Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Memory and Language 48 (2003) 304–319
Journal of Memory and Language
www.elsevier.com/locate/jml

Attributive and relational processes in nominal combination

Zachary Estes

University of Georgia, Department of Psychology, Athens, GA 30602, USA

Received 19 October 2001; revision received 15 March 2002

Abstract

The dual process theory of nominal (noun–noun) combination posits a relational process, in which a relation between concepts is inferred, as well as an attributive process, in which a property of one concept is attributed to another. According to dual process theory, these attributive and relational processes occur in parallel. A relational theory claims instead that attributive and relational comprehension result from the same process, and assumes that relational comprehension will occur serially prior to attributive comprehension. Experiment 1 used a priming paradigm to test whether the relational and attributive processes occur serially or in parallel. Target combinations were more likely to be comprehended, and were comprehended more quickly, when preceded by a prime combination that used the same attribution or relation than when preceded by a prime combination that did not engage the same attributive or relational process. Critically, the patterns of facilitation and interference were virtually identical across the attributive and relational target-types, suggesting that the processes occur in parallel. Experiment 2 showed that particular attributes and relations were primed, rather than the attributive or the relational process more generally. Results of both experiments supported the dual process theory. The emergence of a general model of nominal combination is discussed.

Nominal (noun–noun) combinations are a concise and ubiquitous linguistic construction. In the English language, they typically assume a syntactic structure in which the first noun (i.e., the modifier) serves to modify the second noun (i.e., the head). By combining noun concepts in this way, one specifies a subclass of referents (e.g., CIRCUS LIONS) from a larger class of possible referents (e.g., LIONS), and in so doing, makes relevant those properties that differentiate members of the combined category from members of the head noun category (Estes & Glucksberg, 1999; Glucksberg & Estes, 2000). For instance, LIONS are ferocious, but CIRCUS LIONS are tame (see Franks, 1995; Gagné & Murphy, 1996; Springer & Murphy, 1992). Nominal combinations also serve as abbreviations for more lengthy phrasal descriptions. For example, one may refer to AN AREA FOR PARKING ONE'S CAR WHILE ATTENDING A FOOTBALL GAME much more concisely with the nominal combination FOOTBALL PARKING (Wisniewski & Love, 1998).

While this capacity for concise yet specific language is of course a virtue, it nonetheless introduces into the comprehension process a significant amount of inferential work. That is, the comprehender of a nominal combination must infer the intended relation between the constituent concepts. For instance, because the relation between FLOOR and TELEVISION is not explicitly stated in the nominal combination FLOOR TELEVISION, the comprehender must infer that relation (i.e., a FLOOR TELEVISION is a television *located on* the floor). Essentially, the comprehender must infer the predicate argument structure underlying the nominal combination.

The present investigation examines the processes by which people comprehend such nominal combinations. Specifically, dual process theory (e.g., Wisniewski, 1997) posits that, in contrast to the relational process described above, some combinations are interpreted by an attributive process in which a property of the modifier is attributed to the head concept (e.g., a ZEBRA CLAM is a

E-mail address: estes@uga.edu.

0749-596X/02/$ - see front matter
PII: S0749-596X(02)00507-7

striped clam). According to dual process theory, property attribution and relational inference are distinct processes, and for any given combination, these attributive and relational processes are attempted in parallel (Wisniewski, 1997; Wisniewski & Love, 1998). The given combination will be interpreted either attributively or relationally, depending on which process produces a plausible interpretation first.

Gagné (2000), however, has argued that, rather constituting a distinct process in itself, attribution is really just a specific type of relation, namely, a *resemblance* relation (e.g., a ZEBRA CLAM is a clam that *resembles* a zebra). Such a relational treatment of attribution is parsimonious in that instead of distinct attributive and relational processes operating concurrently, a single relational process could account for both types of nominal combinations. Gagné (2000) proposes a serial model in which the various relations (e.g., *location, resemblance*, etc.) are attempted in order of frequency; more frequent relations are attempted first, less frequent relations are attempted only later. Critically, because *resemblance* is an infrequent relation (Downing, 1977; Gagné, 2000; Warren, 1978), attributive combinations should take longer to comprehend.

Thus, dual process theory claims that attributive and relational comprehension occur in parallel, whereas relational theory claims that relational comprehension (e.g., *location*, etc.) should serially precede attributive comprehension (i.e., *resemblance*). I will present two experiments that address this question of whether attributive and relational comprehension occur serially or in parallel, and will consequently discriminate between the two models of nominal combination.

Dual process theory

The central tenet of dual process theory is that property attribution and relational inference are carried out by two distinct processes in nominal combination. According to Wisniewski (1997, 2000), attributive combinations are comprehended via comparison and mapping, while relational combinations are comprehended by thematic role assignment. For instance, in comprehending the attributive combination CACTUS CARPET, a comparison of CACTUS and CARPET reveals that cacti are prickly while carpets typically are soft. Such a comparison is necessary for attribution because it indicates which property to map, and to where it should be mapped: The most salient properties of a concept are the most likely candidates for attribution, and those properties should be mapped to the dimension with which they are aligned. Thus, as a result of this comparison, the *prickly* property of CACTUS is attributed to the textural dimension of CARPET, yielding the interpretation that a CACTUS CARPET is a prickly carpet. In contrast to this attributive process, understanding the relational combination FLOOR TELEVISION requires a very different process. Here the constituent concepts are assigned different thematic roles in a conceptual relationship: FLOOR fulfills the role of location, while TELEVISION is the located, producing the interpretation that a FLOOR TELEVISION is a television *located on* the floor. Thus, attribution entails comparison and mapping, while relational comprehension entails thematic role assignment. These two processes are claimed to operate concurrently (Wisniewski, 1997; Wisniewski & Love, 1998).

The dual process theory has intuitive appeal. At face value, attributive and relational comprehension do appear to be fundamentally different. In relational comprehension, the two constituent concepts are retained in their entirety. For instance, the representation of FLOOR TELEVISION includes both floor and television (i.e., a *television* located on the *floor*). In attributive comprehension, on the other hand, the modifier is represented in the combination by only a subset of its many properties (Wisniewski, 1996). For example, in the representation of ZEBRA CLAM, only the zebra's stripes are reconstructed in the clam. In terms of representation, then, the distinction between attributive and relational comprehension is incontrovertible.

A key prediction of Wisniewski's dual process model is that attribution should be more frequent in similar concepts (e.g., ZEBRA HORSE) than in dissimilar concepts (e.g., ZEBRA STABLE). This prediction arises from the fact that relational combination requires the two constituent concepts to perform different thematic roles in the relational argument. And because similar concepts tend to assume the same thematic roles, similar concepts tend not to be interpreted relationally. For instance, zebras and horses generally have the same sorts of properties, and therefore, compete for the same thematic roles. Thus, it is difficult to interpret the combination ZEBRA HORSE relationally, and one tends instead to interpret it attributively as "a striped horse." The less similar concepts ZEBRA and STABLE, on the other hand, have different sorts of properties, and perform different, complementary sorts of roles. As a consequence, ZEBRA and STABLE readily assume different thematic roles in a relational argument: a ZEBRA STABLE is "a stable that shelters zebras." Consistent with this prediction, several studies have found a positive correlation between the similarity of the constituent concepts and the likelihood of attributive interpretation (Markman & Wisniewski, 1997; Wilkenfeld & Ward, 2001; Wisniewski, 1996; Wisniewski & Love, 1998; but see General discussion).

Although the correlation between similarity and attribution supports Wisniewski's model in general, it is not particularly convincing evidence that the attributive and relational processes operate concurrently. For, as

Wisniewski and Love (1998) point out, the attributive combinations used in those previous studies might have lacked any plausible relation between their constituent concepts, and therefore, the attributive process might have been evoked only *after* the relational process failed to yield an interpretation. This explanation of the data is perfectly consistent with a serial model of nominal combination. In a test of this serial explanation, Wisniewski and Love (1998, Experiment 1) demonstrated that similar combinations tend to be interpreted attributively even when a plausible relation does exist between the concepts. For example, participants in their study agreed that a BOOK MAGAZINE could plausibly be a magazine *about* books, which is a relational interpretation. Another group of participants, though, tended to interpret such combinations attributively (e.g., "a thick magazine"). This result argues against the serial model—if attribution were attempted only after the relational process failed, then the presence of a plausible relation should have obviated the attributive process entirely. But this clearly was not the case; the similar combinations were interpreted attributively 66% of the time, despite the existence of a plausible relation.

In a second experiment, Wisniewski and Love preceded ambiguous combinations with attributive combinations or relational combinations. To illustrate, the ambiguous combination BOOK MAGAZINE was preceded either by a series of 10 attributive combinations (e.g., UMBRELLA TREE, RAZOR INSULT, etc.), by a series of 10 relational combinations (e.g., BIRTHDAY DINNER, KIDNEY SURGEON, etc.), or by no prime combinations at all (i.e., the "neutral" condition). The ambiguous target combinations were interpreted attributively 45% of the time in the neutral condition. But when preceded by attributive primes, those same target combinations were interpreted attributively 68% of the time. They were interpreted attributively only 37% of the time when preceded by relational primes. Wisniewski and Love argued that if the relational process were always attempted prior to the attributive process, as the serial model claims, then interpretation of the ambiguous combinations should have been unaffected by the preceding combinations. Thus, the priming of attributive interpretation suggests that it is not always a strategy of last resort.

Finally, Wisniewski and Love sought to demonstrate that attributive interpretation is not only different from relational interpretation, but also that attributive interpretation is ubiquitous. To this end, Wisniewski and Love (Experiment 3) determined the relative proportion of attributive combinations in a large sample of nominal combinations. They found that 29% of the combinations in their sample were attributive. Wisniewski and Love thus suggested that the attributive process is not rare, and therefore, models of nominal combination must account for it.

In sum, then, nominal combination appears to involve dual attributive and relational processes. However, the relational theory (described below) suggests that this duality may be more apparent than real. As Gagné (2000) argues, the question of serial primacy is better addressed by on-line comprehension measures than by the off-line interpretation task used by Wisniewski and Love (1998). Such off-line tasks leave open the possibility that relational comprehension is in fact carried out serially prior to attributive comprehension, though the attributive interpretation may sometimes be selected as the more likely interpretation.

Relational theory

According to the relational theory, nominal combinations are understood by inferring some relation that is purported to exist between the constituent concepts. For example, the combination WATER BALLOON is understood by inferring the *containment* relation, as in "a balloon that *contains* water." There are many such conceptual relations. In addition to *containment*, other relations include *causation* (e.g., STRESS HEADACHE), *location* (e.g., FLOOR TELEVISION), *composition* (e.g., BRICK HOUSE), *part* (e.g., HORSE HOOVES), *about* (e.g., RODEO MAGAZINE), and so on. Critically, the relational theory rejects the distinction between attributive and relational processes. Relational theorists do recognize the occurrence of attributive combinations (e.g., CACTUS CARPET). However, by this account, attributive interpretation simply entails the use of a *resemblance* relation, with the particular properties of resemblance specified later in a subsequent stage of processing. That is, in order to comprehend CACTUS CARPET as "a prickly carpet," one must first infer the *resemblance* relation and only later determine that *prickliness* is the basis of the resemblance. Thus, attribution is just another conceptual relation. Gagné exemplifies this position: "it is not necessary to posit that relation interpretations and property interpretations arise from two distinct sets of processes. Both relation interpretations and property interpretations can be accounted for within a relation-based framework if one assumes that the selection of a relation is followed by an elaboration process in which the properties/features of the newly formed combinations are derived" (2000, p. 384). The claim of the relational theory, then, is that attributive and relational comprehension result from the same relational process.

Citing several linguistic studies (i.e., Downing, 1977; Gleitman & Gleitman, 1970; Levi, 1978; Warren, 1978) as their foundation, Gagné and Shoben (1997) developed a psychological model of relational comprehension (see also Coolen, van Jaarsveld, & Schreuder, 1991). They argued that people have (presumably implicit) statistical knowledge of which semantic relations are likely to occur with which concepts, and that this statistical knowledge affects comprehension. For example, the concept CHOCOLATE, when used as a modifier, tends to instantiate the *composition* relation with most head nouns (e.g., a CHOCOLATE BALL is a ball *made of* chocolate). The *sell* relation, on the contrary, is less frequent for the modifier CHOCOLATE (e.g., a CHOCOLATE SHOP is a shop that *sells* chocolate). Importantly, each modifier may have several relations that differ in frequency of use, and hence different relations are dominant for different modifiers. Gagné and Shoben proposed the CARIN model (*C*ompetition *A*mong *R*elations *i*n *N*ominals), which predicts that the different relations of a modifier compete serially during comprehension, with the most frequent relations being attempted first and the least frequent relations being attempted last. As a consequence, nominal combinations using a relation that is frequent for the modifier (e.g., CHOCOLATE BALL) will be comprehended faster than combinations using a less frequent relation for that particular modifier (e.g., CHOCOLATE SHOP). But because modifiers precede head nouns in the English language, the frequency with which a relation is associated with the head noun will not significantly affect comprehension speed. That is, the fact that the *composition* relation is relatively infrequent for the head concept BALL is irrelevant, because it is the frequency of that relation for the modifier CHOCOLATE that determines comprehension speed.

Data obtained by Gagné and Shoben (1997) supported these predictions. In their study, participants read nominal combinations and judged whether they were comprehensible. The combinations were of three different types, varied as to whether the relation between concepts was of high (H) or low (L) frequency for the given nouns. In one condition, the relation between the concepts was highly frequent for both the modifier and the head noun (i.e., the HH condition). For example, the *composition* relation occurs frequently for both the modifier PLASTIC and the head noun TOY, creating the HH combination PLASTIC TOY. In the HL condition, the relation was highly frequent for the modifier, but infrequent for the head. The *composition* relation rarely occurs with the head concept EQUIPMENT, and thus PLASTIC EQUIPMENT constitutes an HL combination. And finally, the *about* relation is infrequent for the modifier PLASTIC, but is frequent for the head noun CRISIS, so the combination PLASTIC CRISIS was in the LH condition. As predicted by the CARIN model, the HH ($M = 1068$ ms) and HL ($M = 1072$ ms) combinations were comprehended faster than the LH combinations ($M = 1152$ ms). That is, the frequency of occurrence of the target relation with the given modifier reliably predicted comprehension speed. Moreover, there was no reliable difference in response times for the HH and HL combinations, suggesting that the frequency of the target relation for the given head noun does not affect comprehension speed.

Gagné (2001) next used a priming procedure to further test the CARIN model. She reasoned that prior exposure to a particular relation should make that relation more "available" for subsequent use, thus facilitating comprehension. In one experiment, target combinations (e.g., GAS LAMP) were preceded by prime combinations that used either the same relation (e.g., GAS CAR), a different relation (e.g., GAS HOSE), or an unrelated combination (e.g., RICE PAPER). The results indicated that, although the same relation and different relation conditions both facilitated comprehension in comparison to the unrelated condition, the same relation condition facilitated comprehension more than the different relation condition. Thus, when the modifier concept was repeated from prime to target, the particular relation used for comprehension was primed. In another experiment, target combinations again used either the same or a different relation, but here the head concept was repeated from prime to target. For example, the target GAS LAMP was preceded by either BATTERY LAMP (same relation), CABIN LAMP (different relation), or RICE PAPER (unrelated). As in the previous experiment, both the same relation and different relation conditions facilitated comprehension, in comparison to the unrelated condition. But unlike the previous experiment, there was no reliable difference in comprehension times in the same relation and different relation conditions. Thus, modifier repetition elicited relation priming, but head repetition did not. In a final experiment, Gagné tested whether relations could be primed when neither concept is repeated from prime to target. Target combinations (e.g., GAS LAMP) were preceded by prime combinations that lacked lexical repetition, but used either the same relation (e.g., BATTERY CAR) or a different relation (e.g., CHOCOLATE SHOP). Comprehension times showed no reliable difference between these conditions, suggesting that relations cannot be primed without lexical repetition from prime to target (but see Gerrig & Murphy, 1992, described below).

These findings support the three main claims of the CARIN model. First, prior use of a conceptual relation affects the subsequent use of that relation. Second, a relation's association with the modifier concept is more influential than its association with the head concept. And third, the lack of relation priming in the absence of modifier repetition suggests that relations are represented with their particular modifier concepts.

Gagné (2000) also conducted a corpus study investigating the relative frequencies of attributive and relational combinations in ordinary discourse. Sampling from the Brown corpus, which is generally accepted as representative of everyday language, Gagné reported that attributive combinations accounted for 1% of all nominal combinations in her sample. She also noted the data of Warren (1978) and Downing (1977), both of whom found attributive combination to be infrequent in their corpora (i.e., less than 2%). Although Downing (1977), Gagné (2000) and Warren (1978) all found attributive combination to be rare in their samples, recall that Wisniewski and Love (1998) found attributive combination to be much more common (29%) in their sample. As Gagné (2000, p. 371) notes, this difference in results might be explained by differences in sampling. The sample used by Wisniewski and Love included a large proportion of natural (i.e., animal and plant) concept combinations sampled from a wildlife book. Critically, natural concepts are more likely to be used attributively than are artifact concepts (Bock & Clifton, 2000; Downing, 1977; Wisniewski & Love, 1998). For example, Wisniewski and Love found that 43% of the natural kind combinations in their sample were attributive, while only 14% of their artifact stimuli were attributive. Wisniewski and Love's sample, then, might have yielded a more robust percentage of attributive combinations as a consequence of its high proportion of natural concepts. Thus, assuming that the Brown corpus is more typical of everyday language than Wisniewski and Love's sample, Gagné (2000) concluded that attributive combination is in fact a rare occurrence in everyday language.

The relative infrequency of attributive comprehension leads the relational theory to a critical prediction: If attribution entails the use of a *resemblance* relation (e.g., Gagné, 2000), and if frequent relations are processed serially prior to infrequent relations (Gagné & Shoben, 1997), then it follows from the infrequency of attribution that attributive comprehension should occur serially subsequent to relational comprehension (Gagné, 2000). Gagné (2000) tested this prediction by presenting attributive and relational combinations to participants in a variety of tasks. Consistent with the CARIN model, Gagné found that attributive combinations took longer to comprehend, were less likely to be comprehended, were less likely to be produced, and were judged less acceptable than relational combinations. This demonstration of the primacy of relational combination over attributive combination therefore suggests a serial comprehension process (Gagné, 2000).

Overview of the present research

It is crucial to note that the conclusion of serial primacy does not necessarily follow from the evidence of faster and easier comprehension (nor does Gagné claim necessity). Indeed, one plausible alternative explanation is that attributive and relational comprehension occur in parallel, with relational comprehension just tending to be faster and easier. This possibility was tested via a priming paradigm in the experiments reported below.

Ample evidence demonstrates that the apprehension of properties and relations of nominal combinations is susceptible to priming effects via contextual manipulations. Several studies have shown that particular properties of combined concepts can be made more or less accessible by context (e.g., Gagné & Murphy, 1996; Glucksberg & Estes, 2000). For instance, the property "white" of the combination PEELED APPLES was verified faster following a context that included DICED APPLES, which are also white, than following a context that included WHOLE APPLES, which are not white (Estes & Glucksberg, 1998; see also Estes & Glucksberg, 1999). Particular relations can also be made more or less accessible by a preceding context (Gerrig & Bortfeld, 1999; Gerrig & Murphy, 1992). For example, the study by Gagné (2001) described above showed that comprehension of a target combination was facilitated by a prime combination that used the same relation (and the same modifier) as that target. And the study by Wisniewski and Love (1998) showed that a preceding context could even influence whether an ambiguous combination would be interpreted attributively or relationally.

Experiment 1, reported below, tested whether the attributive and relational comprehension processes occur serially or in parallel. Experiment 2 further tested whether the *particular* attribute or relation is activated by the prime combination (Gagné, 2001), or whether the *process* of attributive or relational comprehension is primed instead (Wisniewski & Love, 1998, 2000). Thus, Experiments 1 and 2 investigated the nature of attributive and relational comprehension in nominal combination.

Experiment 1

The present experiment used a priming paradigm to test whether relational comprehension is serially prior to attributive comprehension, or whether they occur in parallel. Attributive target combinations (e.g., SHARK POLITICIAN) were preceded by prime combinations that used either the same attribution (e.g., PIRANHA LAWYER) or a relation (e.g., SWIMMING FLIPPERS). And likewise, relational target combinations (e.g., RUGBY SHOES) were preceded by prime combinations that used either the same relation (e.g., SWIMMING FLIPPERS) or an attribution (e.g., PIRANHA LAWYER). Both target types were also preceded by a baseline prime (i.e., SOMETHING SOMETHING).

If relational comprehension is serially prior to attributive comprehension, then Prime-type and Target-type should interact. To see why this prediction follows from relational theory, assume that attribution is a *re-*

semblance relation. Because it occurs infrequently, it is attempted only after other more frequent relations are attempted, in a serial manner. Call a more frequent relation here "R_1," and the less frequent *resemblance* relation "R_2." If R_1 is ordinarily attempted before R_2 (as the relational theory claims), then presenting a R_2 prime before a R_1 target should induce additional processing—R_2 would not have been activated without the prior presentation of that R_2 prime. The R_2 prime would thus induce additional processing, thereby slowing comprehension of the R_1 target. But presenting a R_1 prime before a R_2 target would not induce any additional processing—R_1 would have been attempted before R_2 anyway. Thus, if relational comprehension serially precedes attributive comprehension, then an interaction should obtain. If, on the contrary, attributive and relational comprehension processes operate concurrently, then preceding a R_1 target with a R_2 prime, and preceding a R_2 target with a R_1 prime, might be expected to produce similar effects. That is, because both processes are claimed to operate simultaneously, the dual process theory predicts no interaction.

Method

Participants. Twenty-eight Princeton University undergraduates participated for course credit. All participants spoke American English as their first language.

Materials. Twenty-four attributive and 24 relational combinations were chosen under the intuition that each had 1 unambiguous meaning. A pre-test confirmed this intuition. In the pre-test, each combination was presented with its expected interpretation (e.g., RUGBY SHOES most likely means "shoes worn while playing rugby"). The pre-test additionally included 10 filler combinations with unlikely interpretations (e.g., SONG BOOK most likely means "a book that sings"). The combinations and their interpretations were presented in random order. Ten participants rated how much they agreed with the given interpretation of each combination. The scale ranged from one (completely disagree) to seven (completely agree). The mean agreement rating for the filler items was 1.39 ($SD = .40$). More importantly, mean agreement ratings did not reliably differ between attributive ($M = 5.69$, $SD = .63$) and relational ($M = 5.82$, $SD = .63$) combinations, $p = .48$. Thus, the pre-test verified the interpretations of the combinations. Of course, in the experiment proper, participants could potentially produce different interpretations than those agreed upon by the participants in the pre-test. But to the extent that the experiment proper produces reliable effects, one may infer that those combinations were interpreted as expected.

Attributive and relational combinations were also matched for length: The mean number of syllables in attributive ($M = 4.33$, $SD = 1.40$) and relational ($M = 4.38$, $SD = 1.21$) combinations was virtually identical. Each of these 48 experimental targets was paired with one prime combination that used the same property or relation (i.e., "Same" condition), one prime combination that used a different property or relation (i.e., "Different" condition), and one uninformative prime (i.e., "Baseline" condition). For instance, the relational target RUGBY SHOES was preceded either by SWIMMING FLIPPERS, which uses the same relation, by PIRANHA LAWYER, which uses an attribution, or by SOMETHING SOMETHING, which is uninformative. Forty-eight nonsense filler targets (e.g., CHOIR PEEL) were also included. The nonsense targets were always preceded by a sensical prime combination (e.g., PEACH SEED). Thus, the prime combinations were always sensical, while the target combinations were sensical only half the time. Since each target appeared with three Prime-types (i.e., baseline, same, different), three lists were constructed, with each target appearing once per list. Each list contained eight targets in each of the six experimental conditions (i.e., three Prime-types $\times$ two Target-types), in addition to the 48 nonsense filler targets. Thus, Prime-type and Target-type were both manipulated within-participants.

The experimental stimuli are presented in Appendix A. As apparent there, some of the prime combinations (e.g., PLASTIC BOWL) may be more familiar than others (e.g., NOSE SOUND). However, regardless of the familiarity of a combination, its comprehension entails apprehending the relation (or attribution) between its constituent concepts. For instance, assuming that PLASTIC BOWL is in fact a relatively familiar combination, comprehension nonetheless requires knowing that it is a bowl *composed of* plastic. Thus, because it did not seem relevant to the present purposes, the familiarity of the prime combinations was not strictly controlled.

Procedure. Each trial began with the prompt "Prepare for the next trial," which appeared in the center of the screen for 1 s. A 500 ms interval occurred prior to presentation of the prime combination, which was presented in black lower-case letters in the center of the screen. The prime remained on the screen until the participant pressed the space bar. Participants were instructed to press the space bar only after they understood the meaning of the prime combination. There was a 500 ms inter-stimulus interval, followed by a 500 ms fixation cross in the center of the screen. Finally, the target combination appeared in red capital letters. The target remained on the screen until the participant responded. Participants were instructed to press the 'j' key if the combination "made sense" or the 'f' key if the combination "did not make sense." The order of presentation of the trials was random for each participant. The 96 trials of the experiment proper were preceded by 12 practice trials. There was one practice trial for each of the three Prime-types (baseline, same, different) for each of the two Target-types (attributive, relational), yielding six sensical practice trials. The other six practice trials had nonsensical target combinations.

Results and discussion

Mean response times and probabilities of sense judgments are presented in Table 1. Notice first that attributive target combinations took longer to comprehend (attributive baseline $M = 1638$, relational baseline $M = 1520$), and were less likely to be comprehended (attributive baseline $M = .63$, relational baseline $M = .81$), than relational target combinations. This finding is consistent with the claim that relational comprehension occurs serially prior to attributive comprehension. However, a more direct and critical test of serial processing is stated in the following prediction: If the relational process were attempted first, then an attributive prime should hinder comprehension of a relational target more than a relational prime should hinder comprehension of an attributive target. Because an attributive prime would induce additional processing (i.e., the attributive process) that would not otherwise have occurred prior to relational comprehension, the attributive prime should hinder comprehension of the relational target. But preceding an attributive target with a relational prime should have caused less interference (if any at all), because the relational processing induced by the prime would have occurred prior to the attributive process anyway. Therefore, the relational theory predicts an interaction of Prime-type and Target-type. Critically, though, this interaction did not obtain, and hence the present result fails to support the relational theory. Instead, the results of Experiment 1 supported the dual process theory of nominal combination. If the attributive and relational processes operate in parallel, then preceding an attributive target with a relational prime should have approximately the same effect as preceding a relational target with an attributive prime. This pattern of results is apparent in Table 1.

Because 2 of the 28 participants judged all targets in one condition to be "nonsensical," the data of those two participants were excluded from statistical analyses. All responses greater than 2.5 standard deviations from a participant's mean response time were excluded from analyses (2% of the data). Response time analyses also excluded trials in which the target combination was judged "nonsensical." The remaining data were submitted to two separate analyses of variance (ANOVAs), one with subjects as a random factor (F_s) and one with items as a random factor (F_i). A three (Prime-type: baseline, same, different) × two (Target-type: attributive, relational) ANOVA revealed a reliable main effect of Prime-type on both response time [$F_s(2, 50) = 4.53$, $p = .02$ and $F_i(2, 92) = 7.35$, $p = .001$] and likelihood of comprehension [$F_s(2, 50) = 20.07$, $p < .001$ and $F_i(2, 92) = 18.48$, $p < .001$]. That is, comprehension of the target was facilitated by the prior presentation of a prime that used the same attribute or relation as that target, regardless of whether the target combination was attributive or relational. This main effect will be the focus of Experiment 2 below. There was also a significant main effect of Target-type on comprehension time [$F_s(1, 25) = 19.02$, $p < .001$ and $F_i(1, 46) = 5.68$, $p = .02$] and comprehension probability [$F_s(1, 25) = 18.42$, $p < .001$ and $F_i(1, 46) = 5.75$, $p = .02$], indicating that comprehension of relational targets was faster and more likely than comprehension of attributive targets. This primacy of relational over attributive comprehension replicates the finding of Gagné (2000), and is consistent with the proposal that relational comprehension might be serially prior to attributive comprehension (Gagné, 2000). However, as described above, more direct evidence of serial primacy would be an interaction of Prime-type and Target-type. And as illustrated in Table 1, this critical interaction did not approach significance in either the subject or the item analyses of either response times or sensicality judgments, all $F < 1$. Thus, Experiment 1 failed to support the claim that attributive and relational comprehension occur serially in nominal combination.

Although the mean proportion of "sense" judgments of the target combinations ($M = .73$) may at first seem low, a review of other comparable studies indicates that this proportion may actually be higher than most. Across all experiments of which I am aware that re-

Table 1
Mean response times (in ms) and proportions of "sense" judgments, Experiment 1

Target-type	Prime-type		
	Baseline	Same	Different
Attributive			
Response time	1638 (109)	1524 (110)	1637 (118)
Sensicality	.63 (.05)	.80 (.03)	.56 (.04)
Relational			
Response time	1520 (121)	1328 (77)	1546 (119)
Sensicality	.81 (.04)	.88 (.02)	.71 (.05)

Note. Standard errors are in parentheses.

sponse times to relational noun–noun combinations were measured in a sense/nonsense judgment task (i.e., Gagné, 2000, Experiments 1 and 2; Gagné, 2001, Experiments 1, 2, 3, and 6; Gagné & Shoben, 1997, Experiments 1 and 3; Gerrig & Bortfeld, 1999, Experiment 2b; Murphy, 1990, Experiment 1), the mean proportion of "sense" judgments was approximately .82, and ranged from .68 (Gagné, 2000, Experiment 2) to .90 (Gagné, 2001, Experiment 1). The mean proportion of "sense" judgments for relational combinations in the present study (baseline $M = .81$) was in line with those previous studies. I am aware of only one study in which attributive combinations served as targets in a timed sense/nonsense judgment task: Gagné (2000) found that 47% of attributive combinations were judged sensical in her Experiment 1 and only 22% in her Experiment 2. Comparatively speaking, then, the proportion of sense judgments for attributive combinations in the present experiment (baseline $M = .63$) may be viewed as high rather than low. It might also be noted that participants in the pre-test of the present study overwhelmingly agreed with the interpretations of the target combinations, which of course presupposes that they are sensible. Thus, it is likely that the failure of participants to comprehend all of the target combinations in the experiment proper was simply due to the time pressure inherent in the task.

Experiment 2

Experiment 1 failed to support the claim that relational comprehension occurs serially prior to attributive comprehension. Instead, the two comprehension processes appeared to operate concurrently. Another finding of Experiment 1 was that target combinations were more likely to be comprehended, and were also comprehended more quickly, when preceded by prime combinations that used the same attribution or relation. This was true even though there was no lexical overlap between prime and target combinations, a finding that is inconsistent with the CARIN model of nominal combination. According to Gagné (2000, 2001; Gagné &Shoben, 1997), conceptual relations are only activated in the context of their particular modifier concepts, and hence relation priming should occur only if the prime and target combinations have the same modifier. For instance, Gagné (2001) found that comprehension of the target combination GAS LAMP was facilitated by the prime GAS CAR but not by the primes BATTERY LAMP or BATTERY CAR, even though all four combinations use the same conceptual relation. In Experiment 1 above, on the contrary, comprehension of RUGBY SHOES was facilitated by the prior presentation of SWIMMING FLIPPERS, even though the two combinations have no lexical redundancy. The CARIN model cannot account for this finding of relation priming in the absence of lexical repetition.

At the same time, the finding of relation priming in the absence of lexical repetition is not without precedent: Gerrig and Murphy (1992) obtained a similar result in an off-line task. They had participants read stories in which the relation used for comprehension of a prime combination was also required for comprehension of a target combination. Critically, there was no lexical overlap between the prime and target combinations. For instance, in one story, participants read about a woman who was carving figures out of olives. The story included the prime combination TRUMPET OLIVE (i.e., a trumpet *carved out of* an olive), and concluded with the target combination KITTEN APPLE (i.e., a kitten *carved out of* an apple). There was also a neutral context that lacked the prime combination. After reading the context stories, participants judged how difficult it was to understand the target combination, and then provided their interpretation of that target. Gerrig and Murphy found that, in comparison to the neutral contexts, target combinations were easier to comprehend, and were more likely to be interpreted correctly, when preceded by prime combinations that used the same relation. Thus, Gerrig and Murphy demonstrated relation priming without lexical repetition, similar to that of Experiment 1 above, albeit in an off-line task. These two results appear to contradict the claim of the CARIN model.

However, it could be argued that the facilitation in Experiment 1 was due to process priming (cf. Wisniewski & Love, 1998, 2000) rather than attribute priming (cf. Estes & Glucksberg, 1998) or relation priming (cf. Gagné, 2001). That is, perhaps a relational prime facilitates comprehension of *any* subsequent relation, not just the relation used by that prime combination. Indeed, Wisniewski and Love (2000) reported that prior presentation of a series of attributive combinations (e.g., SKUNK CIGAR) facilitated the subsequent comprehension of metaphors (e.g., THAT JOB IS A JAIL), even though the combinations and metaphors lacked lexical repetition. Similarly, Wisniewski and Love (1998) showed that ambiguous target combinations were more likely to be interpreted attributively when preceded by a series of attributive prime combinations, and were more likely to be interpreted relationally when preceded by a series of relational prime combinations, despite the lack of lexical repetition.

In terms of Experiment 1 above, then, comprehension of SWIMMING FLIPPERS might simply have primed the relational *process* rather than the *particular* relation involved in comprehension of RUGBY SHOES. If this were the case, then Experiment 1 would not necessarily provide evidence against the CARIN model. Process priming may be considered a sort of processing bias induced by the experimental context, analogous to the "mental set" effects illustrated most famously by the

Luchins (1942) water jug problems. And that may be an entirely different phenomenon from the retrieval and use of specific conceptual relations. In other words, process priming in no way implicates the priming of particular relations. Importantly, Gagné's (2000, 2001; Gagné & Shoben, 1997) claim is that *particular* relations are not primed in the absence of lexical repetition.

Therefore, to test whether the facilitation obtained in Experiment 1 was due to process priming, relational target combinations (e.g., RUGBY SHOES) were preceded by prime combinations that used either the same relation (e.g., SWIMMING FLIPPERS) or a different relation (e.g., ROAD CONSTRUCTION), and attributive targets (e.g., SHARK POLITICIAN) were preceded by primes that used either the same attribution (e.g., PIRANHA LAWYER) or a different attribution (e.g., MIRROR POND) in Experiment 2. So whereas the "Different" condition in Experiment 1 used a different comprehension process, the "Different" condition in Experiment 2 used the same attributive or relational process, but the particular attribute or relation differed. An uninformative "Baseline" prime condition (i.e., SOMETHING SOMETHING) was also included. If Experiment 1 revealed only process priming, then there should be no difference between the Same and Different conditions here in Experiment 2, because both primes use the same process (Gagné, 2001). But if Experiment 1 demonstrated the priming of particular attributes and particular relations, then a main effect of Prime-type should obtain: Comprehension of attributive and relational targets should be facilitated by primes that use the same attribute or relation, but should not be facilitated by primes that use a different attribute or relation. Such a result would be inconsistent with the CARIN model.

The dual process theory clearly predicts that attributive and relational comprehension can be primed. But it makes no clear prediction about the Prime-type variable. On the one hand, Wisniewski and Love's (1998, 2000) process priming results suggest that the Same and Different prime-types should both facilitate comprehension, because both use the same process. But on the other hand, if one assumes that process priming requires some critical number of primes (i.e., greater than 1) to occur, then a failure to obtain process priming in the present experiment might just indicate that there was an insufficient number of prime combinations to induce process priming. Thus, the dual process theory makes no clear prediction here, and hence Experiment 2 was intended primarily as a further test of the CARIN model.

Method

Participants. Thirty-three Princeton University undergraduates participated for partial course credit. None had participated in Experiment 1.

Materials. Target combinations were preceded by prime combinations that used either the same attribution or relation (i.e., "Same" condition), a different attribution or relation (i.e., "Different" condition), or an uninformative prime (i.e., "Baseline" condition). Importantly, the primes in the Same and Different conditions always used the same process (i.e., attributive or relational) as their target combinations. The materials are presented in Appendix A.

Procedure. The procedure was identical to that of Experiment 1.

Results

Results are presented in Table 2. First note that the response times and sensicality judgments of Experiment 2 were very similar to those of Experiment 1. In particular, the procedure and materials of the Baseline and Same conditions of Experiment 2 were an exact replication of those conditions in Experiment 1. And as the table shows, mean response times and probabilities of "sense" judgments were similar between experiments. Notice also that primes that used the same attribute or relation as their targets ("Same" $M = 1456$) once again facilitated comprehension of those targets ("Baseline" $M = 1578$), thus replicating the finding of Experiment 1. The novel result of Experiment 2 is that primes that used a different attribute or relation from their targets interfered with comprehension of those targets ("Different"

Table 2
Mean response times (in ms) and proportions of "sense" judgments, Experiment 2

Target-type	Prime-type		
	Baseline	Same	Different
Attributive			
Response time	1668 (93)	1539 (89)	1733 (102)
Sensicality	.68 (.04)	.75 (.03)	.62 (.04)
Relational			
Response time	1487 (77)	1374 (72)	1707 (126)
Sensicality	.79 (.03)	.91 (.02)	.67 (.04)

Note. Standard errors are in parentheses.

$M = 1720$), even though the primes and targets used the same attributive or relational process. A three (Prime-type: baseline, same, different) × two (Target- type: attributive, relational) ANOVA indicated that this main effect of Prime-type was significant for both response times [$F_s(2,64) = 11.70$, $p < .001$ and $F_i(2,90) = 8.28$, $p < .001$] and sensicality judgments [$F_s(2,64) = 20.36$, $p < .001$ and $F_i(2,90) = 21.27$, $p < .001$]. Clearly, the particular attributes and relations were primed, rather than the attributive and relational processes. This result is inconsistent with the CARIN model, which claims that relation priming will not occur in the absence of lexical repetition of the modifier. The main effect of Target-type was significant in the participant analyses [latency $F_s(1,32) = 5.51$, $p < .05$ and judgment $F_s(1,32) = 14.32$, $p = .001$], but not in the item analyses [latency $F_i(1,45) = .91$, $p > .30$ and judgment $F_i(1,45) = 2.26$, $p > .10$]. The interaction of Prime-type and Target-type was not reliable, all $p > .10$.

As Gagné (2001) pointed out, one tenable explanation of priming effects in this paradigm is simple associative priming. That is, the strength of the semantic associations between prime and target combinations might have differed across the different experimental conditions. Specifically, it might have been the case that the primes in the Same condition (e.g., SWIMMING FLIPPERS) were simply more strongly associated with the targets (e.g., RUGBY SHOES) than were the primes in the Different condition (e.g., ROAD CONSTRUCTION). If so, then this differential associative priming might have facilitated comprehension of the targets in the Same condition, thereby producing a pattern of results that *appeared* to reflect relation priming, but was instead due to associative priming. Therefore, a post-test used the double lexical decision task to investigate the possibility of differential associative priming between conditions (cf. Gagné, 2001). In this task, participants read combinations of letter strings and indicate whether both letter strings are words, or whether one is a nonword. If the results of the experiment proper can be explained by differential associative priming, then one should find a main effect of Prime-type similar to that in the experiment proper.

Post-test. Twenty-six Princeton University undergraduates participated in the post-test. The materials were identical to those of the experiment proper, with the exception that the 48 nonsense fillers (e.g., CHOIR PEEL) were replaced with 48 nonword fillers (e.g., CHOIR KIGRO). The fillers all contained one word and one orthographically legal nonword. The nonword appeared as the first letter string in half of the filler combinations (e.g., FANINK PHONE), and as the second letter string in the other half (e.g., CHOIR KIGRO). The procedure was similar to that of Experiment 1, except that participants pressed one or another key to indicate whether or not both letter strings were words in the English language.

The data were submitted to an initial three (Prime-type: baseline, same, different) × two (Target-type: attributive, relational) ANOVA. Incorrect trials, and trials in which response times were greater than 2.5 standard deviations from a participant's mean, were excluded from response time analyses. The Prime-type × Target-type interaction did not approach significance in either latency or accuracy (both $p > .69$). There was a significant main effect of Target-type on response time [$F(1,25) = 18.04$, $p < .001$], such that attributive targets ($M = 968$) were responded to more slowly than relational targets ($M = 915$). Although this main effect might account for the main effect of Target-type in the experiment proper, this effect was not central to the present concerns or the concerns of the experiment proper. Most importantly, the main effect of Prime-type was nonsignificant (latency $p = .12$ and accuracy $p = .10$). That is, there was no reliable difference amongst the Baseline (latency $M = 962$, accuracy $M = .98$), Same (latency $M = 925$, accuracy $M = .98$), and Different (latency $M = 937$, accuracy $M = .96$) prime-types. However, because these probabilities were "marginally" significant, paired comparisons were conducted on the Same and Different prime-types. These comparisons revealed that the difference between those two prime-types did not approach reliability. For the attributive combinations, latency $t(25) = .76$, $p = .46$ and accuracy $t(25) = 1.00$, $p = .33$. And for the relational combinations, latency $t(25) = .34$, $p = .74$ and accuracy $t(25) = 1.22$, $p = .23$. Clearly, then, the main effect of Prime-type obtained in Experiment 2 was not due to differential semantic associations between conditions, and hence was not due to an item artifact. Rather, specific attributes and relations were primed in the experiment proper, even in the absence of prime-target lexical redundancy.

Discussion

Gerrig and Murphy (1992) obtained relation priming without lexical repetition in an offline task, but Gagné (2001) failed to find relation priming without lexical repetition in an on-line task. Wisniewski and Love (1998), meanwhile, demonstrated that the process of attributive or relational comprehension can be primed. The present experiment has shown relation priming in the absence of lexical repetition in an on-line task, and has shown furthermore that this relation priming was specific to the particular relation, and was not simply due to priming of the relational process more generally, nor to associative priming. Although I do not claim to provide a full reconciliation of all of these disparate results, I will nonetheless attempt to clarify them.

First, by ruling out process priming as an explanation of the present results, I do not intend to imply that priming the attributive or the relational process is impossible. On the contrary, it was precisely because Wisniewski and Love (1998) have shown that process priming can occur that I sought to rule it out as an explanation of the present results. Notice, however, an important difference between the present study and that by Wisniewski and Love: Target combinations in the present experiment were preceded by only one prime combination, whereas target combinations in Wisniewski and Love's (1998) experiment were preceded by 10 prime combinations. The most parsimonious and intuitive explanation is that process priming simply requires more than one prime combination. The lack of process priming in the present experiment, therefore, was likely due to an insufficient number of prime combinations (i.e., one) that failed to induce the processing bias. If so, the present results are entirely compatible with those of Wisniewski and Love (1998).

Second, the present results are also compatible with those of Gerrig and Murphy (1992). In that study, prime combinations such as TRUMPET OLIVE facilitated comprehension of target combinations such as KITTEN APPLE. Although there is no way to determine from Gerrig and Murphy's experiment, the fact that their context stories used only one prime combination suggests that their result was likely due to the priming of particular relations rather than process priming. If this is true, then the present results and those of Gerrig and Murphy provide converging evidence of relation priming without lexical repetition.

It is less clear how Gagné's (2001) result comports with the present results or those of Gerrig and Murphy (1992). Recall that Gagné failed to find relation priming without lexical repetition, while Gerrig and Murphy did obtain that effect. The experiment by Gerrig and Murphy differed from Gagné's in two potentially relevant respects. Gagné's experiment lacked context stories, and it was an on-line task rather than an off-line task. These two factors constitute the most obvious explanations of the different results between those experiments. The present experiment, though, rules both of those explanations out. The present experiment was in fact modeled after Gagné's—it also lacked context stories, and used an on-line task. Yet, relation priming was obtained here in the absence of lexical repetition. One feasible explanation of this discrepancy is that the similarity of the relations between the prime and target combinations in the present study and in Gerrig and Murphy (1992) might have been greater than that in Gagné's (2001) study. For instance, in the present experiment, the conceptual relation required to comprehend the prime SWIMMING FLIPPERS is quite specific and similar to the relation used in the target RUGBY SHOES. And in Gerrig and Murphy (1992), the relation used in TRUMPET OLIVE is very similar that used in KITTEN APPLE. The similarity of the relations in Gagné's study, however, might not have been as great. In her study, for instance, the combination BABY CRY served as a prime for the target STUDENT VOTE. My intuition is that these two combinations have less similar relations. Unfortunately, because a complete list of Gagné's materials was not published, I cannot say whether this speculation generalizes to the rest of her items. But if it does, then it might simply be that conceptual relations are more specific than Gagné assumed, and hence some of her primes might not have activated the relation used by the target. To be sure, however, this explanation remains purely speculative, and more research must be conducted to properly resolve this disparity.

General discussion

Experiment 1 provided a critical test of the dual process theory and the relational theory of nominal combination. According to relational theory, the infrequency of attribution should have led attributive primes to interfere with comprehension of relational targets, because the attributive process would not have been attempted prior to the relational process otherwise. The attributive prime would therefore induce additional processing. But preceding an attributive target with a relational prime should not interfere, because the relational process would have been attempted serially prior to the attributive process anyway, and hence no additional processing would be induced by the prime. Thus, if relational comprehension is serially prior to attributive comprehension, then an interaction should have obtained. If, on the other hand, attributive and relational processes occur in parallel, then one would expect similar patterns of facilitation and interference across combination types. The data reported in Table 1 support such a parallel model of nominal combination (i.e., Wisniewski, 1997, 2000).

In Experiment 1, target combinations were more likely to be comprehended, and were comprehended more quickly, when preceded by prime combinations that used the same attribution or relation, despite the absence of lexical overlap between prime and target combinations. For instance, comprehension of MOUNTAIN SNAKE was facilitated by the prior presentation of JUNGLE BIRD. This finding is inconsistent with the CARIN model of nominal combination, which claims that conceptual relations are only activated in the context of their particular modifiers. If relations are modifier-specific, then relation priming should occur only when the modifier is repeated from prime to target (Gagné, 2000, 2001; Gagné & Shoben, 1997). The results of Experiment 1 appeared to contradict this prediction. However, because Experiment 1 manipulated whether the prime and target combinations used the same process or a different process, it might have been the case that the differential priming observed in the experiment was due to the priming of the attributive or relational *process* rather than a *particular* attribute or relation. Therefore, Experiment 2 held constant the process engaged by the prime and target combinations, and manipulated whether the particular attribution or the particular relation was repeated from prime to target. For instance, the target COPPER HORSE was preceded either by GLASS ROSE, which uses the same relation, or by SIGN POST, which uses a different relation. The results showed that target combinations were again comprehended faster and more easily when following a prime combination that used the same attribution or relation than when following a prime that used a different attribution or relation, even though prime and target combinations engaged the same process. Thus, Experiment 2 demonstrated that the prime combinations activated particular attributes and relations, rather than the attributive process or the relational process more generally. So in addition to the lack of evidence for the claim that the relational process serially precedes the attributive process, there was also evidence against the claim that relation priming occurs only when the modifier is repeated from the prime to the target.

Furthermore, Experiment 2 may be viewed as a test of whether a *resemblance* relation can plausibly account for attributive comprehension. For, if all attributive interpretations derive from the *resemblance* relation, then both the Same and Different prime combinations in Experiment 2 should have facilitated comprehension of their attributive targets, because both prime-types would reputedly involve the same general *resemblance* relation. Therefore, the fact that the primes in the Same condition facilitated comprehension of the attributive targets, while the primes in the Different condition failed to do so, suggests that the Same and Different prime combinations were not comprehended by the same relation (i.e., *resemblance*).

It should be noted that although the data failed to support the CARIN model (Gagné, 2000, 2001), they do not undermine the importance of conceptual relations in nominal combination (e.g., Gagné & Shoben, 1997). Rather, the current data simply emphasize that the relational process may not account for all nominal combination. There appear to be distinct attributive and relational processes, and these two processes appear to operate concurrently.

Although the focus of this research has been on processing, the present results also have implications for representation. The finding of relation priming in the absence of lexical repetition strongly suggests that conceptual relations are not "stored with" or represented with their specific modifier concepts, as Gagné (2000, 2001; Gagné &Shoben, 1997) has argued. If the comprehension of a target combination is facilitated by the prior presentation of a prime combination that uses the same conceptual relation, but lacks any lexical overlap with that target, then there is little reason to believe that the conceptual relation is concept-specific. On the contrary, this result suggests that the conceptual relation is an independent representational structure, one that can be activated by a prime and utilized by a target. This account neatly explains how a particular relation can be primed without priming the entire relational process: Only the specific relation activated by the prime is made more accessible for subsequent use. This account might also explain Gagné's (2001) failure to obtain relation priming in the absence of lexical redundancy: Her prime (e.g., BABY CRY) and target (e.g., STUDENT VOTE) combinations might not have used the same specific relation. Clearly, this is speculative, and more research needs to address this issue. But nonetheless, the results do suggest that the representation of semantic relations is concept-independent rather than concept-specific.

The emergence of a general model of nominal combination

As a consequence of the accumulation of data, and despite the ongoing theoretical debate in this field, it appears that a general model of nominal combination is beginning to emerge. Concerning the nature of the attributive process, there is strong empirical support for Wisniewski's (1997, 2000) comparison-and-mapping model. In a recent study, Wisniewski and Middleton (2002) had participants draw pictures depicting the referents of nominal combinations. For instance, one group of participants was told that "a PORCUPINE MUSHROOM is a prickly mushroom," while another group was told that "a ROSE MUSHROOM is a prickly mushroom." Interestingly, even though both groups were instructed to draw a prickly mushroom, the drawings produced by the two groups of participants differed systematically. When drawing a PORCUPINE MUSHROOM, participants tended to draw thorns on the top of the mushroom, akin to the quills on the back of a porcupine. But when drawing a ROSE MUSHROOM, participants located the thorns on the stem of the mushroom, as they are on the stem of a rose. This finding provides convincing evidence that when comprehending attributive combinations, people compared the constituent concepts and used this comparison as a basis for mapping a property of the modifier onto the head concept in a perceptually analogous manner.

However, some details of the model are in need of specification or revision. Recall, for instance, that Wisniewski's model predicts a positive correlation between constituent similarity and the likelihood of attribution (Wisniewski, 1996). That is, because similarity facilitates comparison (Gentner & Markman, 1994), and because comparison is necessary for attribution (Wisniewski, 1997, 2000), similarity should facilitate attribution (Wisniewski, 1996). Indeed, several studies have borne out this prediction (Markman & Wisniewski, 1997;

Wilkenfeld & Ward, 2001; Wisniewski, 1996; Wisniewski & Love, 1998). But several others have failed to support it (Bock & Clifton, 2000; Costello & Keane, 2001; Estes & Glucksberg, 2000a; see also Gagné, 2000). Those studies that have accurately predicted deviations from the correlation have all found that modifier concepts with particularly salient (or diagnostic) properties tend to produce attributive interpretations, regardless of their similarity (or alignability) to the head concept. For instance, CACTUS CARPET is most frequently interpreted as "a prickly carpet," despite the fact that cacti and carpets are judged highly dissimilar (Estes & Glucksberg, 2000a). Although the comparison-and-mapping model may account for this finding (Wisniewski, 2000), the studies question whether alignment and comparison are *necessary* to account for attribution, or whether a salience-based model might be more parsimonious (Estes & Glucksberg, 2000b). Nonetheless, the comparison-and-mapping model generally provides a sound basis for models of property attribution.

Wisniewski's thematic role-filling hypothesis also provides a sound basis for models of the relational process. It is clear that when a combination is interpreted relationally, the two concepts fulfill different functional roles in that relationship (see also Murphy, 1990). But because most of the support for the dual-process model comes from investigations of the attributive process, the model remains underspecified in its claims about this more ubiquitous relational process. One specific aspect of this relational process is addressed by Gagné and Shoben's (1997) model. In particular, the CARIN model attempts to explain the selection of an appropriate relation between concepts. And in fact, convincing experimental results support the model's general claims that relational comprehension entails selection from a small set of extant conceptual relations (Coolen et al., 1991), and that this selection of relations is sensitive to prior use (Gagné, 2000, 2001; Gagné & Shoben, 1997). Where the model requires revision, according to the present results, is in that conceptual relations appear to be represented independently of any particular concept, and hence can be primed in the absence of lexical redundancy.

A general model of the relational process, then, begins to emerge. First, the most accessible conceptual relation is selected. The accessibility of a given relation, in turn, may be determined by that relation's frequency of use with the particular constituent concepts (Gagné, 2000, 2001; Gagné & Shoben, 1997), or by the recency of activation of that relation (Experiments 1 and 2 above). Once a relation is selected, it is then subject to the satisfaction of thematic role constraints. That is, the constituent concepts must plausibly perform different functional roles in the relationship defined by the selected relation (Murphy, 1990; Wisniewski, 1997). If this constraint is satisfied, then an interpretation has been reached. If not, then another relation is selected and subjected to the constraint (either iteratively or simultaneously). To illustrate, suppose that the target CHOCOLATE BALL is preceded by the prime TIRE SHOP. Relational comprehension might proceed as follows. First, because the *sell* relation has been activated by the prime (i.e., TIRE SHOP), it will be highly accessible, and therefore, may be selected immediately. Next, CHOCOLATE and BALL are considered to determine whether they could plausibly play complementary roles in the *sell* relation. However, because BALL cannot *sell* CHOCOLATE, and CHOCOLATE cannot *sell* BALL, the *sell* relation is rejected as inappropriate for comprehension. Thus, a new relation must be considered. Now, because the *composition* relation is highly frequent for the modifier CHOCOLATE, that relation might also be highly accessible, and consequently might be selected for constraint satisfaction. Because BALL can be *composed of* CHOCOLATE, that relation satisfies the thematic constraints, and an interpretation has been reached. In this way, it is conceivable that a modified version of Gagné's model of relation selection could act in conjunction with Wisniewski's role-filling model to provide a fuller explanation of the relational process.

If the attributive and relational processes do occur in parallel, as suggested by Wisniewski (1997, 2000; Wisniewski & Love, 1998) and as supported by the experiments reported above, then yet another process becomes necessary: Supposing that the attributive and relational processes both produce an interpretation of a nominal combination, an interpretation selection process will be needed. That is, if the processes were carried out serially, then the search for an interpretation could simply be terminated once an interpretation had been reached. But if the processes occur in parallel, then there may be cases in which both processes achieve an interpretation. Indeed, Murphy (1990) and Costello and Keane (1997) have shown that many nominal combinations are polysemous, eliciting a multitude of interpretations. In such cases, one must select the most appropriate interpretation for comprehension. How might this interpretation selection be accomplished? Clearly, context will play an important role. Costello and Keane (2000) suggest that the interpretation of a nominal combination is subject to the satisfaction of pragmatic constraints (cf. Grice, 1975); namely, the best interpretation of a combination is the one that is most plausible, informative and diagnostic (i.e., contains important features of both constituent concepts; see also Hampton, 1988). By "plausible" Costello and Keane do not mean that the two concepts can play different thematic roles in the functional relationship, as described above. Rather, Costello and Keane mean that the interpretation must be consistent with the pragmatic context in which the nominal combination occurs. As an example, in one experiment, Gerrig and Bortfeld (1999) established that one group of nominal combinations (e.g., DOLL SMILE) was comprehended faster than another group of combinations (e.g., BASEBALL SMILE). But in another experiment they showed that, when preceded by a context paragraph that implied a novel relation for the target combinations (e.g., DOLL SMILE as "a smile *elicited by* a doll"), both types of combination were comprehended at about the same rate. In this case, there is a more accessible meaning of DOLL SMILE out of context (i.e., "a smile on the face of a doll"). But in the context of a story in which a doll—given as a gift—has elicited a smile from its recipient, the novel relation is more plausible. Thus, pragmatic constraints may affect which interpretation is chosen as the appropriate one.

In conclusion, then, research on nominal combination is beginning to converge on a general model. It may well be that attribution often entails comparison and mapping (Wisniewski, 1997), which is driven by salient or diagnostic properties of the modifier (Bock & Clifton, 2000; Costello & Keane, 2001; Estes & Glucksberg, 2000a), and that relational comprehension entails thematic role-filling (Wisniewski, 1997), which is driven by the prior use of particular conceptual relations (Gagné & Shoben, 1997). Finally, interpretations produced by these processes are undoubtedly subject to the satisfaction of pragmatic constraints (Costello & Keane, 2000; Gerrig & Bortfeld, 1999). The contribution of the present research has been to emphasize this distinction between the attributive and relational processes by demonstrating their parallel occurrence (Experiment 1), and to illustrate the specificity of property attribution and relational inference (Experiment 2).

Acknowledgments

I am indebted to Sam Glucksberg, James Hampton and Uri Hasson for their insightful comments throughout the course of this research. I also thank three anonymous reviewers for their cogent comments on the manuscript.

Appendix A. Stimuli, Experiments 1 and 2

Target-type	Prime-type		
	Same	Different (Exp. 1)	Different (Exp. 2)
Attributive			
cactus carpet	needle grass	thank-you note	mirror pond
canary crayon	lemon paint	lounge beer	razor insult
rocket sprinter	bullet train	table vase	skunk cigar
sandpaper skin	brillo beard	grill steak	alarm-clock rooster
blimp belly	balloon pregnancy	sailing hat	catapult promotion
zebra clam	tiger paper	bacon tongs	medicine music
rock bagel	iron fist	wood stove	iceberg symptom
sponge memory	warehouse brain	motorcycle documentary	dinosaur computer
pig-sty bedroom	junkyard desk	gambling policy	boomerang lie
sleeping-pill personality	sedative voice	council mandate	magnet smile
butterfly ballerina	swan athlete	student artwork	time-bomb temper
gymnast squirrel	acrobat cat	expert advice	handlebar mustache
shark politician	piranha lawyer	swimming flippers	feather luggage
molasses traffic	turtle jogger	glass rose	chameleon ninja
strawberry ink	tomato dye	plastic bowl	twig legs
pillow lips	silk hair	ocean nausea	mushroom cloud
leech boyfriend	parasite citizen	mustard spot	umbrella tree
vampire insect	mosquito fly	baseball injury	volcano pimple
octopus chair	centipede table	gold flakes	funnel tornado
ice marriage	refrigerator parents	raisin cereal	gem idea
inferno radiator	fire coffee	laser surgery	hourglass body
thunder applause	megaphone talker	jungle bird	elbow macaroni
see-saw relationship	roller-coaster emotions	peanut oil	soldier ant
butter grip	grease fish	video surveillance	pie chart
Relational			
sympathy bouquet	thank- you note	needle grass	family vacation
patio cigarette	lounge beer	lemon paint	servant scandal
floor television	table vase	bullet train	cream sauce
microwave sandwich	grill steak	brillo beard	tax form
bowling sweater	sailing hat	balloon pregnancy	coal town

Appendix A. (*continued*)

Target-type	Prime-type		
	Same	Different (Exp. 1)	Different (Exp. 2)
pancake spatula	bacon tongs	tiger paper	city riots
gas oven	wood stove	iron fist	party dance
rodeo magazine	motorcycle documentary	warehouse brain	cracker tray
battle theory	gambling policy	junkyard desk	sex hormone
employee vote	council mandate	sedative voice	sugar tea
prisoner graffiti	student artwork	swan athlete	song book
doctor testimony	expert advice	acrobat cat	finger nerve
rugby shoes	swimming flippers	piranha lawyer	road construction
copper horse	glass rose	turtle jogger	sign post
concrete fountain	plastic bowl	tomato dye	evening prayer
onion tears	ocean nausea	silk hair	winter breeze
burrito stain	mustard spot	parasite citizen	chest rash
boxing bruise	baseball injury	mosquito fly	milk virus
chocolate chunks	gold flakes	centipede table	nose sound
honey soup	raisin cereal	refrigerator parents	gender comedy
scalpel incision	laser surgery	fire coffee	flower season
mountain snake	jungle bird	megaphone talker	stereo headphone
daisy extract	peanut oil	roller-coaster emotions	conference chairman
telescope observation	video surveillance	grease fish	pasta jar

Note. The Different (Exp. 1) primes used a different process (i.e., attributive or relational) than their targets, while the Different (Exp. 2) primes used the same process as their targets, but the particular attribution or relation was different.

References

Bock, J. S., & Clifton, C. (2000). The role of salience in conceptual combination. *Memory & Cognition, 28*, 1378–1386.

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

Costello, F., & Keane, M. T. (1997). Polysemy in conceptual combination: Testing the constraint theory of combination. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 137–142). Hillsdale, NJ: Erlbaum.

Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science, 24*, 299–349.

Costello, F. J., & Keane, M. T. (2001). Testing two theories of conceptual combination: Alignment versus diagnosticity in the comprehension and production of combined concepts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 255–271.

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

Estes, Z., & Glucksberg, S. (1998). Contextual activation of features of combined concepts. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 333–338). Mahwah, NJ: Erlbaum.

Estes, Z., & Glucksberg, S. (1999). Relevance and feature accessibility in combined concepts. In *Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society* (pp. 149–154). Mahwah, NJ: Erlbaum.

Estes, Z., & Glucksberg, S. (2000a). Interactive property attribution in concept combination. *Memory & Cognition, 28*, 28–34.

Estes, Z., & Glucksberg, S. (2000b). Similarity and attribution in concept combination: Reply to Wisniewski. *Memory & Cognition, 28*, 39–40.

Franks, B. (1995). Sense generation: A "quasi-classical" approach to concepts and concept combination. *Cognitive Science, 19*, 441–505.

Gagné, C. L. (2000). Relation-based combinations versus property-based combinations: A test of the CARIN theory and 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., & Murphy, G. L. (1996). Influence of discourse context on feature availability in conceptual combination. *Discourse Processes, 22*, 79–101.

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.

Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science, 5*, 152–158.

Gerrig, R. J., & 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). Conceptual influences on the comprehension of complex concepts. *Language and Cognitive Processes, 7*, 205–230.

Gleitman, L., & Gleitman, H. (1970). *Phrase and Paraphrase: Some Innovative Uses of Language*. New York: Norton.

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

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics: Speech Acts* (Vol. 3). New York: Academic Press.

Hampton, J. A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*, 12–32.

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

Luchins, A. (1942). Mechanization in problem solving: The effect of Einstellung. *Psychological Monographs, 54*(1), entire.

Markman, A. B., & Wisniewski, E. J. (1997). Similar and different: The differentiation of basic level categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 54–70.

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

Springer, K., & Murphy, G. L. (1992). Feature availability in conceptual combination. *Psychological Science, 3*, 111–117.

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

Wilkenfeld, M. J., & Ward, T. B. (2001). Similarity and emergence in conceptual combination. *Journal of Memory and Language, 45*, 21–38.

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 and Review, 4*, 167–183.

Wisniewski, E. J. (2000). Similarity, alignment, and conceptual combination: Comment on Estes and Glucksberg. *Memory & Cognition, 28*, 35–38.

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

Wisniewski, E. J., & Love, B. C. (2000). A dual-process account of conceptual combination and metaphor. Paper presented at the 41st Annual Meeting of the Psychonomic Society, New Orleans.

Wisniewski, E. J., & Middleton, E. L. (2002). Of bucket bowls and coffee cup bowls: Spatial alignment in conceptual combination. *Journal of Memory and Language, 46*, 1–23.