et al., 1988).

These results signal a rejection of all models in which relational representations are bound to specific concepts (e.g., Gagné, 2001; Gagné et al., 2005). According to the bound model, relation priming should obtain only if the same concept occurs in the prime and target combinations. GLASS EYE may facilitate comprehension of GLASS HORSE, because the *composed of* relation would be pre-activated by the prime. But GLASS EYE should not prime comprehension of COPPER HORSE, because relations are presumed to be concept-specific, and in this case no concept is repeated from prime to target.

The previous finding that relation priming can occur with similar but nonidentical concepts (Estes, 2003; Gagné, 2002) appears to contradict the bound model of representation—given the lack of lexical repetition, relation priming should not occur (Gagné, 2001). In attempt to accommodate this finding within the framework of a bound model, Gagné (2002) proposed that comprehension of a phrase may involve retrieval of other phrases that have a similar modifier, and that the relations used by those retrieved phrases may influence comprehension of the target (see also van Jaarsveld, Coolen, & Schreuder, 1994). For instance, when interpreting the phrase STUDENT VOTE, the modifier STUDENT triggers the retrieval of other similar concepts (e.g., SCHOLAR), which in turn elicits retrieval of phrases using those concepts (e.g., SCHOLAR ACCUSATION). Finally, the relation instantiated by the retrieved phrase (i.e., Y *by* X) is applied to the original phrase (i.e., STUDENT

[5] The use of a one-tailed test is justified by (1) the directional nature of the hypothesis, and (2) the results of Study 1.

VOTE). If the relation provides an appropriate interpretation, then the target is comprehended.

We find this explanation of relation priming implausible on both a priori and empirical grounds. To begin with, this process appears at face value to be quite time-consuming, so it is not at all clear how it could speed comprehension of the target. Second, this explanation cannot account for the present results. Because the prime and target phrases of Study 2 were lexically dissimilar, the similarity-based retrieval process suggested by Gagné (2002) would fail to elicit retrieval of the prime. In contrast, the independent model simply claims that the most recently activated relation (i.e., the prime) is applied to the target combination, regardless of their lexical similarity. This process would appear to be faster and easier than the proposed similarity-based retrieval process, and it would provide a more parsimonious and complete explanation of the observed data.

This research was motivated by an empirical puzzle, specifically, that Estes (2003) obtained relation priming whereas Gagné et al. (2005) did not. The present results offer a resolution, namely, that the primes and targets used by Gagné and colleagues were not as relationally similar as those used by Estes (2003). But this empirical resolution raises a further theoretical question: Why were the primes and targets of Gagné et al. so low in relational similarity? We believe that their use of relationally dissimilar primes and targets may be attributed to their adoption of Levi's (1978) taxonomy of general relations. Several linguistic analyses of various text corpora (e.g., Lees, 1963; Levi, 1978; Ryder, 1994; Warren, 1978) have converged on the conclusion that a small set of general conceptual relations (roughly, between six and 16) is sufficient to characterize the vast majority of nominal combinations that occur in ordinary discourse (for review see Ryder, 1994). The stimuli used by Gagné and colleagues are based on one such general linguistic taxonomy (see also Coolen, van Jaarsveld, & Schreuder, 1991; Gagné, 2001; cf. Spellman et al., 2001).

Levi's (1978) taxonomy has often been criticized for the vagueness of its relations (e.g., Clark, 1983; Downing, 1977; Murphy, 1988, 2002). For instance, according to Levi (1978), NUT BREAD, VEGETABLE SOUP, and FRUIT TREE all instantiate a general *have* relation. Alternatively, these combinations may utilize specific and distinct relational representations. The relation in NUT BREAD indeed seems relatively similar to that of VEGETABLE SOUP (i.e., *contain*), but intuitively, it differs markedly from that of FRUIT TREE (i.e., *produce*). Downing (1977) provides another example: "[Levi] analyses both *headache pills* and *fertility pills* as deriving from underlying structures containing FOR—though headache pills are designed to eliminate headaches, while fertility pills are intended to enhance fertility" (p. 814).

Because Gagné et al. based their stimuli on Levi's taxonomy, many of their primes and targets that were assumed to instantiate the same relation instead utilize different relations. For example, Gagné and colleagues used TIRE RIM as a Same Relation prime for the target FAMILY COW. Indeed, both of these phrases could be comprehended via a general *of* relation, as the rim *of* a tire and a cow *of* the family. By a more specific cognitive analysis, however, TIRE RIM might instantiate something more akin to a *part-of* relation whereas FAMILY COW instantiates a very different *possession* relation. Thus, the failure of Gagné et al. (2005) to obtain relation priming might ultimately be attributable to their assumption that relational representations are general. In contrast, when the prime and target phrases instantiate the same specific relation, then their relational similarity will be high, and hence relation priming may occur. Determining exactly what relations are represented, and how specific those representations are, is an onerous but important task for future investigation. One conclusion from the present research is that linguistic theory may guide the development of a more comprehensive taxonomy of relations, but ultimately any such taxonomy must be empirically validated.

Acknowledgments

We are grateful to Christina Gagné, who graciously provided her data to us for re-analysis. We also thank several anonymous reviewers for their suggestions, as well as Adam Pope, Bonney Reed, Patrick Thompson, and Kevin Wilson for their assistance in data collection.

Appendix A

Stimuli of Study 2

Target	Same relation	Different relation
aluminum cage	butter milk	butter dish
beef broth	charcoal smoke	charcoal pencil
brass frame	flour tortilla	flour bag
brick patio	gold bracelet	gold price
bronze medal	peach daiquiri	peach pit
caramel candy	leaf nest	leaf blower
cardboard box	snow fort	snow sled
ceramic jar	tomato soup	tomato garden
cheese sauce	plastic toy	plastic mold
clay pottery	ice sculpture	ice tray
cloth napkin	log cabin	log truck
copper money	wheat bread	wheat field
cotton socks	silver bell	silver mine
crystal vase	wood balcony	wood polish
feather pillow	pumpkin pie	pumpkin patch
foam cushion	rubber gloves	rubber plant
fur coat	chain fence	chain lock
glass bottle	leather sofa	leather catalog
grape juice	concrete bridge	concrete truck

Appendix A (*continued*)

Target	Same relation	Different relation
gravel road	chocolate cake	chocolate shop
iron gate	rice paper	rice cooker
linen pants	sand castle	sand shovel
nylon jacket	olive oil	olive branch
polyester shirt	apple cider	apple core
satin sheets	cinnamon stick	cinnamon spoon
straw hat	steel scissors	steel factory
styrofoam cup	stone house	stone cutter
tin can	wintergreen mint	wintergreen leaf
vinyl floor	silk tie	silk worm
wool sweater	marble stairs	marble sculptor

References

Chaffin, R., & Herrmann, D. J. (1984). The similarity and diversity of semantic relations. *Memory & Cognition, 12*, 134–141.

Chaffin, R., & Herrmann, D. J. (1988). Effects of relation similarity on part-whole decisions. *Journal of General Psychology, 115*, 131–139.

Chaffin, R., Herrmann, D. J., & Winston, M. (1988). An empirical taxonomy of part-whole relations: Effects of part-whole relation type on relation identification. *Language and Cognitive Processes, 3*, 17–48.

Clark, H. H. (1983). Making sense of nonce sense. In G. B. Flores d'Arcais & R. J. Jarvella (Eds.), *The process of language understanding* (pp. 297–331). New York: John Wiley.

Coolen, R., van Jaarsveld, H. J., & Schreuder, R. (1991). The interpretation of isolated novel nominal compounds. *Memory & Cognition, 19*, 341–352.

Downing, P. (1977). On the creation and use of English compound nouns. *Language, 53*, 810–842.

Estes, Z. (2003). Attributive and relational processes in nominal combination. *Journal of Memory and Language, 48*, 304–319.

Fenker, D. B., Waldmann, M. R., & Holyoak, K. J. (2005). Accessing causal relations in semantic memory. *Memory & Cognition, 33*, 1036–1046.

Gagné, C. L. (2001). Relation and lexical priming during the interpretation of noun–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 236–254.

Gagné, C. L. (2002). Lexical and relational influences on the processing of novel compounds. *Brain and Language, 81*, 723–735.

Gagné, C. L., Spalding, T. L., & Ji, H. (2005). Re-examining evidence for the use of independent relational representations during conceptual combination. *Journal of Memory and Language, 53*, 445–455.

Gerrig, R. J., & Murphy, G. L. (1992). Contextual influences on the comprehension of complex concepts. *Language and Cognitive Processes, 7*, 205–230.

Howard, M. W., & Kahana, M. J. (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language, 46*, 85–98.

Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review, 10*, 785–813.

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*, 211–240.

Lees, R. B. (1963). *The grammar of English nominalizations*. Bloomington, IN: Indiana University Press.

Levi, J. N. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review, 7*, 618–630.

McKoon, G., & Ratcliff, R. (1995). Conceptual combinations and relational contexts in free association and in priming in lexical decision and naming. *Psychonomic Bulletin & Review, 2*, 527–533.

McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 558–572.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90*, 227–234.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science, 12*, 529–562.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Ryder, M. E. (1994). *Ordered chaos: The interpretation of english noun–noun compounds*. Berkeley, CA: University of California Press.

Spellman, B. A., Holyoak, K. J., & Morrison, R. G. (2001). Analogical priming via semantic relations. *Memory & Cognition, 29*, 383–393.

Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications: Festschrift in Honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer*. Washington, DC: American Psychological Association.

Tagalakis, G., & Keane, M. T. (2005). How understanding novel compounds is facilitated by priming from similar, known compounds. *Proceedings of the Cognitive Science Society*.

van Jaarsveld, H. J., Coolen, R., & Schreuder, R. (1994). The role of analogy in the interpretation of novel compounds. *Journal of Psycholinguistic Research, 23*, 111–137.

Warren, B. (1978). *Semantic patterns of noun–noun compounds*. Goteborg, Sweden: Acta Universitatis Gothoburgensis.

Winston, M. E., Chaffin, R., & Herrmann, D. (1987). A taxonomy of part-whole relations. *Cognitive Science, 11*, 417–444.

Wisniewski, E. J., & Love, B. C. (1998). Relations versus properties in conceptual combination. *Journal of Memory and Language, 38*, 177–202.