Available online at www.sciencedirect.com
ScienceDirect
Journal of Memory and Language 57 (2007) 380–395
Journal of Memory and Language
www.elsevier.com/locate/jml

Priming the interpretation of noun–noun combinations

Claudine N. Raffray [a,*], Martin J. Pickering [b], Holly P. Branigan [b]

[a] Department of Linguistics, University of KwaZulu-Natal, South Africa
[b] Department of Psychology, University of Edinburgh, Scotland

Received 28 June 2006; revision received 6 June 2007
Available online 13 August 2007

Abstract

Noun–noun combinations like *dog scarf* are common in everyday discourse but often have more than one interpretation. How do language users arrive at an interpretation of the relationship between the two nouns? This paper reports three expression-picture matching experiments that used a priming paradigm to investigate the influence of modifier and head constituents on the comprehension of novel ambiguous noun–noun combinations. Experiment 1 examined the effects of lexical repetition and semantic relation. We found reliable relation priming, regardless of whether the modifier or head was repeated between prime and target: Participants tended to choose target pictures involving the same relation as a preceding prime picture. Experiment 2 demonstrated significant relation priming when neither constituent was repeated. Experiment 3 showed significant relation priming when each picture contained both possible semantic relations, arguing against a possible visual-priming account of the effect. We interpret the findings in light of competing models of conceptual combination.

Keywords: Conceptual combination; Noun–noun combinations; Relation priming; Semantic priming; Thematic relations; Compound nouns

How do people determine the meaning of complex expressions such as *tourist castle*? Even when the meaning of the individual words is not in question, their meanings may be combined in many different ways. Thus, *tourist castle* appears to mean a castle for tourists, but the apparently similar *mountain castle* appears to mean a castle located on a mountain. Still other combinations can mean more than one thing—is a *dog scarf* a scarf possessed by a dog or a scarf with images of dogs on it?

This process of *conceptual combination* is highly productive and serves many communicative functions (Downing, 1977). Some combinations (or *combined concepts*, *compounds*, or *complex concepts*) are well established (e.g., *computer chip*), but others are novel (e.g., *dog scarf*). The meaning of the product of this process is unconstrained in some ways (so that many different relationships between the individual concepts are possible) but constrained in others. In English, the second noun in the phrase is typically the head noun and denotes the category; the modifier precedes the head and serves to specify the way in which the noun differs from other members of its category (Berman & Clark, 1989; Clark & Berman, 1987; Glucksberg & Estes, 2000). In this paper we are concerned with the way in which prior experiences may influence how people interpret such nominal (noun–noun) combinations. So how do people combine such information to produce appropriate interpretations?

* Corresponding author. Fax: +27 0 31 260 1253.
E-mail address: Raffray@ukzn.ac.za (C.N. Raffray).

0749-596X/$ - see front matter
doi:10.1016/j.jml.2007.06.009

Two competing theoretical approaches to conceptual combination dominate the recent literature. One approach is relation-based, known as the Competition Among Relations In Nominals (CARIN) model (Gagné & Shoben, 1997). According to CARIN, interpreting a combination involves identifying a relation that links the two concepts. Thus, *tourist castle* is formed using the relation *head* FOR *modifier* (as it means a castle for tourists), *mountain castle* is formed using the relation *head* LOCATED AT *modifier* (as it means a castle located on a mountain), and *stone castle* is formed using the relation *head* MADE OF *modifier* (as it means a castle made of stone). The theory makes two basic claims. First, the relation is a *bound representation* (Estes & Jones, 2006), with no independent existence. It is captured within the representation of the modifier, but not the representation of the head noun. For example, the FOR relation is associated with *tourist* in *tourist castle*; the LOCATED AT relation is associated with *mountain* in *mountain castle*; and so on. Second, several relations become activated and compete for selection during comprehension (Gagné & Shoben, 1997). The modifier is primarily responsible for interpretation, with more frequent relations for a given modifier receiving higher activation than less frequent relations for that modifier. For example, the modifier *mountain* is more frequently associated with the LOCATED AT relation than the ABOUT relation. Hence, *castle* LOCATED AT *mountain* would be more highly activated than *castle* ABOUT *mountain*. In contrast, the head plays a secondary role; it is used to decide whether a particular relation is plausible (i.e., compatible with existing knowledge; Gagné, 2002). Gagné and Shoben proposed that knowledge of a concept includes information about the frequency of these relations, such that the conceptual representation of every noun indicates the relations that it prefers when it serves as a modifier, but not the relations that it prefers when it serves as a head.

The relations assumed by CARIN form a small set of highly generalized types, adopted from the linguistic typology of Levi (1978). The most appropriate taxonomy of basic linking relations has been the subject of ongoing debate, with Levi's approach being criticized on the grounds that the relations are vague and underspecified (see Downing, 1977; Estes, 2003; Estes & Jones, 2006; Maguire, Devereux, Costello, & Cater, 2007; Murphy & Wisniewski, 2006; Wisniewski & Murphy, 2005). For example, the combination *daisy chains* purportedly instantiates the relation MADE OF, but this phrase could equally be connected by the relation IS. Moreover, Downing (1977) argued that Levi's thematic relations do not capture the non-linguistic knowledge that listeners bring to bear when interpreting noun–noun combinations. We return to the issue of relations in the General Discussion.

Gagné and Shoben (1997) provided some support for the CARIN model. Selecting 91 modifiers and 91 heads from Levi (1978), they judged 3239 combinations (of the 8281 possible pairings) as having sensible meanings. They ascertained which of the set of 15 relations were more or less likely to occur with particular modifiers and heads, deriving the relation frequency for each modifier and each head by calculating the percentage of combinations involving that noun and relation. They then had participants evaluate the sensicality of combinations instantiating relations of high (H) versus low (L) frequency for modifier and head respectively. For example, *chocolate utensils* is HH (i.e., high modifier, high head), because its interpretation is *utensils* FOR *chocolate*, and both *chocolate* as a modifier and *utensils* as a head frequently instantiate the FOR relation. In contrast, *chocolate rabbit* is HL, because its interpretation is *rabbit* MADE OF *chocolate*, and the MADE OF relation is frequent for the modifier *chocolate* but infrequent for the head *rabbit*. Finally, *chocolate plant* is LH, because its interpretation is *plant* MAKES *chocolate*, and the MAKES relation is infrequent for the modifier *chocolate* but frequent for the head *plant*. Gagné and Shoben found faster response latencies for HH and HL combinations than LH combinations, but no reliable difference between the HH and HL combinations. Thus, the frequency of the relation associated with the modifier affected ease of processing, but the frequency of the relation associated with the head noun did not. More recently, Gagné and Shoben (2002) suggested that the activation of relations does not depend on the association with the modifier, contrary to CARIN (Gagné & Shoben, 1997). Importantly, however, Gagné and Shoben (2002) did not present a revision of CARIN that makes predictions about the effects of relations. Hence, we interpret CARIN as referring to Gagné and Shoben's (1997) proposals.

CARIN contrasts with *schema* theories, including the *selective modification model* (Smith, Osherson, Rips, & Keane, 1988), the *concept specialization model* (Murphy, 1988, 1990), and *dual-process theory* (Estes, 2003; Wisniewski, 1996, 1997). In these accounts, a concept is seen as a structured set of dimensions (or slots) and values (or fillers) for those dimensions (Rumelhart, 1980). In the experiments below, we focus on the process of relation linking, by which different thematic roles are assigned to the modifier and head constituents in a combination: In *mountain castle*, *mountain* fulfils the role of location, whereas *castle* is the located object.

In schema theories, modification of a concept involves altering the schema by instantiating a new value for a given slot. In adjective–noun combinations such as *green apple*, the adjective *green* straightforwardly matches a slot (i.e., COLOR) in the head noun's schema and instantiates the value for that slot. In contrast, relation linking in noun–noun combinations is heavily

dependent on general knowledge and local context, as is apparent when considering the possible interpretations of *pie apple* (Gerrig & Bortfeld, 1999; Gerrig & Murphy, 1992; Medin & Shoben, 1988; Murphy, 1988, 1990). Because nouns are conceptually richer than adjectives, they do not have a single salient dimension to automatically match a slot of the head noun's schema. Listeners therefore need to integrate a nominal modifier with its head to arrive at the most plausible interpretation for a given combination.

According to Estes and Jones (2006), interpreting a conceptual combination involves activating a linking relation from a set of possible relations that form part of the semantic network. Crucially, in their model the relations are not bound to modifiers (or indeed to heads) but are *independent representations*. They constitute representational structures in their own right, which are applied to specific combinations as appropriate. Neither the modifier nor the head is privileged in the process of combining two concepts. This contrasts with CARIN, which proposes that relations are specifically bound to modifiers and have no independent existence.

Distinguishing accounts of conceptual combination using priming

A number of studies have used priming to examine the interpretation of combinations with fixed interpretations (i.e., where only one interpretation is plausible), with mixed results. Gagné (2001) reported evidence for relation and modifier-repetition priming in a task where participants judged whether combinations made sense: Participants comprehended target combinations such as *murder film* faster and more accurately following combinations using the same modifier and relation (*murder investigation*) compared to the same modifier and a different relation (*murder attempt*), and following combinations using the same modifier (*murder attempt*) compared to a different modifier (*vocal range*). But a comparable experiment using head repetition revealed repetition priming but no relation priming: *Poverty film* and *foreign film* primed *murder film* equally. These data are clearly compatible with relations being specifically associated with modifiers (rather than heads).

However, there is also evidence of priming in the absence of head or modifier repetition. Gagné (2002) found relational priming when prime and target used semantically related modifiers (e.g., *scholar accusation—student vote*) but no such priming when they did not (*surgery remedy—oil treatment*). This demonstrates that priming does not require modifier repetition. In addition, Gagné and Shoben's (2002) results go against CARIN's predictions. They had participants verify definitions of ambiguous combinations and found that repeating the head and repeating the modifier facilitated response times and response accuracy to an equivalent extent. They pointed out that these data contrast with Gagné and Shoben's (1997) data on frequency effects, and suggested that different mechanisms may underlie long-term versus short-term effects of head and modifier repetition in relation selection. Gagné and Shoben (2002) concluded that their data are not consistent with CARIN and that the theory needs modification to deal with short-term effects based on head repetition. But they did not specify how such a modification might be realized (and in any case this would not explain the pattern of effects reported in Gagné, 2001). Most importantly, any such account would abandon the central representational claim of CARIN, that relations are specifically associated with modifiers.

Other studies have shown that participants interpret novel combinations using the same relation as they have used to interpret a previous combination that involved a different head and modifier (Estes, 2003; Gerrig & Murphy, 1992; Spellman, Holyoak, & Morrison, 2001; Wisniewski & Love, 1998). For example, Estes found that combinations such as *pancake spatula* (meaning *spatula* FOR *pancake*) were interpreted more quickly following *bacon tongs* (meaning *tongs* FOR *bacon*) than following *city riots* (not meaning *riots* FOR *city*). However, Gagné (2000) failed to replicate Wisniewski and Love's (1998) priming effects using the same materials. Additionally, Gagné, Spalding, and Ji (2005) argued that Estes's (2003) results might reflect semantic priming rather than relation priming, due to the semantic similarity between prime and target combinations in the same-relation condition (e.g., *pancake* and *bacon* are more similar than *pancake* and *city*; similarly, *spatula* and *tongs* are more similar than *spatula* and *riots*). In accord with this, Gagné et al. found no priming in the absence of repeated constituents when the semantic similarity of the constituent words was controlled. However, Estes and Jones (2006) demonstrated relation priming between semantically dissimilar phrases (e.g., *wheat bread—copper monkey*). Moreover, they found that Gagné et al.'s same-relation conditions often used primes and targets whose relations differed considerably (e.g., *bear paw* served as a prime for *honey soup*), and argued that relation priming could not be explained as semantic priming.

In summary, there is contradictory evidence concerning the determinants of relation selection in concept combination and the putative representational asymmetry between the modifier and head noun. Some of the evidence suggests that concept combination involves selecting conceptual relations that are in some sense anchored to the modifier, and hence supports the CARIN model; other evidence suggests that it involves selecting conceptual relations that are independent of the modifier, and hence may be compatible with schema-based accounts such as that of Estes (2003).

Priming with expression-picture matching

Because previous priming studies have produced conflicting results, we turn to expression-picture matching to help us distinguish schema- and relation-based accounts of concept combination. Rather than consider the time-course of comprehension of (relatively) unambiguous expressions, we consider the process by which people select between interpretations for expressions that could instantiate different relations. Studies of syntactic comprehension (parsing) have shown that the interpretation of syntactically ambiguous sentences can be affected by prior comprehension of sentences with one or other syntactic structure (Carey, Mehler, & Bever, 1970). More recently, Branigan, Pickering, and McLean (2005) presented participants with an ambiguous prime expression, such as *The policeman prodding the doctor with the gun*, which is ambiguous between meaning that the policeman used the gun to prod the doctor (*verb-attachment analysis*) or the policeman prodded the doctor who had the gun (*noun-attachment analysis*). Participants then saw two pictures, one which matched one or other interpretation and one which matched neither interpretation, and selected the appropriate picture. They were then presented with a target expression, such as *The waitress prodding the clown with the umbrella*, which contains the same ambiguity as the prime. They then saw two pictures, but this time each picture corresponded to one interpretation of the target sentence. When prime and target shared the same verb, participants tended to choose the picture that matched the analysis (i.e., verb or noun attachment) assigned to the prime sentence. In contrast, when prime and target used different verbs, participants did not (significantly) tend to choose the picture that matched the analysis assigned to the prime (and priming was significantly stronger when the verb was shared than when it was not).

We applied the same method to examine whether the process of conceptual combination could be primed. Participants were presented with a prime expression that was consistent with two interpretations involving different relations; they then saw two pictures, one which matched one or other interpretation and one which matched neither interpretation, and selected the appropriate picture. They were next presented with a target expression that was consistent with two interpretations involving different relations; after this they saw two pictures, one of which corresponded to one relation, while the other corresponded to the alternative interpretation. We limited our stimuli to combinations that clearly involved a relation between the two nouns. There is ongoing debate as to whether property compounds (e.g., *zebra clam*) are interpreted using a different set of processes than relational compounds (Estes, 2003; Wisniewski, 1996, 1997) or whether both kinds of interpretations can be accounted for within a relation-based framework (Gagné, 2000). Hence, we did not address property compounds in this paper.

Our experiments focused on two of the 15 relations that are assumed by CARIN (numbers 3 and 4 in the table in Gagné & Shoben, 1997, p.72), which are themselves derived from Levi's (1978) typology, namely *head* HAS *modifier* (e.g., *picture book*) and *modifier* HAS *head* (e.g., *lemon peel*). Using these relations we constructed combinations that could plausibly instantiate both relations. For example, the combination *dog T-shirt* could be interpreted as meaning *dog* HAS *T-shirt*, analogous to *book* HAS *picture* in the example above, or as *T-shirt* HAS *dog*, analogous to *lemon* HAS *peel*. These two distinct interpretations of *dog T-shirt* could be depicted respectively as a dog wearing a T-shirt or as a T-shirt decorated with a picture of a dog (see Fig. 1). In the former case, the modifier is in a possessor relation with the head; in the latter, the modifier acts as descriptor of the head. Since these instantiations are subclasses of CARIN's more general *head* HAS *modifier* and *modifier* HAS *head* relations, we refer to them henceforth as POSSESSOR and DESCRIPTOR respectively. Notice that although the English paraphrases for both relations rely on the verb *has*, this is an idiosyncrasy of the English language. The paraphrases do not make reference to a single HAS relation: In German, for example, the paraphrases are *Hund* HAT *T-shirt* AN (verb: *anhaben*) and *T-shirt* HAT *Hund* DRAUF (verb: *draufhaben*).

One important reason for using two relations was that the form of the target could not be predicted from the form of the prime. When using 15 relations with half of the targets having the same relation as the prime, a target involves the same relation as the prime 50% of the time but any other specific relation only 3.6% (i.e., 50/14) of the time. But when using two relations with half of the targets having the same relation as the prime, a target involves the same relation as the prime 50% of the time and the alternative relation 50% of the time.

Let us now consider how processing of a combination such as *dog scarf* might be affected by prior processing in the models of conceptual combination outlined above. In CARIN, relations are exclusively linked to modifier concepts, so relational priming is predicted to occur when the modifier is repeated. Thus, people should be more likely to interpret *dog scarf* as a scarf decorated with a picture of a dog (i.e., *dog* DESCRIBES *scarf*) after interpreting *dog T-shirt* as a T-shirt decorated with a picture of a dog (i.e., *dog* DESCRIBES *T-shirt*) than after interpreting it as a T-shirt worn by a dog (i.e., *dog* POSSESSES *T-shirt*). In contrast, the CARIN model predicts no tendency to repeat relations when the head is repeated but not the modifier. Thus, people should not be more likely to interpret *dog scarf* as a scarf decorated with a picture of a dog (i.e., *dog* DESCRIBES *scarf*) after interpreting *rabbit scarf* as a

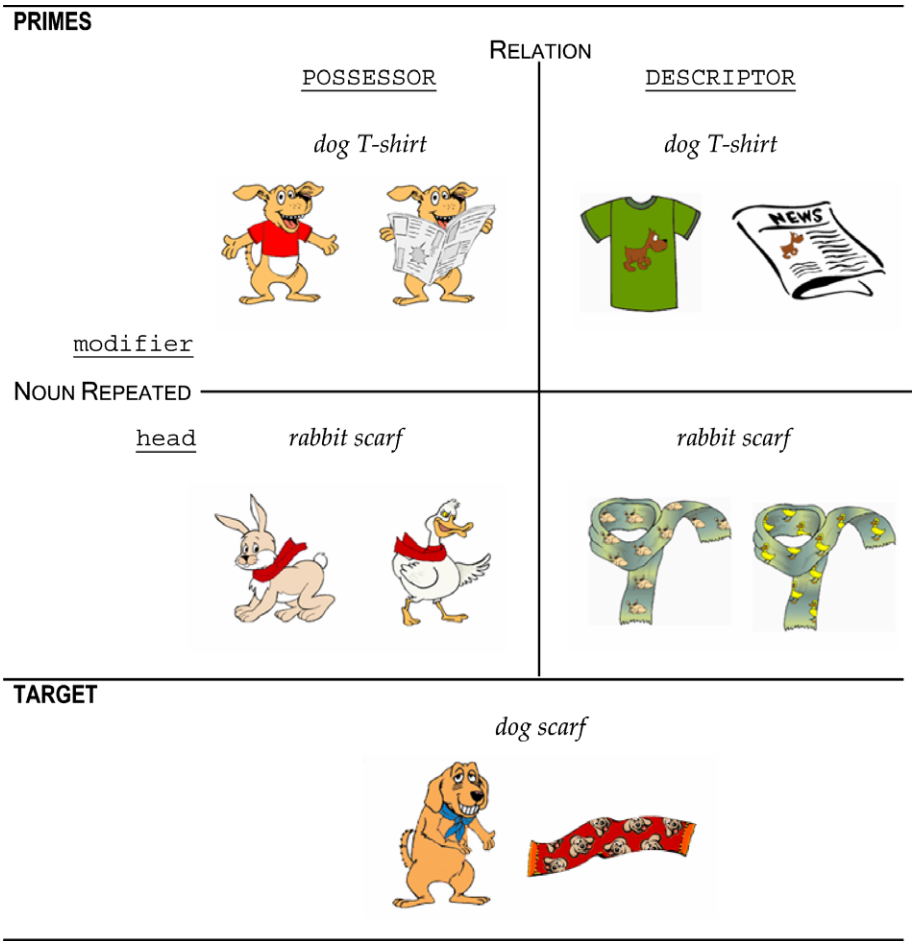

Fig. 1. Sample stimuli for Experiments 1 and 2.

scarf decorated with a picture of a rabbit (i.e., *rabbit* DESCRIBES *scarf*) than after interpreting it as a scarf worn by a rabbit (i.e., *rabbit* POSSESSES *scarf*). Similarly, there should be no tendency to repeat relations when neither the head nor the modifier is repeated (e.g., *rabbit hat* should not relationally prime *dog scarf*).

In contrast, early schema models (e.g., Murphy, 1988, 1990; Smith et al., 1988; Wisniewski, 1996, 1997) do not make clear predictions about relation priming. However, Estes and Jones (2006) represent relations independently of concepts. Their model predicts that priming should occur whenever a relation is repeated, because prior activation of a relation should facilitate its subsequent selection. Hence, unlike CARIN, their model predicts that people should be more likely to interpret *dog scarf* as a scarf decorated with a picture of a dog (i.e., *dog* DESCRIBES *scarf*) after interpreting *rabbit scarf* as a scarf decorated with a picture of a rabbit (i.e., *rabbit* DESCRIBES *scarf*) than after interpreting it as a scarf worn by a rabbit (i.e., *rabbit* POSSESSES *scarf*). Furthermore, this tendency to repeat relations should occur when neither the head nor the modifier is repeated (e.g., *rabbit hat* should relationally prime *dog scarf*), and should be as strong when neither the head nor the modifier is repeated as when either the head or the modifier is repeated.

Although Branigan et al. (2005) did not find significant priming when prime and target used different verbs, studies in language production have shown syntactic priming in the absence of verb repetition (e.g., Bock, 1986). For example, participants tend to produce passives more often after passive primes than after denotationally equivalent active primes. Indeed, there is some evidence that abstract aspects of semantic structure can also be primed in production and comprehension with little or no lexical repetition (Bock, Loebell, & Morey, 1992; Garrod & Anderson, 1987; Watson, Pickering, & Branigan, 2004). Such data are compatible with the existence of independent syntactic and semantic representations, just like the independent

relational representations postulated by Estes and Jones (2006). But although syntactic priming in production does not require lexical repetition, it is considerably enhanced by such repetition (Branigan, Pickering, & Cleland, 2000; Cleland & Pickering, 2003; Pickering & Branigan, 1998). This phenomenon is referred to as the *lexical boost*. One explanation for why there is clear evidence of relational priming when the modifier is repeated, but less clear evidence of relational priming in the absence of repetition, is that relational priming occurs in both cases but repetition enhances priming, just as it appears to do in the syntactic domain. In short, an account based on the analogy to syntactic priming predicts priming whether prime and target involve same or different constituents, but enhanced priming (i.e., a lexical boost) when one or other constituent is repeated.

We now report three experiments that tested these accounts. On each trial, participants first read a noun–noun phrase (e.g., *rabbit scarf*), then saw two pictures and had to select which picture (left or right) matched the phrase. On prime trials, one picture depicted either the POSSESSOR or DESCRIPTOR interpretation, while the other picture depicted neither interpretation. The 'correct' picture on prime trials thus disambiguated the appropriate analysis of the prime phrase as *head* DESCRIBES *modifier* (a rabbit wearing a scarf) or as *head* POSSESSES *modifier* (a scarf decorated with rabbits). To select the appropriate picture, participants therefore had to resolve the prime phrase as involving the POSSESSOR interpretation (for POSSESSOR primes) or as involving the DESCRIPTOR interpretation (for DESCRIPTOR primes). Target trials did not disambiguate the associated phrase; both pictures depicted possible interpretations of the combination, with one corresponding to the POSSESSOR interpretation and the other to the DESCRIPTOR interpretation. Our dependent measure was choice of target picture. Hence we examined whether participants interpreted the target phrase in the same way in which they had interpreted the prime phrase.

Experiments

Experiment 1

Experiment 1 investigated whether relational priming occurred in the context of head repetition, modifier repetition, or both, and allowed direct comparison of the effects of head and modifier repetition.

Methods

Participants. Thirty-two undergraduate students from the University of Edinburgh community were paid to participate. They were all native English speakers and had no reading difficulties.

Items

We constructed 32 item sets (see Appendix). Each comprised a prime combination instantiating a POSSESSOR or DESCRIPTOR relation (such as *dog T-shirt*), a pair of prime pictures, a target combination (such as *dog scarf*), and a pair of target pictures. The first noun in the phrase was always animate; the second noun was always inanimate. Of the prime pictures, one matched either the POSSESSOR or the DESCRIPTOR interpretation of the relation between the two constituents, while the other matched neither interpretation. In addition, primes contained either the same modifier or the same head as the target expression. Thus primes were in four conditions, as shown in Fig. 1. The target pictures corresponded to each interpretation of the expression; one was the POSSESSOR depiction, the other the DESCRIPTOR depiction.

In the PM (i.e., POSSESSOR, modifier repeated) condition, a prime expression such as *dog T-shirt* was displayed with the two pictures in the top-left quadrant of the primes in Fig. 1. In the PH (POSSESSOR, head repeated) condition, a prime expression such as *rabbit scarf* was displayed with the pictures in the bottom-left quadrant of the primes in Fig. 1. Pictures in the DM (DESCRIPTOR, modifier repeated) and DH (DESCRIPTOR, head repeated) conditions are shown on the top-right and bottom-right quadrants respectively of the primes in Fig. 1. In each quadrant, the match picture is shown on the left and the distracter picture is shown on the right (this was of course counterbalanced in the experiment). For both modifier-repeated conditions the distracter pictures contained the same modifier as the correct prime picture but a different head (in this case *dog newspaper*). Similarly, distracters in the head-repeated conditions contained the same head as the prime but a different modifier (in this case *duck scarf*). The target expression *dog scarf* was matched with the two pictures at the bottom of Fig. 1, corresponding to the POSSESSOR and DESCRIPTOR interpretations of the phrase.

In addition, we constructed 96 unambiguous filler expressions and pairs of pictures, one of which matched the expression presented. The set of fillers comprised 32 singular nouns such as *dancer*, *goat*, and *candle*, 32 plural nouns such as *lobsters*, *drums*, and *staplers*, and 32 conjoined singular noun phrases such as *teapot and bucket*, *architect and butcher*, and *mask and sharpener*. Distracter pictures were of the same type, also depicting singular, plural, or conjoined noun phrases. Half of the filler objects were animate and the other half were inanimate. In addition, half of the filler objects appeared as objects in the experimental item set and the other half did not. Those that appeared in the experimental set

appeared in distracter (i.e., non-match) pictures in the filler set. The filler objects that did not appear in the experimental set appeared as matches half of the time and as distracters the other half of the time.

The experimental items were organized into four lists, each including eight items per condition, such that one version of each item appeared in each list. Each list of 160 noun phrases and pictures (32 primes, 32 targets, and 96 fillers) was individually randomized for each participant, with the constraints that each prime immediately preceded the associated target and that at least two filler trials intervened between experimental trials.

Procedure

Participants viewed the materials on a computer screen, which was connected to a PsyScope button box. E-Prime software was used to present the experiment and record the data. There were seven practice trials before the beginning of the experiment, comprising two disambiguated (prime-type) POSSESSOR and DESCRIPTOR combinations and five filler-type phrases. A single trial (prime, target, or filler) comprised a phrase displayed in the centre of the screen for 2000 ms, a blank screen for 500 ms, and then a pair of pictures. Participants were instructed to match each phrase with the corresponding correct picture. This involved deciding on the appropriate match and pressing the left and right keys on the button box, which signaled the end of that trial and the beginning of the next. Each phrase was viewed once. For each list, the matching picture for the phrase appeared on the left for half the trials and on the right for the other half. Similarly, the POSSESSOR picture for the target appeared on the left for half of the trials and on the right for the other half.

Coding and analysis

Each participant was presented with 32 targets, comprising eight in each of the four priming conditions (PM, PH, DM, DH). Each target expression and corresponding pair of pictures were presented to all 32 participants, such that eight participants saw any one version of an experimental item.

Any trial on which a participant selected the incorrect prime picture was eliminated from the analysis. The remaining responses were coded according to whether the corresponding target selected was a POSSESSOR picture or a DESCRIPTOR picture. We calculated the probability of choosing a POSSESSOR target picture after choosing a POSSESSOR prime picture, and the probability of choosing a POSSESSOR target picture after choosing a DESCRIPTOR prime picture. (Note that the choice of POSSESSOR responses is arbitrary because the proportion of DESCRIPTOR responses is complementary: Participants always selected either a POSSESSOR or a DESCRIPTOR picture.) Hence we calculated proportions of targets in both the participants and items analyses by dividing the number of POSSESSOR pictures selected following a POSSESSOR prime by the sum of POSSESSOR targets following POSSESSOR primes and POSSESSOR targets following DESCRIPTOR primes. This measure is analogous to that used by Branigan et al. (2005). We used this measure because it allowed us to compare priming between conditions in cases where the proportions of correctly selected prime pictures were not equivalent and obviated computing separate analyses for the (non-independent) POSSESSOR and DESCRIPTOR targets. Table 1 reports the proportions of POSSESSOR targets in each condition for all experiments. ANOVAs were carried out on these data with Relation (POSSESSOR vs. DESCRIPTOR) and Noun (modifier-vs. head-repeated) as factors, based on participant ($F1$) and item ($F2$) variability. Both factors were treated as within participants and items. The analyses were combined to compute the MinF' statistic (Clark, 1973). Table 2 reports the results of the analyses for all experiments. Confidence intervals (based on participant variability) for the differences between means are reported for all significant effects (following Masson & Loftus, 2003).

Results

Correct prime pictures were selected on 982 trials (96%). Of these, 241 (25%) were PM trials, 254 (26%) were PH trials, 240 (24%) were DM trials, and 247 (25%) were DH trials. On the corresponding target trials for these 982 trials, participants chose 309 (31%) POSSESSOR and 673 (69%) DESCRIPTOR pictures respectively.

Table 1 shows the proportions of POSSESSOR target responses for all three experiments, represented graphically in Fig. 2. In Experiment 1, two-way ANOVAs revealed a main effect of relation. There was a significant overall tendency to repeatedly interpret combinations in the same way, manifested as a relation priming effect of 21%; that is, participants made 21% more target responses that were of the same type (POSSESSOR or

Table 1
Mean POSSESSOR target proportions for responses in each condition for all experiments (based on participants analyses)

	POSSESSOR prime	DESCRIPTOR prime
Experiment 1		
Modifier repetition	0.38	0.19
Head repetition	0.46	0.24
Experiment 2		
No repetition	0.24	0.18
Experiment 3		
No repetition	0.20	0.12

Table 2
Analysis of variance summary for all experiments

Effect	By participants ($F1$)			By items ($F2$)		Min F'	
	df	$F1$	*MSE*	*df*	$F2$	*df*	min F'
Experiment 1							
Relation	1, 31	24.02**	0.06	1, 31	49.68**	1, 55	16.19**
Noun	1, 31	7.59*	0.02	1, 31	6.54*	1, 62	3.51
Relation × Noun	1, 31	<1		1, 31	<1	1, 38	<1
Experiment 2							
Relation	1, 31	4.51*	0.01	1, 31	4.50*	1, 62	2.26
Experiments 1 and 2							
Relation	1, 62	28.50**	0.02	1, 31	41.71**	1, 91	16.93**
Experiment	1, 62	4.91*	0.08	1, 31	18.15**	1, 62	3.86^
Relation × Experiment	1, 62	8.05**	0.02	1, 31	9.67**	1, 87	4.39*
Experiment 3							
Relation	1, 23	5.93*	0.01	1, 31	8.84**	1, 48	3.55

* $p < .05$.
** $p < .01$.
^ $p < .06$.

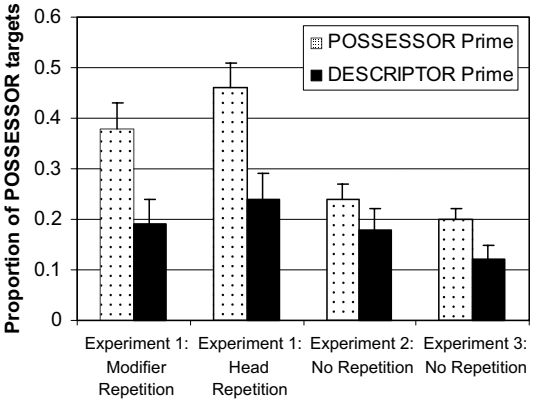

Fig. 2. POSSESSOR target proportions across all experiments and across all conditions, with error bars representing the standard error of the mean (based on participants analyses).

DESCRIPTOR) as the prime response than target responses that were of the alternative type to the prime response (CI = 0.07). In addition, there was a significant effect of noun; participants chose 7% more POSSESSOR interpretations when the head noun was repeated than when the modifier was repeated (CI = 0.04). There was no interaction between noun and relation (both Fs < 1), indicating that priming did not differ whether the modifier or the head noun was repeated. Priming occurred both when the modifier was repeated (CI = 0.06) and when the head was repeated (CI = 0.07).

It is conceivable that participants might have detected the relationship between matching prime-target pairs and hence begun to employ a strategy in choosing between the two pictures depicting the ambiguous target combinations. If they did so, we would predict that the tendency to repeatedly choose the same relation would increase through the experiment. To test this, we compared the magnitude of priming in the two halves of the experiment by introducing half (first vs. second half) as a within-participants and -items factor. Half was defined as the first or last 16 prime-target trials seen by a participant. For each half, the proportion of targets that were interpreted in the same way as the preceding prime was calculated by first excluding trials on which the prime picture was incorrectly selected, and then dividing all target responses that involved the same relation as the immediately preceding prime expression by the sum of all responses. In fact, the analysis revealed that the magnitude of priming was stable across the two halves of the experiment (both Fs < 1), with 59% of target responses involving the same relation as prime responses in the first half of the experiment, and 59% of target responses involving the same relation as prime responses in the second half of the experiment. Hence we can be confident that our pattern of effects does not reflect the development of a strategy for approaching the task over the course of the experimental session.

Discussion

In Experiment 1, participants were presented with novel ambiguous nominal combinations, disambiguated by pictures to involve either a *modifier* HAS *head* (POSSESSOR) or *head* HAS *modifier* (DESCRIPTOR) interpretation. When selecting between target pictures, participants chose DESCRIPTOR relations more frequently than POSSESSOR relations. More importantly, however, participants tended to interpret combinations

in the same way as they had interpreted a preceding (disambiguated) prime. That is, they were more likely to select a POSSESSOR picture after seeing a POSSESSOR picture than after seeing a DESCRIPTOR picture. These results are consistent with previous findings of relation priming (e.g., Estes, 2003; Estes & Jones, 2006; Gerrig & Murphy, 1992; Wisniewski & Love, 1998) and extend demonstrations of relation priming to an expression-picture matching methodology.

Experiment 1 showed that relation priming occurred regardless of whether it was the head or the modifier that was repeated. Moreover, there was no indication that priming in the context of modifier repetition was any stronger than priming in the context of head repetition. These findings are incompatible with the predictions of CARIN, according to which relational information is associated with the modifier but not with the head. In contrast, they are compatible with schema-based models such as that of Estes and Jones (2006), in which relations are represented independently of the head and modifier concepts. There was also a reliable tendency to use more POSSESSOR descriptions when the head noun was repeated than when the modifier was repeated. Neither account specifically predicts this effect, but it is not inconsistent with them.

A further test of the two accounts would investigate priming in the absence of lexical repetition between prime and target. According to CARIN, priming should not occur under these circumstances because relation retrieval is dependent on modifier retrieval. According to Estes and Jones (2006), priming should occur because relation retrieval is independent of head or modifier retrieval. In Experiment 2, we therefore presented participants with primes that involved one of two semantic relations (POSSESSOR vs. DESCRIPTOR) followed by a choice of two target pictures, one of which corresponded to the same relation and the other of which corresponded to the alternative relation, but neither of which contained either of the two nouns that had appeared in the prime.

Experiment 2

Experiment 2 investigated whether priming occurred in the absence of repeated constituents between prime and target. In other respects, it was identical to Experiment 1.

Methods

Thirty-two further participants from the same population were paid to participate. The pictures (including the fillers) for this experiment were identical to those in Experiment 1. However, we rotated the assignment of prime to target combinations such that neither the head nor the modifier was repeated across prime and target. Each item comprised a prime combination such as *dog T-shirt* together with a pair of pictures, one of which depicted either the POSSESSOR or DESCRIPTOR interpretation of the expression and one of which depicted neither interpretation of the expression; and a target combination such as *dinosaur flag*, together with pictures depicting each of the two interpretations of the phrase. The procedure, coding, and analysis were the same as in Experiment 1.

Results and discussion

Correct prime pictures were selected on 982 trials (96%). Of these, 489 (50%) were POSSESSOR trials and 493 (50%) were DESCRIPTOR trials. On the corresponding target trials for these primes, participants chose 201 (20%) POSSESSOR and 781 (80%) DESCRIPTOR pictures respectively. One-way ANOVAs revealed that participants selected 6% more POSSESSOR target pictures following POSSESSOR prime pictures than following DESCRIPTOR prime pictures, with a confidence interval of 0.03 (see Tables 1 and 2).

As in Experiment 1, participants chose DESCRIPTOR relations more frequently than POSSESSOR relations. More importantly, we found evidence for relational priming in the interpretation of nominal combinations: Participants tended to select the same interpretation for an ambiguous combination that they had just used on a preceding prime trial. However, Experiment 2 demonstrated that these effects occurred in the absence of any lexical repetition between prime and target. That is, interpretation was facilitated on the basis of repetition of the semantic relation alone, when neither the modifier nor the head was repeated. This result is inconsistent with CARIN's assumption of modifier-dependent relational representation (Gagné & Shoben, 1997). Rather, it offers support for Estes and Jones's (2006) independent representation model.

Comparison of Experiments 1 and 2

Estes and Jones (2006) did not directly address the issue of the effects of lexical repetition (of the head or modifier) on relation priming, but their account provides no reason to assume a lexical boost. In contrast, if relation priming is similar to syntactic priming (e.g., Cleland & Pickering, 2003), lexical repetition should enhance relation priming. To contrast these accounts, we conducted a combined analysis of Experiments 1 and 2 in order to determine whether relation priming was enhanced by lexical repetition. Given that modifier-repeated and head-repeated conditions in Experiment 1 showed nearly identical results and Experiment 2 contained no comparable distinction, we collapsed across these two conditions for the purposes of the following analysis. Accordingly, Relation (POSSESSOR vs.

DESCRIPTOR prime) × Experiment (repetition vs. no repetition) ANOVAs were carried out on the POSSESSOR target proportions (see Tables 1 and 2). Relation was within participants and items; Experiment was between participants but within items. The analysis revealed a main effect of relation: Across experiments there was a 12% greater tendency to select target responses that were of the same type (POSSESSOR or DESCRIPTOR) as the prime response than target responses that were of the alternative type to the target response. More importantly, there was an interaction between Relation and Experiment: The 21% priming effect in Experiment 1 (i.e., when head or modifier was repeated) was significantly larger than the 6% priming effect in Experiment 2 (i.e., when neither head nor modifier was repeated). In other words, lexical repetition enhanced priming.

This result is not predicted by Estes and Jones's (2006) independent representation model, according to which there should be no difference in priming effects in the presence of repetition compared with priming effects in the absence of repetition. Since conceptual relations are independently primed in this account, the level of activation that a given relation receives should be the same whether or not there is overlap of lexical items between prime and target. Hence, although the results of Experiments 1 and 2 argue against CARIN, they also suggest that Estes and Jones's (2006) independent model does not fully account for processes of conceptual combination.

Experiments 1 and 2 demonstrated a reliable tendency for participants to choose a picture to match an ambiguous combination that involved the same relation as they had used to match a combination to a picture on a preceding trial. We have interpreted this finding in terms of a linguistic priming effect centered on repetition of a particular relation. However, in both experiments the different interpretations of the ambiguous target combination were associated with visually distinct pictures; one in which the modifier was the visually most salient (e.g., in most cases, the largest) element in the picture (the POSSESSOR interpretation), and one in which the head was the visually most salient element (the DESCRIPTOR interpretation). Depending on the prime condition, the match (i.e., 'correct') prime picture would similarly contain either a visually salient modifier (for POSSESSOR primes) or a visually salient head (for DESCRIPTOR primes). Recall that in all experimental materials the modifier was animate and the head was inanimate. Hence, the priming effect found in Experiments 1 and 2 could reflect facilitation of perceptual processing by prior processing of a visual stimulus with similar salient visual characteristics: Participants might have preferred to choose similar over dissimilar pictures, without accessing the conceptual relation instantiated by the combination and depicted in the associated picture. For example, they might have preferred to choose a picture with a salient animate entity (modifier) after choosing a picture with a salient animate entity than after choosing a picture with a salient inanimate entity (head), or vice versa. That is, they would show a greater preference for choosing a picture of a tortoise after choosing a picture of a dog than after choosing a picture of a hat.

If such a visual priming explanation is correct, then the effect is dependent upon the match prime picture depicting a single relation. Hence this explanation predicts that the priming effect should disappear if the match prime picture depicts both relations, such that both alternative target pictures share common visual elements with the match prime. If, however, the priming effect is a linguistic effect based upon retrieval of the activated linking relation, then exposure to pictures containing both (POSSESSOR and DESCRIPTOR) relations *within the same picture* should yield relation priming. Experiment 3 was designed to distinguish between these two accounts. We used the conditions in which neither head nor modifier was repeated between prime and target, so that we could be more confident that relational priming in the absence of lexical repetition is robust.

Experiment 3

In Experiment 3 we used prime pictures that contained both POSSESSOR and DESCRIPTOR relations within the same picture. Fig. 3 shows that the picture associated with the POSSESSOR interpretation of *dog T-shirt* depicted a dog wearing a T-shirt with a picture of a rabbit on it. In the DESCRIPTOR prime the identical picture appeared, this time matching the phrase *rabbit T-shirt*. In other words, a dog wearing a T-shirt with a picture of a rabbit on it depicted the phrase *dog T-shirt* and the phrase *rabbit T-shirt*, for POSSESSOR and DESCRIPTOR interpretations respectively. The alternative relation in each case was also included in the distracter picture. Prime combinations were the same as those used in Experiments 1 and 2, with the exception of 12 items. For these items, we used the combination that had appeared as the distracter in the first two experiments (e.g., *flamingo plate*, which had been a distracter in Experiments 1 and 2, replaced *flamingo ice-cream*), on the grounds that in these cases the distracter combinations allowed for easier depiction of the two relations within the same picture. Target picture pairs were the same as in Experiment 1. As in Experiment 2, the item set was organized so that there was no lexical repetition between prime and target phrases. Our dependent measure was again the proportion of POSSESSOR targets selected on target trials.

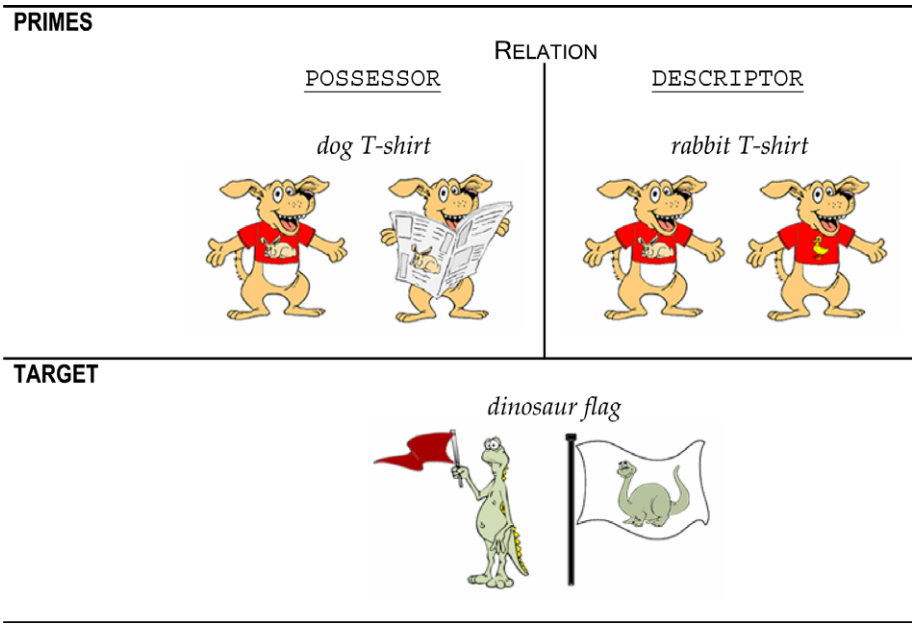

Fig. 3. Sample stimuli for Experiment 3.

Methods

Thirty-two further participants from the same population were paid to participate. The procedure, coding, and analysis were the same as in Experiments 1 and 2.

Results and discussion

Correct prime pictures were selected on 737 trials (96%). Of these, 362 (49%) were POSSESSOR trials and 375 (51%) were DESCRIPTOR trials. In these 737 trials, participants chose 116 (16%) POSSESSOR and 621 (84%) DESCRIPTOR pictures respectively.

As in Experiments 1 and 2, participants chose DESCRIPTOR relations more frequently than POSSESSOR relations (see Table 1). One-way ANOVAs on the proportions of POSSESSOR targets in each condition revealed a reliable effect of Prime (CI = 0.04); see Table 2. Participants produced reliably more POSSESSOR target interpretations after a POSSESSOR prime than after a DESCRIPTOR prime. The magnitude of priming was 8%.

In sum, there was significant relation priming when the prime picture contained both relations, arguing against a possible visual priming account. In selecting the correct prime picture to match the given phrase, participants were forced to assess both POSSESSOR and DESCRIPTOR relations; moreover, they could not choose a target picture based on visual similarity to the prime picture that they had previously chosen, because both alternative target pictures were visually similar (though in different ways) to the chosen prime picture. Hence our results suggest that the effect is linguistic in nature. Activation of the relevant relation, rather than visual properties of the pictures, led participants to reliably interpret the ambiguous target combination in the same way as they had interpreted the prime.

General discussion

Three forced-choice expression-picture matching experiments examined the interpretation of novel noun–noun combinations. They all showed reliable relation priming. Experiment 1 demonstrated that participants tended to interpret an ambiguous novel combination as involving the same relation as a disambiguated novel combination that they had just comprehended; this effect was stable over the course of the experiment. The priming effect was reliable and comparable in magnitude both when the head was repeated and when the modifier was repeated. Experiment 2 extended these findings by showing that priming occurred when both nouns in the prime differed from both nouns in the target. Experiment 3 demonstrated that priming without noun repetition did not depend upon visual similarities between the prime and target pictures.

Hence our results provide evidence of relation priming in the absence of lexical repetition. Previous empirical demonstrations of relation priming in the absence of lexical repetition have been open to alternative explanations. Gerrig and Murphy (1992)

demonstrated relation priming between novel prime and target compound nouns embedded in the context of narratives. However, Gagné et al. (2005) suggested that participants' interpretations of the novel targets might not have depended on the prime, but rather on the content of the story. Thus their effects need not have been due to the repetition of a conceptual relation. Gagné et al. also argued that Estes's (2003) apparent relation priming of novel combinations could be due to a design confound, because prime and target modifiers and heads were not controlled for semantic similarity across experimental conditions. They therefore suggested that what was reported as relation priming could not be distinguished from semantic priming.

The design of the current set of experiments addressed these two points. The semantic relations expressed between modifier and head in the novel combinations were depicted in pictures. Hence the possible interpretations were clear and there was no possibility of discourse cues affecting comprehension. Furthermore, by constructing the materials set around ambiguous combinations (word pairs that allow both the *head* POSSESSES *modifier* interpretation and the *head* DESCRIBES *modifier* interpretation, with the appropriate interpretation depending on the associated picture), identical phrases appeared across experimental conditions, ruling out the possibility of one condition being linguistically more similar to the target than the other. Hence we can be sure that the effects that we found reflect priming of relations and not of lexical representations.

The current findings also demonstrated relation priming both when the head and when the modifier was repeated, and suggested that priming was very similar in both cases. Although reliable differences in magnitude of priming occurred in the contrast between repetition and no repetition, they did not occur in the contrast between repeated modifier versus repeated head. This suggests that modifier concepts do not have a special role in the processing of combinations but rather play a comparable role to that of head concepts.

Taken together, these experiments cast new light on the influence of modifier and head constituents in the interpretation of novel combinations. Our results are not compatible with the CARIN model as detailed in Gagné and Shoben (1997). According to CARIN, modifiers have a privileged status: Relations are stored as part of each modifier's representation, such that processing a combination involves mediated retrieval of the relevant relation via the modifier with which it is associated; relations have no independent existence in the semantic network. Thus priming of a semantic relation should only occur if the modifier constituent, which acts as the key to relation retrieval, appears in both prime and target. But we found reliable priming when prime and target contained different modifiers. Our results therefore show no evidence for the special status of the modifier, suggesting that modifiers do not have greater access to relational information than do head nouns.

How can our results be reconciled with existing previous data on conceptual relations? First, our review of priming experiments shows many inconsistencies (e.g., Gagné, 2001 vs. Gagné & Shoben, 2002), so our results would necessarily not be compatible with all previous data. These inconsistencies may be due to problems with the set of relations assumed by CARIN. Estes (2003; Estes & Jones, 2006) argued that the generality and vagueness of the CARIN relations account for Gagné's (2001; 2002; Gagné et al., 2005) pattern of priming results. Gagné (2001) failed to show relation priming without lexical repetition, but Estes (2003) argued that the absence of these effects was due to the fact that prime and target relations in her Same Relation condition were not sufficiently similar. Supporting this explanation, relation priming was found in Gagné's (2002) study, in which prime and target relations were more closely matched. Furthermore, Estes and Jones (2006) argued that the same design flaw underlay the null result reported in Gagné et al. (2005), and concluded that relation priming depends in part on similarity of prime and target relations. A related critique is that the relations assumed by CARIN do not give credence to the richness and detail of the interpretations people derive from combinations (Downing, 1977; Murphy & Wisniewski, 2006; Wisniewski & Murphy, 2005). Our study avoided these issues by making use of narrow instantiations of just two of the CARIN relations. It also avoided the concern that participants could predict the form of the target from the form of the prime (see Introduction). Whereas using just two relations clearly limits the generality of our results, it sheds light on a confusing pattern of data with a novel technique.

Because of the inconsistencies in previous findings, it is perhaps not surprising that our results also clash with Gagné and Shoben's (1997) frequency data, which appear to show effects of the modifier but not effects of the head. However, recent work has questioned their results on methodological grounds (Maguire et al., in press; Murphy & Wisniewski, 2006; Storms & Wisniewski, 2005; Wisniewski & Murphy, 2005). One issue that all of these studies raised is that Gagné and Shoben's relation frequency measure for modifiers and heads is not representative of the true distributions of the relations that were used with the words in the stimulus set. Thus, Wisniewski and Murphy (2005) analyzed Gagné and Shoben's materials for plausibility (using participant ratings) and familiarity

(measured in numbers of Google search hits) and found that these two factors were correlated with the RT patterns in their results. That is, familiarity and plausibility were confounded with frequency of relation type.

In some respects, our results are compatible with Estes's (2003; Estes & Jones, 2006) model, which accords no privileged status to the modifier, and proposes that relations are represented independently in the semantic network. Processing a combination involves direct (unmediated) retrieval of the relevant relation. Because relations are independent, they can be primed through prior use in the absence of repetition of modifier or head, in keeping with our findings. Our results therefore suggest that when people process nominal combinations they retrieve a conceptual relation, and that this process of retrieval is at least partly independent of the retrieval of information associated with the modifier and the head. Critically, however, one aspect of our results is not predicted by Estes's model: We found stronger priming when either the head or the modifier was repeated than when neither the head nor the modifier was repeated. If relations were simply independently represented, there would be no reason to expect such a difference.

One possible explanation for the above finding is that the repetition of one of the constituent concepts between prime and target causes the instantiation of the linking relation (POSSESSOR or DESCRIPTOR) to be more similar than when neither concept is repeated. The combination *dog T-shirt* (depicted as a dog wearing a T-shirt) is different from *frog hat* (depicted as a frog wearing a hat), in as far as a T-shirt is worn differently to a hat. On such an account, the linking relations might not constitute a discrete set, but rather each instance of a relation would be more or less similar to other instances. The use of a set of relations in a given experiment or model would constitute an approximation to the real state of affairs. Such an account seems plausible for heads when paired with POSSESSOR relations, but is perhaps less plausible in other cases. For example, the modifier does not generally affect the meaning of the relation. In other words, the relations instantiated by *dog T-shirt* and *duck T-shirt* (depicted respectively as a dog wearing a T-shirt and a duck wearing a T-shirt) appear to be more similar than the relations instantiated by *dog T-shirt* and *dog scarf* (depicted respectively as a dog wearing a T-shirt and a dog wearing a scarf). Likewise, it is not clear whether the same clear differences hold for the DESCRIPTOR relation, for example, whether *bull mittens* (depicted as mittens with pictures of bulls on them) are different from *ladybird trainers* (depicted as trainers with pictures of ladybirds on them), in as far as decorating knitted mittens is different to printing pictures on trainers.

The findings of greater relation priming in the presence of lexical repetition converge with the findings for syntactic priming in language production, and suggest that a similar type of explanation may be appropriate. Such studies show priming in the absence of lexical repetition but enhanced priming in the context of such repetition (e.g., Branigan et al., 2000; Cleland & Pickering, 2003; Pickering & Branigan, 1998). Pickering and Branigan interpreted such results in terms of Levelt, Roelofs, and Meyer's (1999) model of lexical representation, in which lexical entries are split into a conceptual stratum (that captures the meaning of words), a lemma stratum (that captures their syntactic properties), and a word-form stratum (that captures their phonological properties). They constructed a network model of the lemma stratum in which the syntactic component of lexical entries (i.e., their lemmas) is associated with nodes corresponding to grammatical constructions. For example, transitive verbs such as *kick* and *drink* are linked to nodes corresponding to the active and passive constructions. Use of *kick* in the passive, for example, activates the *kick* node and the passive node, and their coactivation leads to a strengthening of the link between them. Activation does not decay immediately. The activation of the passive node thereby increases the chances of subsequently using a passive, and the strengthening of the link leads to a particularly strong tendency to subsequently use the passive with the verb *kick*.

We therefore propose a comparable account of relation priming, with the main difference being that we localize the account at the conceptual stratum rather than the lemma stratum. We assume that the concepts associated with each word are linked to particular relations. To capture the overall preference for DESCRIPTOR over POSSESSOR relations, the initial activation of the DESCRIPTOR(X, Y) node is greater than the activation of the POSSESSOR(X, Y) node (just as the initial activation of the ACTIVE node is greater than the initial activation of the PASSIVE node in English). Because relations constitute components of the meaning of lexical entries, the network forms part of the conceptual stratum rather than the lemma stratum, but the organizational principles are similar. Note that lexical entries are not decomposed into semantic features in Levelt et al. (1999) account. The associations between concepts and relational nodes hold equally for heads and modifiers, because the association between concepts such as DOG(X) or SCARF(X) and relations such as POSSESSSOR or DESCRIPTOR does not depend on how those concepts are used within an expression (i.e., whether they are used as heads or modifiers). Interpretation of *dog scarf* as meaning a dog wearing a scarf activates the DOG(X) node, the SCARF(X) node, and the POSSESSOR(X, Y) node.

The simultaneous activation of the three leads to a strengthening of the connections between DOG(X) and POSSESSOR(X, Y) and between SCARF(X) and POSSESSOR(X, Y). Activation of the POSSESSOR(X, Y) node increases the likelihood of using the POSSESSOR relation to link a subsequent combination (of two new concepts). The strengthening of the connections between DOG(X) and POSSESSOR(X, Y) further increases the likelihood of using the POSSESSOR relation if the DOG(X) concept is activated again, as in the subsequent comprehension of *dog T-shirt*. Likewise, the strengthening of the connections between SCARF(X) and POSSESSOR(X, Y) further increases the likelihood of using the POSSESSOR(X, Y) relation if the SCARF(X) concept is activated again, as in the subsequent comprehension of *rabbit scarf*.

In conclusion, relations appear to be represented in a partly independent and partly lexically bound manner, just as syntactic information appears to be (Pickering & Branigan, 1998). The lexically bound component does not appear to be modifier-specific. Instead, the associations between heads and relations and between modifiers and relations appear to be equally strong. We have argued that these findings can be integrated with accounts drawn from the literatures concerned with conceptual combination and psycholinguistics.

Appendix A. Items for Experiments 1 and 2

Targets for each experiment are separated by slashes. The first target was used in Experiment 1; the second target was used in Experiment 2.

Primes		Target
Modifier repeated	Head repeated	
Dog T-shirt	Rabbit scarf	Dog scarf/alligator trainers
Bear mug	Doll umbrella	Bear umbrella/ant chair
Dinosaur lamp	Porcupine flag	Dinosaur flag/bear umbrella
Frog bucket	Baby hat	Frog hat/cat bottle
Mouse bed	Cyclist fan	Mouse fan/caterpillar wallet
Pig teapot	Grandmother book	Pig book/cowboy saw
Gorilla sign	Fox tent	Gorilla tent/dinosaur flag
Elephant parachute	Penguin boots	Elephant boots/eagle shorts
Ant shoes	Goat chair	Ant chair/dragon guitar
Eagle handbag	Monkey shorts	Eagle shorts/dog scarf
Cat purse	Pelican bottle	Cat bottle/elephant boots
Snake table	Octopus cap	Snake cap/fish balloon
Seal bracelet	Rat cookie	Rat bracelet/flamingo bell
Fish car	Crane balloon	Fish balloon/frog hat
Hippo bath	Turtle suitcase	Hippo suitcase/giraffe tie
Flamingo ice-cream	Grasshopper bell	Flamingo bell/gorilla tent
Lion pipe	Butterfly can	Lion can/hippo suitcase
Dragon dress	King guitar	Dragon guitar/kitten mittens
Wizard iron	Elf sweet	Wizard sweet/leopard socks
Raccoon backpack	Parrot horn	Raccoon horn/lion can
Puppy cot	Cow jug	Puppy jug/lizard gloves
Cowboy bowl	Fairy saw	Cowboy saw/mouse fan
Turkey apron	Angel vase	Turkey vase/pig book
Kitten bow	Bull mittens	Kitten mittens/pilot jacket
Caterpillar cushion	Bee wallet	Caterpillar wallet/puppy jug
Giraffe earring	Ape tie	Giraffe tie/raccoon horn
Shark bandage	Mosquito belt	Shark belt/rat bracelet
Leopard poncho	Donkey socks	Leopard socks/shark belt
Tiger pyjamas	Ostrich trousers	Tiger trousers/snake cap
Alligator coat	Ladybird trainers	Alligator trainers/tiger trousers
Lizard brush	Mole gloves	Lizard gloves/turkey vase
Pilot jeans	Monk jacket	Pilot jacket/wizard sweet

Appendix B. Items for Experiment 3

Primes		Target
POSSESSOR	DESCRIPTOR	
Dog T-shirt	Rabbit T-shirt	Dinosaur flag
Frog boat	Ladybird boat	Elephant boots
Eagle handbag	Mermaid handbag	Hippo suitcase
Pig teapot	Penguin teapot	Dog scarf
Tiger pyjamas	Cow pyjamas	Bear umbrella
Gorilla basket	Panda basket	Frog hat
Dragon dress	Elf dress	Mouse fan
Hippo plane	Kangaroo plane	Pig book
Ant flippers	Butterfly flippers	Gorilla tent
Kitten bow	Baby bow	Ant chair
Cat purse	Sheep purse	Eagle shorts
Rat cupcake	Rooster cupcake	Cat bottle
Fish car	Grasshopper car	Snake cap
Alligator cardigan	Monkey cardigan	Rat bracelet
Flamingo plate	Owl plate	Fish balloon
Bear mug	Parrot mug	Flamingo bell
Mouse trophy	Cyclist trophy	Lion can
Pilot jeans	Fairy jeans	Dragon guitar
Dinosaur sofa	Grandmother sofa	Wizard sweet
Giraffe pendant	Monk pendant	Raccoon horn
Snake desk	Stork desk	Puppy jug
Shark bandage	Boy bandage	Cowboy saw
Leopard poncho	Ostrich poncho	Turkey vase
Puppy cot	Donkey cot	Kitten mittens
Raccoon backpack	Panther backpack	Caterpillar wallet
Caterpillar cushion	Rhino cushion	Giraffe tie
Lizard brush	Mosquito brush	Shark belt
Turkey apron	Pirate apron	Leopard socks
Lion pipe	Fox pipe	Tiger trousers
Cowboy bowl	Hedgehog bowl	Alligator trainers
Elephant parachute	Mole parachute	Lizard gloves
Wizard teacup	Chicken teacup	Pilot jacket

References

Berman, R. A., & Clark, E. V. (1989). Learning to use compounds for contrast: Data from Hebrew. *First Language, 9*, 247–270.

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology, 18*, 355–387.

Bock, J. K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review, 99*, 150–171.

Branigan, H. P., Pickering, M. J., & Cleland, A. A. (2000). Syntactic coordination in dialogue. *Cognition, 75*, B13–B25.

Branigan, H. P., Pickering, M. J., & McLean, J. F. (2005). Priming prepositional-phrase attachment during comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 468–481.

Carey, P., Mehler, J., & Bever, T. G. (1970). Judging the veracity of an ambiguous sentence. *Journal of Verbal Learning and Verbal Behavior, 9*, 243–254.

Clark, E. V., & Berman, R. A. (1987). Types of linguistic knowledge: Interpreting and producing compound nouns. *Journal of Child Language, 14*, 547–567.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior, 12*, 355–359.

Cleland, A. A., & Pickering, M. J. (2003). The use of lexical and syntactic information in language production: Evidence from the priming of noun-phrase structure. *Journal of Memory and Language, 49*, 214–230.

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*, 299–349.

Estes, Z., & Jones, L. L. (2006). Priming via relational similarity: A copper horse is faster when seen through a glass eye. *Journal of Memory and Language, 55*, 89–101.

Gagné, C. L. (2000). Relation-based combinations versus property-based combinations: A test of the CARIN theory and the dual-process theory of conceptual combination. *Journal of Memory and Language, 42*, 365–389.

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., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 71–87.

Gagné, C. L., & Shoben, E. J. (2002). Priming relations in ambiguous Noun–noun combinations. *Memory & Cognition, 30*, 637–646.

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.

Garrod, S. C., & Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition, 27*, 181–218.

Gerrig, R., & Bortfeld, H. (1999). Sense creation in and out of discourse contexts. *Journal of Memory and Language, 41*, 457–468.

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

Glucksberg, S., & Estes, Z. (2000). Feature accessibility in conceptual combination: Effects of context-induced relevance. *Psychonomic Bulletin & Review, 7*, 510–515.

Levelt, W. J. M., Roelofs, A. P. A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences, 22*, 1–37.

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

Maguire, P., Devereux, B., Costello, F. J., & Cater, A. (2007). A re-analysis of the CARIN theory of conceptual combination. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 811–821.

Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology, 57*, 203–220.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20*, 158–190.

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

Murphy, G. L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language, 29*, 259–288.

Murphy, G. L., & Wisniewski, E. J. (2006). Familiarity and plausibility in conceptual combination: Reply to Gagné and Spalding (2006). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 1438–1442.

Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic persistence in written language production. *Journal of Memory and Language, 39*, 633–651.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 33–58). Hillsdale, NJ: Erlbaum.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selection modification model. *Cognitive Science, 12*, 485–527.

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

Storms, G., & Wisniewski, E. J. (2005). Does the order of head noun and modifier explain response times in conceptual combination? *Memory & Cognition, 33*, 852–861.

Watson, M.E., Pickering, M.J., & Branigan, H.P. (2004). An empirical investigation into spatial reference frame taxonomy using dialogue. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, Chicago, July 2004.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language, 35*, 434–453.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review, 4*, 167–183.

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

Wisniewski, E. J., & Murphy, G. L. (2005). Frequency of relation type as a determinant of conceptual combination. A reanalysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 169–174.