JOURNAL OF MEMORY AND LANGUAGE **38,** 177–202 (1998)
ARTICLE NO. ML972550

Relations versus Properties in Conceptual Combination

Edward J. Wisniewski and Bradley C. Love

Northwestern University

We evaluate the hypothesis that people interpret noun-noun combinations by first attempting to find a plausible relation in which the constituents play different functional roles (as in ''a hawk that *preys* on robins'' for *robin hawk*). If they do not succeed, they then attempt as a last resort to derive an interpretation based on properties of the modifier (''hawk with a *red breast*''). In one study, combinations with plausible relations were more often interpreted with properties when their constituents were highly similar. In a second study, property and relation interpretations of the same combination were selectively primed. In a final study, we show that property interpretations characterize the meanings of a number of familiar combinations in the English language. Taken together, the results of these studies run counter to the view that concepts are combined by first attempting to link them by a relation. We describe other reasons to doubt this view and discuss the implications of our findings for models of conceptual combination.

People frequently combine concepts to produce new coherent representations, as in understanding a sentence. In cognitive psychology, there has been much recent interest in one particular aspect of conceptual combination: how people interpret novel noun-noun combinations (Coolen, van Jaarsveld, & Schreuder, 1991; Gerrig & Murphy, 1984; Gagne & Shoben, 1997; Hampton, 1987; Murphy, 1988, 1990; Shoben, 1993; Shoben & Gagne, 1997; Smith, Osherson, Rips, & Keane, 1988; Wisniewski, 1996, 1997). A general aspect of language is that people create novel combinations and listeners have little trouble comprehending them. For example, one of the authors easily understood the novel combinations *ostrich steak*, *ostrich burger*, *ostrich meat*, and *ostrich ranch*, which were recently mentioned in a brief news story about the increasing popularity of *ostrich farming* in this country.

Studying how concepts combine is important for at least two reasons. First, it provides a methodology for investigating the nature of concepts (Hampton, 1987; Markman & Wisniewski, 1997; Medin & Shoben, 1988; Murphy, 1988). For example, studies of conceptual combination have identified ways in which prototype theories need to be extended (Medin & Shoben, 1988) as well as differences between the structure of superordinates and basic level concepts (Markman & Wisniewski, 1997).

Second, novel combinations typically occur in communicative contexts and serve a variety of functions (see Downing, 1977; Wisniewski, 1997, for detailed discussions). They are used to designate significantly new categories, as in the examples involving *ostrich* above. These phrases use new combinations of old terms to capture changes in reality which may become commonplace. Combinations also convey information in a concise and efficient way. For example, *football parking* designates an area for parking one's car while attending a football game. Even though this phrase is elliptical, readers generally understand what it means. Finally, combinations function as anaphora in that they are used to refer back to a previous referent. In doing so, they help link information to the appropriate referent

The second author was supported by the Office of Naval Research under the National Defense Science and Engineering Graduate Fellowship Program. The first author thanks Bob Dylan for providing part of the inspiration for this work.

Address correspondence and reprint requests to Edward J. Wisniewski, Department of Psychology, Northwestern University, Evanston, IL 60208.

0749-596X/98 $25.00

and establish cohesion in discourse contexts (Garrod & Sanford, 1994). Studying novel combinations can illuminate how these functions are achieved.

In order to develop a processing-level (algorithmic) account of *how* people combine concepts, one must identify the important phenomena of conceptual combination. These phenomena define the generality to be attained by a model and help to constrain its processing assumptions. Recently, Wisniewski (1996) examined a large sample of people's interpretations of novel combinations with the goal of identifying the significant phenomena associated with conceptual combination. People's interpretations fell into two basic categories. One kind of interpretation involved a *thematic relation* between the referents of the modifier and head concepts. For example, a *robin hawk* could mean ''a hawk that preys on robins.'' In *property* interpretations, people asserted that one or more properties of the modifier concept apply in some way to the head concept, as in ''hawk with a red breast,'' for *robin hawk.* Sometimes these interpretations refer to an entity which shares many properties of *both* constituents (e.g., a *robin hawk* could refer to a bird that is cross between a robin and a hawk).

Wisniewski (1996) also noted that there were important conceptual distinctions between these kinds of combinations. For example, in *robin hawk* interpreted with a thematic relation as ''a hawk that preys on robins,'' the modifier refers to a robin and the head noun refers to a type of hawk. The meaning of the combination specifies a relationship in which the entities play different functional roles. In this example, *robin* would play the role of the *prey* and *hawk* the role of the *predator*. However, in *robin hawk* interpreted with a property as ''a hawk with a red breast,'' the modifier refers to a property of a robin rather than to a robin. There is no relationship between a robin and a hawk in which they play different functional roles. Rather, the meaning refers to a hawk that resembles a robin in some respect (i.e., in having a red breast like a robin does).

Although subjects produce these two types of interpretations in the laboratory, researchers have assumed that only thematic relation interpretations are psychologically significant. In particular, a popular view is that people first attempt to derive an interpretation that involves a thematic relation and only derive a property interpretation if there is no plausible relation between the usual referents of the modifier and head noun concepts (Downing, 1977; Gagne & Shoben, 1997; Shoben & Gagne, 1997; Wisniewski & Gentner, 1991). On this view for example, people would interpret *robin hawk* by first attempting to find a thematic relation in which *robin* and *hawk* plausibly play different functional roles (as in ''hawk that preys on robins'') and *only if* there was no plausible relation would they attempt a property interpretation (as in ''hawk with a red breast'').

There are two arguments that have led to this hypothesis. One is derived from the claim that people rarely produce property interpretations in communicative contexts. Downing (1977) apparently did not find property interpretations among novel combinations that she sampled from several kinds of discourse materials. However, when she gave subjects novel combinations, they interpreted a number of them in this manner. To explain this asymmetry, Downing suggested that an interpretation involving a property:

> . . . may be a *last-resort* [italics added] or default relationship, considered by the addressee when no other useful underlying relationship seems plausible; and it also suggests a difference in speaker and hearer strategies. (p. 830)

Like Downing, Shoben and Gagne (1997) doubt that occurrences of property interpretations are very common in everyday speech and writing. They come to similar conclusions about these interpretations:

> We do not dispute that property matches occur; however, we contend that property matches are the interpretation of *last resort*. . . . [italics added] We also doubt that the occurrence of these property matches is very common in everyday speech and

> writing. None for example, appeared in our corpus. (p. 35)
>
> . . . these kinds of interpretations are secondary to the strategy of relation assignment. That is, people will *always* [italics added] try to assign a thematic relation to a combination; however if they are unable to assign a relation, they seek a metaphorical or property matching solution. (p. 47)

Wisniewski and Gentner (1991) put forth a different argument for the primacy of relation interpretations. They noted that relation interpretations generally preserve the meaning of each noun in the combination whereas property interpretations use the modifier in an extended sense (i.e., the modifier refers to one of its properties). For example, the relation interpretation of *book magazine* as ''a magazine that reviews books'' preserves the usual referents of *book* and *magazine.* In contrast, a property interpretation of *book magazine* such as ''a thick magazine'' refers to a property of books rather than to a book. Given the general goals of communication, people may initially assume that a combination involves the usual meanings of its constituents and thus first look for a relation interpretation. Like other researchers, Wisniewski and Gentner (1991) suggest that interpreting combinations by relations is ''a very natural strategy for combining concepts'' and that ''people prefer this strategy'' (p. 266).

This view of the primacy of relation interpretations also manifests itself in processing models of conceptual combination. Several approaches exclusively focus on how a relation is determined which plausibly holds between the usual referents of the constituents of a combination. For example, in the concept specialization model (Cohen & Murphy, 1984; Murphy, 1988) the modifier and head nouns are represented as schemata or frames. One interprets a noun-noun combination by filling a slot of the head noun with the modifier noun. Thus, one might interpret *robin hawk* by filling the *prey* slot in *hawk* with the modifier concept *robin* in deriving the meaning ''a hawk that preys on robins.'' In effect, this slot-filling process captures a relation between the usual referents of the constituents. However, this approach does not address the interpretations in which the modifier noun is taken to refer to a property.

In another approach, called the ''competition among relations in nominals'' (CARIN) model, people use distributional knowledge of thematic relations associated with how a noun has previously combined to interpret a combination containing that noun (Gagne & Shoben, 1997; Shoben & Gagne, 1997). In addition, the modifier's combinatorial history has a greater influence on interpretation (see Gagne & Shoben, 1997, for evidence). For example, when *mountain* is used as a modifier it typically instantiates a locative relation (as in *mountain stream, mountain resort, mountain goat,* etc.) and is only rarely involved in other types of relations (*mountain range*). Therefore, people are biased to interpret a novel combination such as *mountain fish* as ''fish found in the mountains'' by using their knowledge that *mountain* has previously combined with other nouns in a similar manner. Like the concept specialization model, the CARIN model does not address interpretations in which the modifier noun is taken to refer to a property but rather captures interpretations involving a relation between the usual referents of the constituents.

Although a number of researchers have argued that relation interpretation is the basic strategy for combining nouns, we are unaware of any research that directly evaluates this claim. Some studies show that property interpretations are not infrequent (e.g., Markman & Wisniewski, 1997; Wisniewski, 1996; Wisniewski & Markman, 1993). On first glance, these findings might argue against the primacy of relation interpretations. However, in all of these studies, subjects were given experimenter-constructed combinations whose constituents were either arbitrarily paired or were highly similar. Both of these factors are very likely to yield combinations which do not have plausible relation interpretations (which could then lead people to generate property interpretations as a last resort). In the case of combinations with highly similar constituents (e.g., *whiskey beer*), it is often the case that a

plausible relation does not exist between the constituents (though there are exceptions, as Experiment 1 will show). In general, highly similar things play the same role in a relation. Because highly similar things are so alike, they also tend to share those characteristics which are required of the potential filler of a role. For example, both *whiskey* and *beer* can play the object role in the drink relation because they are both beverages that one can drink. However, in order to derive a plausible relation interpretation, the constituents must play different roles in the relation. Thus, *whiskey beer* cannot be interpreted with the drink relation because both constituents are plausible fillers for the (same) object role. In contrast, *wine drinker* does have this relation interpretation because *wine* and *drinker* can play different roles: *wine* can play the object role whereas *drinker* can play the agent role.

To examine this view, we re-examined Wisniewski's (1996; Experiment 2) stimuli and intuitively assessed whether there were reasonably plausible relations between the constituents. In our judgment, there was no plausible relation between the constituents of 31 of the 32 combinations that were highly similar (e.g., *pistol rifle, organ piano, apple pear, robin canary*) and no plausible relation between the constituents of 16 of the 32 combinations which were arbitrarily paired (e.g., *fork scarf, stool hotel, saxophone couch, pineapple piranha*). We then compared Wisniewski's classification of the interpretations for the combinations to our judgments about the constituents of these combinations. For those constituents we judged not to have a plausible relation, only 12.3% of the interpretations of their corresponding combinations involved relations (the others involved properties). For constituents that we judged to have plausible relations (e.g., those which formed *book pamphlet, apartment piano, motorcycle screwdriver, cow cabbage*), 66.5% of the interpretations of their corresponding combinations involved relations. Thus, although Wisniewski (1996) found that subjects readily produce property interpretations, our analysis suggests that the combinations with these interpretations did not have plausible thematic relations between their constituents. As a result, these interpretations could have been generated as a last resort when a plausible relation did not exist. At the very least, it is unclear from previous research whether property interpretation is an important strategy for combining nouns, outside of the laboratory.

In this paper, we evaluate the assumptions that (a) relation interpretation is the basic strategy used to understand combinations and that property interpretations are derived only as a last resort, and (b) people have a strong bias to produce combinations involving relations in which the constituents play different, functional roles. Together, we refer to these assumptions as the *last resort hypothesis.* In evaluating this hypothesis, we will examine the possibility that both relation and property strategies are important and useful in combining concepts and that people are more flexible in both their interpretation and production of combinations. Resolving these issues has important theoretical implications for research on conceptual combination. Finding that property interpretation is a strategy of last resort would imply that it is not a significant phenomenon to be seriously addressed by theories or models of conceptual combination. Rather, it tends to occur under very special circumstances (e.g., when researchers ask subjects to interpret experimenter-selected novel combinations that do not have plausible relation interpretations). Alternatively, finding that property interpretation is an important and useful strategy suggests a number of important issues to be addressed by theories and models. For example, what factors influence the tendency to use one or the other strategy? Why do people use different strategies?

We conducted several experiments to examine these issues. In the first study, we revisited the effect of high constituent similarity on interpretation. Even though high similarity typically prevents relation interpretations, a number of studies indirectly suggest that high similarity could also facilitate the use of the property interpretation strategy. Thus, we attempted to create combinations whose constit-

uents were both highly similar *and* could plausibly play different roles in a thematic relation. We used these combinations to assess whether interpretation is also a function of the similarity between the constituents of a combination and is not determined solely by whether there is a plausible relation (as suggested by the last resort hypothesis). In Experiment 2, we examined the effects of prior use of relation versus property interpretation strategies on their subsequent use. If strategies are flexibly applied, prior use of one strategy may influence the tendency to use that strategy later. Finally, in Experiment 3, we re-examined the relative prevalence of property and relation meanings associated with combinations that are part of the English language. As noted, Downing (1977) found an asymmetry between production and interpretation which formed the basis for her view that property interpretation was a last resort strategy.

EXPERIMENT 1

Previous research indirectly suggests that high similarity between the constituents of a combination facilitates the tendency to produce property interpretations. Many studies suggest that when people are explicitly instructed to compare mental representations, they use the commonalities between these representations to find their differences (Markman & Gentner, 1993ab; Markman & Wisniewski, 1997; Gentner & Markman, 1994). Several studies show that subjects not only list more commonalities for similar than for dissimilar concepts but also more differences. For example, people list more commonalities and differences for *kitten* and *cat* than for *kitten* and *newspaper*. Importantly, the differences are conceptually related to the commonalities, suggesting that finding commonalities leads to the finding of differences. For example, when subjects list ''has wheels'' as a commonality of *car* and *motorcycle* they also tend to list ''has four wheels versus two wheels'' as a difference (Markman & Gentner, 1993a). Even though virtually any property of the items in a dissimilar pair will be a difference, people find it easier to generate differences between similar concepts and more easily determine these differences (Gentner & Markman, 1994).

In regard to conceptual combination, some evidence suggests that people generate property interpretations by comparing the modifier to the head noun and looking for a difference between them that could form the basis of the interpretation (Wisniewski, 1996, Experiment 2). For example, people might interpret *zebra horse* by comparing *zebra* and *horse* and noting that a zebra has stripes but a horse does not. They then could use this difference to produce the interpretation ''horse with stripes'' for *zebra horse.* As previous research implies, it should be easier for people to find such differences (and more of them) for combinations with highly similar constituents. For example, it should be easier to find the difference ''has stripes versus does not have stripes'' between *zebra* and *horse* than between *zebra* and *clam.*

We can take advantage of this difference in the ease of finding differences between concepts to examine whether property interpretation is a last resort strategy. On this view, interpretation depends on there being a plausible relation in which the constituents can play different roles. In cases in which a plausible relation does exist between constituents, the influence of similarity on property interpretation should not be a factor because this strategy will only be triggered if a plausible relation does not exist. For example, people should interpret *book magazine* with a relation as in ''a magazine that reviews books'' or ''a magazine that sells books'' (assuming that these relations are plausible) and thus similarity will not affect property interpretation. Alternatively, if property interpretation is not a last resort strategy, then interpretation should also be a function of the similarity between the constituents and thus similar combinations should have property interpretations despite having plausible relations between their constituents.

Note that this alternative hypothesis is based on the important premise that plausible relations exist between the constituents of a

similar combination. As noted in the Introduction, plausible relations generally do not exist between highly similar entities. However, there are some highly similar nouns whose referents nevertheless can play different roles in the same thematic relation (they are difficult to find). For example, a musician could play music for dancers, with *musician* filling the *agent* role and *dancer* filling the *object* role. Thus, *dancer musician* could plausibly be interpreted as ''a musician who plays music for dancers.'' In the first study, we use such nouns to test the last resort hypothesis. Specifically, we created pairs of novel combinations with three important characteristics. First, in each pair one combination was similar (e.g., *dancer musician*) and the other was dissimilar (e.g., *mourner musician*). Second, both combinations could involve the same relation (e.g., ''plays for''). Third, according to our intuitions, it was at least as plausible that the referents of the highly similar constituents could be involved in that relation as those of the less similar constituents.

To test the last resort hypothesis, one group of subjects judged the plausibility of the similar entities being involved in the relation relative to the dissimilar entities being involved in that relation. For example, they were asked to judge the relative plausibility of ''a vulture that preys on dead geese'' (similar entities) compared to that of ''a vulture that preys on dead fish'' (dissimilar entities). This task allowed us to assess the relative plausibility of entities being involved in a relation, independent of the interpretation process (which we are suggesting is also influenced by the similarity of the constituents of a combination). The goal was to verify that there were plausible thematic relations between similar entities. A different group of subjects then interpreted the combinations whose constituents named those entities (e.g., *goose vulture, fish vulture*). We then examined if constituent similarity influenced interpretation even though plausible relations existed between similar entities named by the constituents.

Method

Subjects. The subjects were undergraduates from Northwestern University who participated as part of a course requirement. 23 subjects participated in the plausibility task and 34 subjects in the interpretation task.

Materials. 10 pairs of novel noun-noun combinations were used for the interpretation task. For each pair, one combination had highly similar constituents (e.g., *goose vulture*) and one combination had constituents that were less similar (e.g., *fish vulture*). According to our intuitions, the constituents of each combination of a pair could be plausibly linked by the same relation. For example, *goose vulture* could describe a vulture that eats a dead goose, and *fish vulture* could describe a vulture that eats a dead fish. Table 1 shows the pairs of combinations and the relation that was assumed to plausibly link the entities named by the constituents.

We also selected relations that people typically use in producing combinations. All of the relations are specific instantiations of the general relations that Levi (1978) proposed as characterizing most familiar combinations (see Table 1). They also fit the categories suggested by Downing (1977). In addition, Table 1 lists some familiar combinations whose meanings involve the same or similar relations.

We also attempted to use the same nouns for both the similar and dissimilar combinations and to use them in the same position as head noun or as modifier noun. We designed the stimuli in this way so as to insure that differences between interpretations were not due to differences between the nouns or their positions used in the conditions. As Table 1 shows, each pair of combinations involved the same head noun and there is 80% overlap between the modifiers of the similar and dissimilar combinations. (These constraints are difficult to achieve given the other constraints that must hold between the stimuli.)

In order to assess the plausibility of similar entities linked to a thematic relation relative to dissimilar entities linked by that relation, we constructed descriptions which linked the entities by the relation. These descriptions were typed onto paper with two descriptions listed side by side, corresponding to each pair

TABLE 1

Similar and Dissimilar Combinations and a Plausible Relation That Links Their Constituents (Experiment 1)

Similar combination	Dissimilar combination	Plausible relation	Familiar combination
dancer musician	mourner musician	plays music for	jazz musician
car truck	yarn truck	transports	garbage truck
yarn string	book string	tied around	neck tie
kidnapper killer	painter killer	kills	cop killer
newspaper brochure	spear brochure	about	college brochure
spear chisel	barrel chisel	chisels	wood chisel
goose vulture	fish vulture	eats	sparrow hawk
book magazine	kidnapper magazine	about	sports magazine
barrel box	goose box	contains	toolbox
painter photographer	car photographer	photographs	wedding photographer

Note. An example of familiar combination is presented which involves that relation or a very similar one.

of combinations (e.g., ''a vulture that eats dead fish'' and ''a vulture that eats dead geese'' for *fish vulture* and *goose vulture*, respectively). Note that the actual noun-noun combinations did not appear with these descriptions. Otherwise, subjects may have interpreted the combinations and then compared their interpretations to the descriptions. Thus, our task would have measured the relative plausibility of interpretations rather than the relative plausibility of thematic relations.

One description was typed on the left side of the paper and the second description was typed on the right side. Below the descriptions was a number scale from 1 to 5, with both 1 and 5 indicating ''much more plausible.'' The 1 was directly below the leftmost description and the 5 was directly below the rightmost description. A rating form was constructed by typing the pairs of descriptions in a random order, with 5 pairs on one sheet and 5 pairs on the second. Whether a description was typed on the left or right side was randomly determined with the constraint that half of the descriptions on each side of the paper corresponded to similar combinations and half to dissimilar combinations. A second rating form was created by interchanging the descriptions typed on the left and right sides. Two more rating forms were created by reversing the order of the two sheets of paper of these forms.

For the interpretation task, the 10 pairs of combinations were randomly divided into two groups of 10, subject to the constraints that only one combination from a pair appear in a group and that each group contain equal numbers (5) of similar and dissimilar combinations. Four interpretation forms were constructed by typing the combinations of each of the two groups on the left side of a piece of paper in two different random orders. A second set of four interpretation forms was created in a manner analogous to the creation of the first set.

Procedure (plausibility task). Subjects read instructions telling them that they would see descriptions of things, presented two a time (one on the left and one the right). Their task was to judge which description names a more plausible thing and to indicate how much more plausible one description is than the other. To illustrate, they were given the descriptions, ''a hatchet for pounding in nails,'' and ''a hatchet for pounding in telephone poles,'' and told that it was the experimenter's intuition that the former description was more plausible.

Subjects were instructed to rate the relative plausibility of the descriptions by circling a number from 1 to 5, with 1 indicating that the description on the left is much more plausible than the one the right and 5 indicating that the description on the right was

much more plausible than the one on the left. Subjects were told to circle 2 if they thought the description on the left was more plausible than the one on the right, but 4 if they thought the one on the right was more plausible than the one on the left. They were to circle 3 if they thought the descriptions were equally plausible.

Finally, subjects were told that some descriptions may name things which they had never seen before (like the descriptions of the hatchets) but that they should not equate unfamiliarity with plausibility. They were also told to make their decisions carefully but not to agonize over them for a long time. The task took about 10 minutes to complete. An equal number of subjects (6) filled out three of the rating forms and 5 subjects filled out the fourth rating form.

Procedure (interpretation task). A different group of subjects read instructions indicating that they would read a few noun-noun phrases which they had probably never seen in print or heard someone say (e.g., *earthquake school*). Subjects were to pretend that they had just heard each phrase in a conversation and to think of what the person meant when they said such a phrase. Subjects were to think of the most plausible meaning of the phrase and to write down this meaning to the right of each noun-noun phrase. They did not have to provide a detailed meaning but the meaning should also not be vague. The task typically took 15 minutes to complete and a subject was randomly assigned to an interpretation form.

Results

Plausibility task. Because of the counterbalancing of the left and right position of a description, the mapping of the rating scale onto the relative plausibility of a description was reversed half of the time. Thus, ratings were transformed to create a consistent mapping with increasing ratings indicating increasing relative plausibility for the descriptions involving similar constituents. The average relative plausibility rating for descriptions involving similar constituents was 3.38. This rating was reliably higher than the midpoint of the scale (3.0); $t(22) = 5.56$, $p < .001$, by a subjects analysis, but fell short of reliability in the item analysis, $t(9) = 1.73$, $p < .10$. For 8 of the 10 pairs, the description involving the similar constituents was rated as more plausible (the binomial probability is less than .055 that 8 or more pairs would show this result by chance). The two exceptions were ''a photographer who photographs cars'' and ''string that is tied around a book'' which were judged more plausible than ''a photographer who photographs painters'' and ''string tied around yarn,'' respectively. In general, these results indicate that subjects rated the descriptions involving relations between similar constituents as more plausible than those involving the corresponding relations between dissimilar constituents.

Overall, 76.5% of the ratings were at the midpoint or higher, indicating that subjects very often viewed the relation between the similar constituents as at least as plausible as the relation between the dissimilar constituents. Table 2 shows the percentage of judgments that were at the midpoint or higher for each pair.

Interpretation task. A research assistant who did not know the purpose of the study examined each interpretation and classified it into one of five categories. An interpretation was scored as a *relevant relation* if it referred to the relation on which a pair of combinations was matched. (The rater used the actual descriptions from the plausibility task to make these judgments.) For example, *goose vulture* and *fish vulture* were matched on the *eats* relation. Therefore, ''a bird that preys on the carcasses of geese'' was scored as a *relevant relation* for *goose vulture.* An interpretation was scored as *other relation* if it referred to relation between the constituents, but not to the one on which a pair of combinations was matched. For example, *newspaper brochure* and *spear brochure* were matched on the *about* relation. However, ''a brochure inside a newspaper'' was scored as an *other relation* for *newspaper brochure.* An interpretation was scored as *property* if it attributed a property of one constituent to another (e.g., ''hard-

TABLE 2

Percentage of Time That Relation between Similar Entities (Left-Hand Side) Was Judged as Least as Plausible as the Relation between Dissimilar Entities (Right-Hand Side) (Experiment 1)

A chisel for chiseling spears/a chisel for chiseling barrels	74
A vulture that eats dead geese/a vulture that eats dead fish	87
A magazine about books/a magazine about kidnappers	100
A box that contains barrels/a box that contains geese	78
A photographer who photographs painters/a photographer who photographs cars	39
A musician who plays music for dancers/a musician who plays music for mourners	100
A truck that transports cars/a truck that transports yarn string that tied around a book/string that is tied around yarn	100 39
A killer who killed a kidnapper/a killer who killed a painter	83
A brochure about newspapers/a brochure about spears	70

bound, thick magazine'' for *book magazine*) or if its referent was something with properties of both constituents (e.g., ''someone who is accomplished in both fields'' for *dancer musician*). Finally, an interpretation was scored as *other* if it did not fit into one of the other four categories. For these interpretations, subjects typically gave vague or uninformative meanings (e.g., ''a vehicle'' for *car truck*). One of the authors also scored the interpretations. Agreement between the research assistant and author was high (92%) and differences in scoring were resolved by discussion. The data from two of the 23 subjects were discarded because half of their interpretations were scored as *other*. In analyzing the data from the remaining subjects, *other* responses were discarded (11.8% for similar combinations and 8.1% for dissimilar combinations).

Figure 1 shows the percentage of relevant relation, other relation, and property interpretations for the similar and dissimilar combinations. The percentage of property interpretations was reliably higher for similar than for dissimilar combinations (65.7% versus 16.4%): $t(31) = 9.70$, $p < .001$ (subjects); $t(9) = 7.50$, $p < .001$ (items). In contrast, the percentage of relevant relation interpretations was reliably higher for dissimilar than for similar combinations (62.2% versus 22.3%): $t(31) = 6.85$, $p < .001$ (subjects); $t(9) = 4.48$, $p < .002$ (items). The proportion of *other relations* was also higher for dissimilar combinations but we had not determined their relative plausibility for similar versus dissimilar constituents. Thus, this finding is not directly relevant to the evaluation of the last resort hypothesis.

Importantly, when subjects did interpret combinations with thematic relations they most often used the relation that we had selected for the plausibility rating task, regardless of whether the combination had similar or dissimilar constituents (compare the proportions of relevant relations versus other relations in Figure 1). In principle, subjects could have used a number of relations to interpret these combinations, but when they did produce relation interpretations, they frequently used the selected relation (65% of the time for similar combinations and 74.4% of the time for dissimilar combinations). Thus, our subjects considered the selected thematic relations to be the most important and useful relations for interpreting the phrases.

Discussion

These results undermine the last resort hypothesis, which suggests that: (a) people interpret a novel combination by first attempting to find a plausible relation between the constituents, and (b) if a plausible relation does not exist, they then derive a property interpretation. In this study, relations existed between similar constituents which subjects judged as more plausible than corresponding relations between dissimilar constituents. Thus, if processing had operated in the manner suggested

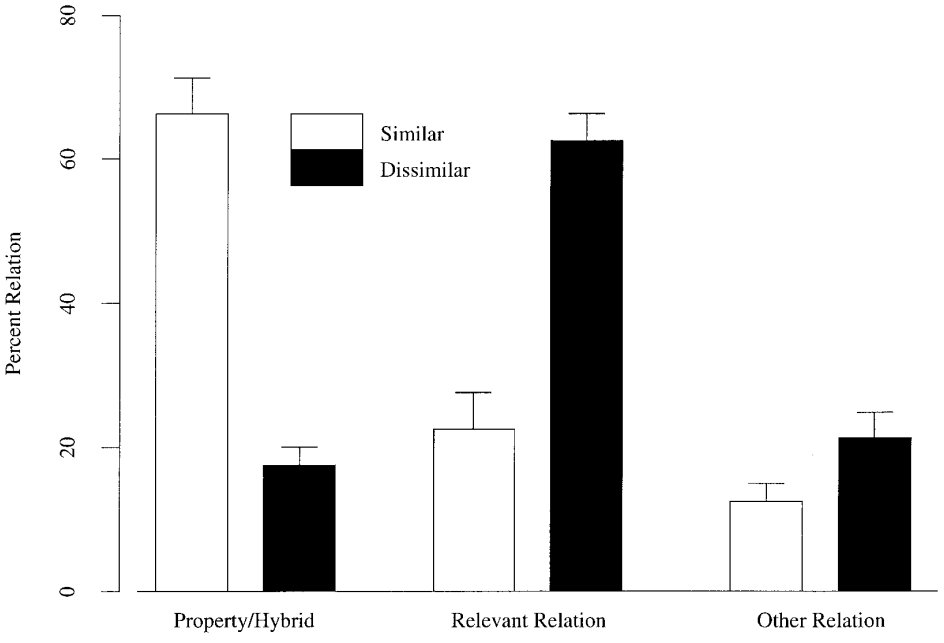

FIG. 1. Percentage of property, relevant relation, and other relation interpretations for similar and dissimilar combinations (Experiment 1).

by the last resort hypothesis, then similar combinations should have yielded *more* interpretations involving these relations.

Instead, a striking finding obtained, opposite to this prediction: a majority of interpretations of similar combinations involved properties (66%), whereas a majority of interpretations of dissimilar combinations (62%) involved the relations that subjects rated as more plausible for the similar constituents. For example, subjects believed that ''a box that contains a barrel'' was more plausible than a ''a box that contains a goose.'' However, other subjects used this relation in their interpretations of *barrel box* only 22% of the time and instead primarily gave property interpretations (78% of the time). In contrast, subjects readily used this thematic relation to interpret *goose box* (71% of the time, using property interpretations 12% of the time). As another example, subjects believed that ''a musician who plays music for dancers'' was more plausible than ''a musician who plays music for mourners.'' However, other subjects never used this relation in their interpretations of *dancer musician* and always gave property interpretations. In contrast, they readily used this relation to interpret *mourner musician* (53% of the time, using property interpretations 29% of the time).

The results also suggest that the thematic relations which we examined were not unusual or atypical of those that people use in interpreting novel combinations. In particular, they were the most common thematic relations used in the relation interpretations of both the similar and dissimilar combinations. Further, the dominant interpretations for the dissimilar combinations generally involved these relations.

EXPERIMENT 2

The last resort hypothesis implies a strict order of processing in which people are strongly biased to first construct relation interpretations and then only attempt property interpretations if there is no plausible relation which holds between the constituents. However, the results of the first study suggest that interpretation is not solely mediated by the

plausibility of relations between constituents but is also affected by the similarity of the constituents. Thus, because interpretation is affected by other characteristics of the constituents, it is more flexible than is implied by the last resort hypothesis.

Another way to explore this flexibility is to examine contextual factors that may affect the tendency to use a strategy. For example, if interpretation is more flexible than implied by the last resort hypothesis, then the use of one strategy may prime its subsequent use. In particular, if one strategy does not dominate over the other, then it should be possible to selectively prime a strategy. In contrast, the last resort hypothesis implies that it should be very difficult to selectively prime a strategy. In effect, it should not be possible to prime relation or property interpretations, because people have a bias to produce relation interpretations which is at ceiling.

Some findings suggest that the use of a specific relation to interpret a combination primes the tendency to interpret another combination using the same relation. Gerrig and Murphy (1988; Experiment 4) had subjects read a discourse context which led them to interpret a novel combination with a particular relation. For example, they read about a woman carving a trumpet out of a stale olive and then read the phrase *trumpet olive.* This interpretation then increased the likelihood of interpreting a subsequent combination (i.e., *kitten apple*) with the same relation (see also Shoben, 1993, for a similar finding).

Unlike this previous work, which primed specific relations, we attempt to differentially prime the tendency to use relation and property interpretation strategies. There were two conditions with one intended to prime property interpretations and one relation interpretations. In the *relation prime* condition, one group of subjects were first presented with combinations that had very plausible relation interpretations but no plausible property interpretations. They generated interpretations for these combinations. Then subjects interpreted test combinations which had both plausible relation and property interpretations. The *property prime* condition was identical to the relation prime condition except that a second group of subjects first interpreted combinations with very plausible property interpretations but no plausible relation interpretations. In a control condition, a third group of subjects interpreted the test combinations without previously interpreting relation or property primes.

This study also differs from the previous studies of conceptual combination in that we did not attempt to prime specific interpretations with those specific interpretations. Rather we intended to prime the general strategy of relation or property interpretation by its previous use. Such priming effects would be closely related to certain problem solving and perceptual set effects (Bugelski & Alampay, 1961; Luchins, 1942). For example, subjects shown a series of animals perceive an ambiguous reversible stimulus as a rat whereas those shown a series of people perceive the same stimulus as a man (Bugelski & Alampay, 1961). In this example, subjects were differentially primed to use a general interpretation strategy (i.e., either interpret stimuli as animals or as people) and were then more likely to apply the primed strategy to a stimulus which had both a plausible interpretation of an animal and of a person.

Method

Subjects. The subjects were 67 undergraduates from Northwestern University. They participated either as part of a course requirement or were paid.

Materials. 10 pairs of novel combinations were selected as primes. Each pair of items had the same head noun but a different modifier. By the judgments of the authors, one item of a pair (the relation prime) had a very plausible relation interpretation but no plausible property interpretation whereas the second item (the property prime) had a plausible property interpretation but no plausible relation interpretation. For example, one pair consisted of *dollar bill beggar* and *skunk beggar.* The former combination has a plausible relation interpretation (‘‘beggar who *begs for* dol-

TABLE 3

Primes and Test Items Used in Experiment 2

Property primes	Relation primes	Test items
bus truck	clothing truck	spear chisel
skunk beggar	dollar bill beggar	encyclopedia writer
motorcycle bicycle	grocery bicycle	ant vegetable
razor insult	girlfriend insult	dinosaur scientist
umbrella tree	fruit tree	snake spear
zebra tablecloth	holiday tablecloth	whale boat
sleeping pill sermon	adultery sermon	painter photographer
bullet sprinter	adidas sprinter	book magazine
roller coaster dinner	birthday dinner	mourner musician
butcher surgeon	kidney surgeon	kidnapper killer

lar bills'') but no plausible property interpretation. On the other hand, the latter combination has no plausible relation interpretation but can be plausibly interpreted by using a property of the modifier (e.g., ''beggar who *smells bad*''). Table 3 shows the pairs.

10 test combinations were selected which had both plausible relation and property interpretations, according to our intuitions. Table 3 also shows the test items. For example, *book magazine* could mean ''a magazine that reviews books'' or ''a magazine that sells books'' (relation interpretations). Or, it could mean a ''thick magazine'' (a property interpretation). *Whale boat* could mean ''a boat used for whale watching'' or ''a boat for hunting whales'' (relation interpretations) or it could mean ''a very big boat'' (property interpretation). As a final example, *spear chisel* could mean ''a chisel used to make spears'' or ''a tool used for sharpening the point of a spear'' (relation interpretations). Alternatively, it could mean ''a chisel that is long'' or ''a sharp chisel'' (property interpretations).[1]

Procedure. Subjects read instructions indicating that they would see phrases presented on a computer screen, one at a time. They were to read each phrase and to think of what someone would plausibly mean by it. Some of the phrases might have obvious meanings and other phrases might have meanings which were not so obvious. Subjects were told that when a meaning came to mind (but not before then) they should press the space bar on the computer keyboard. They were to try to think of the meaning as quickly as they could. After they had pressed the space bar, they would be prompted to type in the meaning of the phrase. The instructions stated that the first part of the task (determining the meaning) would be timed but not the part involving the typing in of that meaning. After subjects typed in the meaning they had thought of, they were to press the tab key. The computer would then instruct them to gently rest an index finger on the space bar and get ready for the next trial.

After reading the instructions, subjects began the task by pressing a key on the computer keyboard. Subjects in the relation prime condition first interpreted the relation primes (the 10 combinations which had plausible relation interpretations but not plausible property interpretations). Then they interpreted the 10 test combinations. Subjects in the property prime condition first interpreted the property primes (the 10 combinations which had plausible property interpretations but not plausible relation interpretations). They then interpreted the same 10 test combinations. Each subject in the prime conditions saw the primes in a different random order. They also were un-

[1] As illustrated by these examples, a combination might have multiple property or relation interpretations. For the purpose of this study, it is only relevant that the combinations have at least one plausible interpretation of each kind.

aware that the stimuli were divided into primes and test items. Subjects in the neutral condition only interpreted the 10 test combinations. Every subject in the three conditions saw the test items in a different random order.

Subjects proceeded through the task at their own pace: they first saw a combination, then pressed the space bar when they had thought of a meaning, then typed in the meaning, and then pressed the tab key to go onto the next trial. The computer recorded the time from the onset of a combination to when the subject pressed the space bar. The task typically took 15 minutes to complete. Subjects were randomly assigned to the relation priming, property priming, and neutral conditions, with approximately equal numbers of subjects participating in each condition.

Results

Scoring. The interpretations of the test items were gathered together and sorted by test item. One of the authors examined each interpretation and classified it into one of three categories: relation, property, or other. He was unaware of whether the interpretation was produced in the relation, property, or neutral condition. An interpretation was scored as *relation* if it explicitly referred to a relation between the constituents (e.g., ''someone who plays an instrument for those in mourning'' for *mourner musician*) or strongly implied a relation (e.g., ''a musician used at funerals'' implies that the musician plays music for mourners). An interpretation was scored as *property* if it attributed a property of one constituent to another (e.g., ''a sad musician'' or ''a musician who writes or sings sad songs'' are both based on the salient property of grief that characterizes mourners) or if it involved combining properties of both constituents (e.g., interpreting *book magazine* as ''a piece of literature that has binding and pages like a book, with the pictures of a magazine''). Finally, an interpretation was scored as *other* if it failed to relate the constituents to each other in some manner (e.g., ''used for meaning of words'' for *encyclopedia writer*) or if it could not be unambiguously scored as a relation or property (e.g., ''a boring person'' for *encyclopedia writer*). In the latter case, the person could be boring because the person writes encyclopedias (a relation interpretation) or because the person talks about lots of tedious facts (a property interpretation). On rare occasions (1.1% of the time), a subject typed two interpretations (e.g., ''either a big boat or a boat used for whaling'' for *whale boat*). In these cases, both interpretations were scored.

TABLE 4

Proportion of Relation and Property Interpretations as a Function of Condition (Experiment 2)

Condition	Relation interpretations	Property interpretations
Relation prime	.628	.372
Neutral	.547	.453
Property prime	.323	.677

As a check for accuracy, the other author scored a randomly selected 15% of the interpretations (he was also unaware of whether an interpretation was produced in the relation, property, or neutral condition). The authors agreed on 97.9% of their classifications.

Results. The data from 3 of the 67 subjects were discarded because at least half of their interpretations were scored as *other*. These subjects were from the neutral condition. For the remaining subjects, interpretations scored as *other* were discarded from analyses (they accounted for a small percentage of responses—less than 4% in each condition). Table 4 shows the proportion of property and relation interpretations for each condition. As predicted, the relation condition had the highest percentage of relation interpretations (62.8%) whereas the property condition had the lowest percentage (32.3%) and thus the highest percentage of property interpretations. In the neutral condition, the percentage of relation interpretations was in between that of the relation and property conditions (54.7%). Using percentage of relation interpretations as the dependent measure, a one-way ANOVA revealed a highly reliable effect for type of condition: $F(2, 62) = 11.71$, MSE = 5023.8,

$p < .001$ (subjects); $F(2, 18) = 52.20$, MSE $= .2405$, $p < .001$ (items). The percentage of relation interpretations was reliably higher for the test combinations when they were interpreted in the relation condition than in the neutral condition, by an item analysis: $t(9) = 2.64$, $p < .03$. This difference did not reach reliability in the subject analysis; $t(43) = 1.35$, $p < .19$. On the other hand, for 17 of 22 subjects in the relation condition the proportion of relation interpretations was greater than the average proportion of relation interpretations in the neutral condition ($p < .04$, by a binomial test). Finally, the percentage of relation interpretations was reliably higher in the neutral condition than in the property condition: $t(40) = 3.14$, $p < .004$ (subjects); $t(9) = 9.72$, $p < .001$ (items).

Comparing the percentage of relation interpretations across test items, all 10 items had lower percentages of relation responding in the property condition compared to their corresponding percentages in both the neutral and relation conditions. 8 of the 10 items had higher percentages of relation responding in the relation condition compared to the neutral condition.

It was also clear that the relation and property primes were interpreted in terms of relations and properties, respectively (thus confirming our intuitions which motivated their selection). By one author's judgment, 90.9% of the relation primes were interpreted by using relations. For example, *dollar bill beggar* was interpreted by 20 of the 22 subjects in the relation condition by using a relation (i.e., beggar that begs for dollar bills). In contrast, 92.1% of the property primes were interpreted by using properties. For example, *skunk beggar* was interpreted by 18 of 19 subjects by using a property (i.e., a foul smelling beggar). When *other* responses are eliminated, these percentages were 92.7% and 94.2%, respectively.

Discussion

The results provide evidence for selective priming of interpretation strategies in conceptual combination. In the neutral (no prime) condition, 45.3% of responses were property interpretations, compared to 67.7% property responding when subjects had been primed to use such a strategy. A smaller but statistically reliable priming effect was also obtained for relation responding. Compared to the neutral condition (which had 54.7% relation responding), there was 62.8% relation responding when subjects had been primed to use such a strategy.

For several reasons, this pattern of findings suggests that relation and property interpretation strategies are on more equal footing than suggested by the last resort hypothesis. First, the last resort hypothesis claims that people first attempt to interpret a combination by determining a plausible relation between constituents. However, this claim does not predict that relation interpretations should *increase* in the relation condition relative to the neutral condition. Presumably, prior to interpreting the test items, subjects in the relation condition were interpreting the primes in the way that they normally interpret combinations. Thus, it is unclear why applying the usual strategy for combining concepts should increase the tendency to apply the usual strategy. Second, the hypothesis claims that property interpretation is a strategy of last resort. However, our simple priming manipulation succeeded in greatly increasing the tendency to produce property interpretations to the point where this strategy was the dominant one applied to combining concepts in the property condition.

A striking finding was the extent to which subjects in the two priming conditions used different strategies and arrived at very different interpretations for the *same* combinations. In the relation condition, subjects produced relation interpretations almost 63% of the time, but in the property condition subjects produced property interpretations almost 68% of the time. For example, more than half the time subjects in the relation condition interpreted *ant vegetable* as ''a vegetable eaten by ants,'' *kidnapper killer* as ''a killer who kills kidnappers,'' and *snake spear* as ''a spear used to kill snakes.'' In contrast, subjects in the property condition interpreted these same combinations more than half the time as ''a

very small vegetable,'' ''a killer who is also a kidnapper,'' and ''a curvy spear,'' respectively. Even when there was an especially strong tendency in one condition to interpret a combination in one way, priming in the other condition counteracted this tendency. For example, virtually every subject in the property condition interpreted *spear chisel* in terms of properties (e.g., ''long, pointy chisel'') but almost a third of the subjects in the relation condition interpreted this combination as a ''chisel for making spears.'' Likewise, every subject in the relation condition interpreted *dinosaur scientist* as ''a scientist who studies dinosaurs'' but almost one-third of the subjects in the property condition interpreted this combination as ''a very old scientist.''

A possible concern about the findings is that they resulted from a demand characteristic. In particular, because the primes could only plausibly be interpreted with one strategy, subjects may have thought that the experimenter intended them to use that strategy for subsequent test combinations as well. For several reasons, we believe that this alternative account is not likely. First, a demand characteristic is more of concern in the case in which the experimenter explicitly provides the interpretations of the primes in contrast to the present study in which subjects generated their own interpretations of the primes. In the latter case, subjects would be less certain about which kinds of interpretations that the ''experimenter wanted'' because they were producing the interpretations themselves. Second, the demand characteristic explanation presupposes that subjects can clearly distinguish property and relation interpretations. However, recent studies suggest that subjects blur this distinction (Wisniewski & Bassok, 1997). In one study, subjects listed commonalities and differences between objects. To illustrate the task, they were shown properties of objects that were identified as commonalities and differences. Nevertheless, subjects also listed thematic relations as commonalities and differences between objects. In another study, subjects judged the thematic relatedness between objects and provided explanations of their ratings. The task was illustrated by presenting several examples of thematic relations between pairs of objects. Nevertheless, subjects listed properties in justifying their ratings.

It also appears that the demand characteristic explanation and the last resort strategy cannot simultaneously apply to the findings. If subjects are already strongly predisposed to interpret combinations with thematic relations (as implied by the last resort hypothesis) then subjects in the relation prime condition should not be influenced by a demand characteristic to produce even more of these interpretations. Yet, there were more relation interpretations in this condition than in the neutral condition. Also, in the property prime condition, one might expect that the demand characteristic would cause subjects to suppress the relation interpretations that they were strongly predisposed to generate.[2] Thus, the last resort hypothesis predicts that subjects in the property prime condition should be slower to generate property interpretations than subjects in the relation prime condition should be to generate relation interpretations. To examine this possibility, we computed the median RT for property interpretations for subjects in the property prime condition (discarding RTs for relation interpretations) and the median RT for relation interpretations for subjects in the relation prime condition (discarding RTs for property interpretations). The average median RT for property interpretations in the property prime condition was actually slightly faster than the RT for relation interpretations in the relation prime condition—2384 msecs versus 2461 msecs, respectively (this difference was not reliable, $t < 1$).

In sum, the findings provide more support for the view that relation and property interpretation strategies are both cognitively useful and important ways of combining concepts. Thus, there is not a strict order of processing with a single, strongly dominant tendency towards relation interpretation followed by property interpretation as a last resort.

[2] We thank an anonymous reviewer for bringing this possibility to our attention.

EXPERIMENT 3

In the previous studies, we evaluated the first part of the last resort hypothesis by examining tendencies to generate relation and property interpretations for experimenter-constructed novel combinations. In this study, we evaluate the second part of the last resort hypothesis by examining the meanings of combinations that people produce. Importantly, Downing's (1977) hypothesis that property interpretation is a last resort strategy was based on a comparison between production and interpretation of novel combinations. In one study, she analyzed the meanings of novel combinations which appeared in three sources of text (written descriptions by subjects of a brief film sequence in which a boy engages in various play and work activities, the general news section of an issue of the San Francisco Chronicle, and two novels by Richard Brautigan). Downing then divided the combinations into five categories on the basis of the referent of the head noun (human, animal, plant, natural object, synthetic object) and within each category ranked the frequency of property meanings and various types of relation meanings (she did not report the actual frequencies of meanings). Downing found that property meanings were absent in three of the categories (plants, natural objects, synthetic objects), ranked last in humans, and ranked second in animals. In contrast, when subjects were given combinations in these categories to interpret, property meanings were absent only in humans, ranked first or second in three categories (plants, animals, synthetic objects), and ranked third in natural objects.[3] She concluded that property meanings were absent in production but favored in interpretation. Based on this asymmetry, she then suggested that property interpretation is a last resort strategy used only when there is no other useful relationship that seemed plausible (Downing, 1977; page 830).

We interpret Downing's view as implying that people do not or only rarely produce combinations with property meanings and that she found property interpretation in combinations presented to subjects because these combinations lacked plausible thematic relations between constituents. Thus, in trying to make sense of the meaning of these combinations, subjects relied on an unusual fallback strategy (i.e., property interpretation). (Downing does not report the combinations that were used or whether she examined them to see if they had plausible relation interpretations.)

In the present study, we revisited the assumption that people have a strong bias to produce combinations whose meanings involve a thematic relationship between their constituents. We followed up Downing's study for several reasons. First, Downing's analysis involved novel combinations coined on one occasion by one speaker whereas we examined the meanings of conventional names that were used frequently enough to have been explicitly encoded in reference manuals. Because they are part of a language, they have been used by multiple speakers on multiple occasions. Thus, we can be more certain that these conventional combinations have a fair degree of functional utility in communication and cognition. As a result, finding property meanings among these combinations would constitute stronger evidence for their importance. Second, the description of Downing's analysis is ambiguous—she does not report the size of her sample or the absolute frequencies of property versus relation meanings. Third, her sample may not have been very representative. One of her sources was apparently descriptions of only a small number of activities performed by a boy. She also relied on several books written by Richard Brautigan who was an unusual person and writer (see Abbot, 1989, for a biography).

In our study we examined a very large sample of familiar noun-noun combinations (more than a thousand). They were selected from a

[3] In Downing's analysis, she refers to property interpretations as comparison relationships. She also had a category for combinations with many properties of both constituents which she refers to as half/half. In her study of production, this category was ranked third in humans, tied for first in animals, and last in synthetic objects. It was absent in plants and natural objects. In interpretation, it was tied for first in humans and plants and was absent in the other categories.

number of sources and spanned a wide range of noun referents, especially animals, plants, and artifacts. The primary goal of the analysis was to obtain a reasonably accurate estimate of the relative frequency of property and relation meanings in the English language. We used these estimates to further assess the last resort hypothesis. A secondary goal was to obtain some idea of the range of the kinds of property and relation meanings that characterize familiar combinations in English. Whereas there have been some very influential analyses of meanings of familiar combinations (e.g., Levi, 1978) we are not aware of analyses which are derived from systematic sampling of a large number of such phrases.

Method

Materials. We examined a large set of combinations that named animals, plants, and artifacts. We also selected a sample of combinations from a dictionary in order to look at a wider range of nouns. Animal and plant terms came from the book *Reader's Digest North American Wildlife* (1992). It is an illustrated guide to 2,000 plants and animals, describing 8 broad categories of wildlife: mammals, birds, reptiles and amphibians, fish, invertebrates, trees and shrubs, wildflowers, nonflowering plants, and mushrooms. For each animal and plant, the book lists its common name, shows a picture of it, and describes its identifying features, habitat, and other important information. We arbitrarily selected for analysis three of the four animal categories (discarding invertebrates) and three of the four plant categories (discarding nonflowering plants). Artifact terms came from two sources: a catalog called the *Office Stores Program Catalog* (1995) which provides a large range of office supplies to Northwestern University faculty and staff. The catalog is divided into many categories of supplies and provides pictures and descriptions of each artifact. From this source, we arbitrarily selected the categories books, clips and clamps, desk accessories, fans and heaters, furniture, janitorial supplies and equipment, knives, lunchroom supplies, scissors, staplers, and waste receptacles. Another source for artifacts was *The Facts on File Visual Dictionary* (1986). As noted in the preface, this dictionary ''lists terms and notions which designate the many elements of everyday life in an industrial, post-industrial, and developing society,'' organized into a wide range of categories. From this source, we selected two large artifact categories (clothing and personal items) which did not overlap with those chosen from the office supply catalog. Finally, we collected a random sample of combinations from the *Random House Webster's College Dictionary* (1995).

Procedure. One of the authors analyzed the animal, plant, and artifact combinations. For the animals and plants, he read through the pages of *Reader's Digest North American Wildlife* (1992) that described the six categories above and examined any noun-noun combinations. Because the category of wildflowers was relatively long (175 pages), he only read every third page for this category. The author picked out a combination and judged whether the corresponding entity was named for a functional/thematic relation that it shared with the referent of the modifier (a relation interpretation) or whether the entity was named for some type of resemblance it had to the modifier (a property interpretation). On a few occasions, there was more than one modifier noun that was part of the combination. For example, in the case of the *California sage thrasher,* there is a functional relation between *thrasher* and *California* (location) and one between *thrasher* and *sage* (location). In these few cases, the author made two judgments. There were also a few occasions in which the head noun that would typically refer to an entity was not present. Instead the entity was given the name of another entity which it resembled in some way. For example, one fungus is called *pig's ear* rather than *pig's ear fungus.* This fungus is perceptually similar to a pig's ear (in color, shape, and size). These combinations were scored as property interpretations as they involve the mapping of properties from one entity (e.g., pig's ear) to a second entity (e.g., fungus). In making his judgments, the author relied on the descrip-

tions of the entity and the color picture. He also consulted other sources which sometimes included more information for making a judgment. These sources included a book on mushrooms and other fungi (Dermek, 1989), a Peterson field guide to Pacific Coast fishes of North America (Eschmeyer, Herald, & Hammann, 1993), the *Reader's Digest Magic and Medicine of Plants* (1986), and the *Random House Webster's College Dictionary* (1995). On a few occasions, the author was unable to determine if a meaning referred to a property or relation interpretation.

The author carried out a similar analysis in regard to the artifacts. For this analysis, he also consulted the *Random House Webster's College Dictionary* (1995) when he was unsure of a judgment. In addition, the combinations found in the *Office Stores Program Catalog* (1995) often contained many noun modifiers. To make the analysis more manageable, the author arbitrarily made a single judgment per combination, based on the head noun and the immediately preceding modifier noun.

In looking over the combinations, the author was also able to derive a set of categories that tended to characterize most of the meanings of the combinations. For example, a very common functional relationship between the constituents of a combination was location. Many combinations referred to entities which were located in or on the referent of the modifier (e.g., mountain sheep, wall clock). As another example, some combinations referred to entities whose overall shape resembled that of the referent of the modifier (e.g., cauliflower fungi, peninsula table). The Appendix describe these categories and provides examples of them from the analysis.

The other author examined noun-noun combinations in the *Random House Webster's College Dictionary* (1995). He opened the dictionary to a randomly selected odd-numbered page and examined the combinations on that page and the following page. He repeated this process until he obtained 200 noun-noun combinations. Based on the interpretation provided in the entry for the combination, he classified its meaning as property or relation (on a few occasions he was unable to determine whether the meaning referred to a property or relation). The sample contained combinations with a wider range of different types of nouns than the artifact, plant, and animal samples and thus appeared to yield a wider range of specific types of relation and property meanings. As a result, the author did not attempt to differentiate the relation and property meanings into more fine-grained categories.

Inter-rater reliability. One author classified a randomly selected subset of 15% of the combinations from each of the three categories animal, plant, and artifact, which had been classified by the other author. There was high agreement on whether a meaning referred to a relation interpretation or property interpretation. The authors agreed on 94.8%, 91.5%, and 93.3% of their classifications for the plant, animal, and artifact combinations, respectively. Overall, there was 93.2% agreement for these combinations. In terms of the more specific property and relation categories (see the Appendix), the authors agreed on 88.7%, 94.4%, and 83.3% of their classifications for the plant, animal, and artifact combinations, respectively.

One author also classified a randomly selected 15% of the combinations from the dictionary sample which had been classified by the other author. Again, there was high agreement between classifications (97.0%).

Results

Meanings that we were unable to classify were discarded from all analyses (5.6%, 8.2%, 2.9%, and 6.8% of the meanings for the animal, plant, artifact, and dictionary samples, respectively). Table 5 shows the proportions of relation and property interpretations for combinations from the various categories. Overall, 70.9% of the combinations had relation meanings and 29.1% had property meanings. Although relation meanings were dominant in production, property meanings were not rare and characterized almost thirty percent of the combinations that people produced. Almost half of the plant combinations had property meanings and there were several sub-

TABLE 5

Proportions of Relation and Property Meanings for Combinations Samples from Animal, Plant, and Artifact Categories, and the Dictionary (Experiment 3)

Category	*n*	Relations	Properties
Animals	298	.61	.39
Mammals	42	.67	.33
Birds	108	.81	.19
Fish	63	.27	.73
Reptiles	85	.60	.40
Plants	311	.54	.46
Wildflowers	186	.55	.45
Trees & shrubs	79	.61	.39
Mushrooms	46	.35	.65
Artifacts	598	.86	.14
Personal items	73	.88	.12
Clothing	104	.62	.39
Furniture	178	.90	.10
Books	19	1.00	0
Clocks	5	.80	.20
Desk accessories	59	.95	.05
Fans & heaters	6	.83	.17
Janitorial supplies	64	.95	.05
Knives	7	.86	.14
Lunchroom supplies	50	.94	.06
Scissors	5	.80	.20
Staplers	9	.78	.22
Clips	8	.63	.38
Waste receptacles	12	.92	.08
Dictionary	194	.67	.33

categories in which property meanings dominated—65.2% of the mushroom combinations and 73% of the fish combinations had property meanings. One third of the combinations in the random sample from the dictionary had property meanings.

There was a fairly large difference between the proportions of property meanings for artifacts versus natural kinds. Only 14.2% of the artifact combinations had property meanings compared to 42.5% for animals and plants. However, one broad category of artifacts (clothing) had a higher than overall average of property meanings (38.5%).

Table 6 presents the proportions of different types of property and relation meanings for the animal, plant, and artifact combinations. The dominant relation was location—animals and plants were often named using modifiers that indicated the places in which they were typically found (e.g., mountain sheep, harbor porpoise, prairie falcon, ground ivy, desert lily). In contrast, artifacts were primarily named for interactions that they had with other artifacts—e.g., about 30% of the combinations involved nominalizations with the modifier and head nouns referring to objects that play different roles in the verb which had been nominalized (e.g., glass cleaner, diving suit, pencil holder, shaving brush, teaser comb, neck tie). Another common relation among the artifact combinations was the contains/holds relation, probably because our sample contained office supplies and furniture which often functioned to organize artifacts and to maximize work space.

In the property meanings, the referent of the combination was often related to the modifier in terms of perceptual properties such as appearance, overall shape, color, and so on. For example, in plants, a distinctive part of the plant often resembled some common entity and was thus named using that entity as the modifier. The Appendix provides many examples of these resemblance relationships.

Discussion

An important claim that led to one version of the last resort hypothesis—that property meanings rarely occur in production—was not supported in our analysis of a large number of combinations taken from the domains of animals, plants, and artifacts. The results show that among combinations used to name animals and plants, over 40% involved property meanings. Although the relative proportion of property meanings was considerably less among combinations used to name artifacts (about 14%), it was not an extremely small proportion. And, for the broad category of clothing, nearly 40% of the combinations involved property meanings. Some notable examples of artifacts that had property meanings included bikini briefs (overall shape), rake comb (function, overall shape), accordion pleat (part-part), leg-of-mutton collar (overall shape), pen knife (function, overall shape), blanket sleeper (function), saddle bag (overall shape), plier stapler

TABLE 6

Proportions of Different Types of Relation and Property Meanings for Animal, Plant, and Artifact Combinations (Experiment 3)

Category	Properties		Relations	Prop
Animals	object-part	.29	location	.63
	color/pattern	.18	named after	.10
	part-part	.14	part	.08
	function/behavior	.13	eats	.07
	overall shape	.10	when active	.05
	size	.06	object-nom	.04
	sound	.03	other	.02
	appearance	.03	human resource	<.01
	hybrid	.02		
	taste	<.01		
	smell	<.01		
	other	<.01		
Plants	object-part	.49	location	.56
	color/pattern	.13	human resource	.13
	overall shape	.09	part	.11
	appearance	.08	named after	.07
	part-part	.07	when active	.04
	smell	.03	eats	.04
	function/behavior	.03	object-nom.	.02
	size	.03	causes	.02
	hybrid	.01	attract	.01
	taste	.01	other	<.01
	small	.01		
Artifacts	overall shape	.32	holds/contains	.19
	hybrids	.20	object-nom	.16
	part-part	.17	location	.15
	function/behavior	.14	part	.14
	resemblance	.14	nom-instrument	.11
	object-part	.02	made of	.08
	size	.01	named after	.05
			collection	.04
			object-instrument	.04
			state-instrument	.02
			when active	.02
			connects to	<.01
			depicts	<.01
			other	<.01

(function), bulldog clip (function), owl clip (part-part), scooter stool (function, part-part), and tower cart (size). (See Appendix for definitions of these classifications.)

GENERAL DISCUSSION

Taken together, the results of these studies suggest that property interpretation is an important strategy in understanding noun-noun combinations and not one that is employed as a last resort. In Experiment 1, one group of subjects judged thematic relations between pairs of similar entities to be more plausible than the corresponding relations between pairs of dissimilar entities. However, a different group of subjects was much more likely to give property interpretations to the combinations involving the similar entities. At the same time, they interpreted a majority of the dissimilar combinations with the thematic re-

lations that they had judged to more plausibly apply to the constituents of the similar combinations. This result is incompatible with the view that subjects only produce property interpretations of combinations when plausible thematic or functional relations do not exist between their constituents.

The results of Experiment 1 go beyond those previous studies that have also showed that highly similar combinations are interpreted with properties (Markman & Wisniewski, 1997; Wisniewski, 1996; Wisniewski & Markman, 1993). In those studies, researchers created novel combinations by pairing highly similar nouns without regard to whether the constituents could plausibly participate together in a thematic relation. As our analysis described in the Introduction suggests, this procedure very often creates combinations which cannot be plausibly linked by a thematic relation. Thus, the previous finding that subjects generate property interpretations for highly similar combinations is consistent with the last resort strategy. However, our findings show that high similarity between constituents facilities property interpretation and does not just prevent relation interpretations (as implied by the previous studies).

In Experiment 2 we examined the effects of prior generation of property and relation interpretations on subsequent interpretations. Subjects interpreted novel combinations which had either highly plausible property or relation interpretations (but not both). As a result, subjects were induced to primarily use either property or relation interpretation strategies. Subjects then tended to interpret combinations which had both property and relation interpretations using the primed strategy. For example, after generating relation interpretations such as ''beggar who asks for dollar bills'' for *dollar bill beggar* and ''chisel for making pottery'' for *pottery chisel,* subjects tended to interpret *book magazine* using a relation as in ''a magazine that discusses books.'' However, after producing property interpretations such as ''begger who smells bad'' for *skunk beggar* and ''very sharp chisel'' for *knife chisel,* subjects tended to interpret *book magazine* using a property as in ''a thick book.''

Importantly, the finding that prior use of the relation interpretation strategy increased its subsequent use is incompatible with the view that people's usual strategy is to attempt a relation interpretation. This view would not predict that previously engaging in the usual strategy will subsequently increase the use of that strategy. The results also suggest that it is relatively easy to increase the tendency to generate property interpretations. Such a tendency seems at odds with the view that property interpretation is a strategy of last resort.

Finally, in Experiment 3, we examined a large sample of combinations which have become part of the English language. In contrast to our previous studies, which constrained subjects to generate interpretations of experimenter-provided combinations, this study investigated combinations which had been produced by speakers. We found that while a majority of the combinations had relation meanings, property meanings were not rare and characterized almost 30% of our sample. This finding does not agree with Downing's (1977) or Shoben and Gagne's (1997) claim that property meanings are not present in combinations that people produce. It was the apparent lack of such meanings in production that led researchers to suggest that they are only generated during interpretation as a last resort.

In sum, our findings suggest that there is not a general, strong bias to interpret combinations by using any particular strategy. We demonstrated that at least some factors—structural properties of constituents (i.e., similarity of their representations) and context (i.e., prior use of a process)—can affect whether property or relation interpretations dominate. In Experiment 1, property interpretations characterized the vast majority of interpretations for similar combinations (even though they had plausible relation interpretations) and relation interpretations characterized the great majority of interpretations of dissimilar combinations. In Experiment 2, the strategy which dominated interpretation de-

pended on whether it had been previously used. Thus, because it is possible to selectively manipulate the dominance of one strategy over the other, our findings suggest that both strategies are important in combining concepts.

Below we describe several other reasons to doubt that people interpret combinations by first attempting to link the constituents by a thematic relation, using property interpretation as a last resort strategy. First, we note that the interpretation of novel combinations is a special case of category formation and that category formation depends both on thematic relations and properties. Second, we suggest that using nouns to refer to their properties plays an important function in communication. In concluding, we discuss the implications of our results for models of conceptual combination.

Categorization

Both relations and properties are important in category representations. On the one hand, these representations capture relations between entities. For example, events represented by scripts typically capture relations in which objects play different functional roles (Bower, Black, & Turner, 1979; Schank & Abelson, 1977). To illustrate, a restaurant script might include the serving relation, which specifies that a waiter or waitress serves food to a customer. In this example, the waiter plays the agent role of the serving relation (i.e., the server), the food the object role (i.e., what is served) and the customer the recipient role (i.e., the person served). At the same time, category representations are also based on overlapping properties or resemblance relationships. For example, there may be a fast food restaurant script and a formal restaurant script which are stored in the same memory location because they share properties which make them restaurant scripts (and not because they have a functional relationship between them). As another example, object categories like *car* represent relations in which entities play different functional roles such as ''a *driver* operates the *car* by pushing the *foot* down on the *pedal.*'' At the same time, *car* is organized together in memory with *truck*, *train*, *plane*, and so on, as subcategories of *vehicle.* This organization is based on overlapping properties rather than on the basis of different functional relationships between these categories.

Thus, both thematic relations between entities and overlapping properties shared by entities are important in forming categories. Of course, conceptual combination is a special case of category formation in which two existing category representations are combined to create a new subcategory of the head noun (Murphy, 1988). For example, *seafood* and *sausage* can be combined to create a new subcategory of *sausage* called *seafood sausage* which differs from other kinds of sausage in a variety of ways, most notably in being made out of seafood rather than some type of meat. Given that conceptual combination is a type of category formation, it seems reasonable that relation and property interpretation strategies would both be important in combining concepts.

The Use of Nouns to Refer to Properties

The present findings suggest that construing a noun to refer to its properties is an important strategy used in conceptual combination. This kind of construal also forms the basis of a major theory of metaphor understanding which assumes that nouns have dual reference (Glucksberg & Keysar, 1990; Glucksberg, Manfredi, & McGlone, 1997). When used literally, a noun refers to its typical category, but when used metaphorically it refers to properties which it exemplifies. For example, in the literal statement ''my job at the jail,'' *jail* refers to a jail, but in the nominal metaphor ''my job is a jail'' it refers to prototypical properties of a jail (e.g., confinement).

It is interesting to speculate about *why* people use nouns to refer to properties. Typically, adjectives play this role in language. However, sometimes a property associated with a noun rather than an adjective may better capture the sense of a property of an entity that needs to be named. For example, a *zebra mus-*

sel is a mollusk that has recently populated the Great Lakes. The outside of its shell is characterized by a pattern of approximately parallel, curved, black stripes on a whitish background. Although this mollusk could be named the *striped mollusk*, its stripes have some striking similarities to those of a zebra's. Thus, *zebra mussel* may be a more accurate name for this mollusk than *striped mussel.* It is easy to find other examples in which a property of a noun's typical referent more specifically characterizes a situation than an adjective referring to the more general property. *Skunk cabbage* could have been called *smelly cabbage* but the smell of this plant is strikingly similar to the smell of a skunk (as those who have stepped on one know). Also, the *swordfish* could have been called the *pointy-nosed fish* but the overall appearance of a *sword* better captures the appearance of the nose of a swordfish.

Nouns may also refer to properties when there is no adjective that appropriately characterizes a situation. (See Glucksberg et al., 1997, for a related discussion.) For example, map turtles have patterns on their backs which are map-like in appearance. There does not seem to be an adjective counterpart which describes this situation. As another example, alligators are the largest reptiles, with big jaws, long tails, and knobby ridges on their skin. In like fashion, *alligator snapping turtles* are the largest freshwater turtles, with big jaws, long tails, and knobby ridges on their shells. Thus, *alligator* is an appropriate term, that succinctly refers to this complex of properties. Again, there is no corresponding adjective which refers to this set of properties, and it seem cumbersome to capture this information with a series of adjectives.

Finally, nouns may be used instead of adjectives because they ground a situation directly in experiences to which we can readily relate. For example, one could refer to a particularly hot desert as an *extremely hot desert.* Alternatively, one might use the phrase *oven desert.* However, the hotness of ovens has been directly experienced by anyone who has opened an oven door and peered in while it is cooking something, or who has been burnt taking a roast out of the oven. The hotness of an oven may conjure up a particularly vivid and visceral experience as opposed to the more abstract phrase *extremely hot.* Note that this reason for using nouns to refer to properties seems different from first one given above. For example, whereas *smelly* fails to capture the sense of bad smell which *skunk* captures, *extremely hot* and *oven* seem to capture a similar degree and sense of hotness.

Models of Conceptual Combination

The present findings have several implications for models of conceptual combination. First, because they suggest that property interpretation is an important and not uncommon phenomenon, a complete model must account for such interpretations and the factors which are likely to lead to their occurrence. However, a number of approaches currently provide single-process accounts of relation interpretation (Cohen & Murphy, 1984; Coolen et al., 1991; Gagne & Shoben, 1997; Murphy, 1988; Shoben & Gagne, 1997; but see Wisniewski, 1997, for a dual-process proposal). The present findings complicate the modelling of how concepts combine because they suggest that there are two significant ways that people use to combine concepts rather than one.

Second, the results from the first and second experiments provide some constraints on a processing account of how nouns combine. In particular, the findings rule out a serial processing model in which people first attempt to interpret a combination with a thematic relation and only then attempt a property interpretation if there is no plausible thematic relation. Experiment 1 provides especially strong evidence against this strict order of processing. In particular, the fact that similar combinations had very plausible relation interpretations yet were frequently interpreted with properties argues against this type of model. Future work needs to develop a processing account in which property interpretation plays a much more significant role in how nouns are combined than has previously been thought.

APPENDIX

Categories of Relation and Property Meanings Used to Characterize Familiar Combinations (Experiment 3)

Relation Categories

Location. For an animal or plant, the modifier indicated its typical habitat (e.g., mountain lion, harbor porpoise, common ground dove, orchard oriole, cavefish, pine mushroom, water celery). For an artifact, the modifier indicated where it was typically used or found (e.g., wall clock, room cabinets, desk stapler, lapel microphone).

Human resource. The modifier and head noun concepts were linked by a relation that described a resource which was important to human beings. Some examples included: *fly agaric,* a fungus from which a poison used to be prepared for killing flies; *birthwort,* a plant whose roots were thought to ease the pain of childbirth; *sugar maple,* a tree whose sap is boiled to produce sugar; and, *northern fur seal,* which is valued for its lustrous pelt.

Part. The modifier named a part of the head noun referent (e.g., fin whale, bonefish, pore fungi, bur cucumber, shagbark hickory).

Named after. The modifier named a person, group of people, or company for which the referent was named after. In the case of an animal or plant, the person was typically its discoverer (e.g., *MacGillivray's warbler* was named for an orinthologist and the *Douglas fir* was named after a botanist). On the other hand, artifacts were sometimes named after their manufacturers (e.g., *lipton tea, Bausch & Lomb magnifiers*) or for people for whom they were intended to be used (e.g., *executive chair, banker's clip*).

Made of. Some combinations referred to artifacts made out of a substance named by the modifier (e.g., *quartz clock, wire sorter, latex gloves, wood furniture, steel scissors*).

When active. For animals and plants, the modifier sometimes indicated when the referent was typically active or when it first appeared during the year (e.g., *winter mushroom, springcress, evening grosbeak, night snake*). For artifacts, the modifier sometimes indicated when the referet was to be used (e.g., nightgown, raincoat).

Eats. The modifier named something that either ate the referent or was eaten by the referent (e.g., *northern grasshopper mouse, acorn woodpecker, gopher snake, cow parsnip, beargrass, butterfly weed*).

Nominalizations (nom-object, instrument-nom). In some combinations, the head noun was derived from a verb with the modifier typically functioning as one of the cases associated with the verb. Usually, the modifier fit the object case (e.g., *sapsucker, flycatcher, staple remover, stationary holder*) though on occasion it fit the instrument case (e.g., *vacuum cleaner*). Sometimes, the modifier was nominalized and the head noun functioned as one of the cases (e.g., in *drafting stool* and *storage hutch, stool* and *hutch* function as the instrument case in the verbs *draft* and *store,* respectively).

Holds/contains. For a variety of artifacts, the referent of the combination functioned to hold or contain the referent of the modifier (e.g., *pencil cup, book rack, coat hook, printer cart*).

Depicts. Some artifacts involved symbolic depictions of the referent of the modifier. For example, *road atlas* is an atlas which shows representations of roads rather than actual roads, and *caution sign* symbolically depicted caution with the words ''caution wet floor'' printed on it.

Property Categories

Object-part. The referent of the modifier resembles the referent of a part of the combination (usually in terms of shape). For example, the *northern pintail* refers to a bird whose tail has the shape of a pin, the *eastern fence lizard* has scales that resemble those of a fence, *spiderflowers* have flowers which are shaped like spiders, the *birdfoot violet* has leaves which have the shape of birdfeet, and the *snowshoe hare* refers to a hare whose feet resemble those of a snowshoe in terms of having dense pads.

Part-part. The referent of the modifier has a part which resembles a part of the referent

of the combination. For example, *porcupine fish* refers to a fish with spines which resemble the quills of a porcupine, the *mule deer* refers to a deer whose ears resemble those of a mule, and the *willow oak* refers to a tree whose leaves resemble the leaves of a willow.

Color/pattern. The color or texture pattern of the referent of the modifier resembles that of the referent of the combination. For example, the *snow goose* has white down which is fluffy like snow, the *leopard lizard* has spots just as leopards do, the *fox grape* refers to a plant with leaves whose undersides are covered with whitish or redish woolly hairs, the *salmon entolomoa* (a mushroom) is salmon-colored, and the *cardinal flower* has red flowers which are the color of the female cardinal.

Overall shape. The shape of the referent of the modifier resembles the overall shape of the referent of the combination (e.g., *candy barrel cactus, guitar fish, peninsula table*).

Appearance. The referent of the combination bears a striking resemblance overall to the referent of the modifier or the usual referent of the combination. For example, the *bird's nest fungus* looks very similar to a bird's nest with eggs in it.

Behavior/function. The referent of the combination shares a behavior or function of the referent of the modifier. For example, the *mole salamander* burrows, the *kangaroo rat* hops, the *surgeon fish* uses a switchblade-like spine as a defense mechanism, the *pen knife* can be clipped on front pocket just like a pen can be, the *bulldog clip* has a strong grip analogous to the one of a bulldog, and the *compass plant* has leaves which point in a north-and-south direction.

REFERENCES

Abbot, K. (1989). *Downstream from trout fishing in America: A memoir of Richard Brautigan.* Santa Barbara, CA: Capra Press.

Bower, G. H., Black, J. B., & Turner, T. F. (1979). Scripts in memory for text. *Cognitive Psychology,* **11,** 177–210.

Bugelski, B. R., & Alampay, D. A. (1961). The role of frequency in developing perceptual sets. *Canadian Journal of Psychology,* **15,** 205–211.

Cohen, B., & Murphy, G. L. (1984). Models of concepts. *Cognitive Science,* **8,** 27–58.

Coolen, R., van Jaarsveld, H. J., & Schreuder, R. (1991). The interpretation of isolated novel nominal compounds. *Memory & Cognition,* **19**(4), 341–352.

Dermek, A. (1989). *The spotter's guide to mushrooms and fungi.* New York: Dorset Press.

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

Eschmeyer, W. N., Herald, E. S., & Hammann, H. (1983). *A field guide to Pacific coast fishes of North America.* Boston: Houghton Mifflin Company.

The facts on file visual dictionary (1986). New York: Facts on File Publications.

Gagne, 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.

Garrod, S. C., & Sanford, A. J. (1994). Resolving sentences in a discourse context: How discourse representation affects language understanding. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics.* New York: Academic Press.

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

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

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons. Beyond similarity. *Psychological Review,* **97,** 3–18.

Glucksberg, S., Manfredi, D. A., & McGlone, M. S. (1997). Metaphor comprehension: How metaphors create new categories. In Ward, T. B., Smith, S. M., & Vaid, J. (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 327–350). Washington, DC: American Psychological Association.

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

Luchins, A. S. (1942). Mechanization in problem solving. *Psychological Monographs,* **54**(6), Whole No. 248.

Markman, A. B., & Gentner, D. (1993a). Structural alignment during similarity comparisons. *Cognitive Psychology,* **23,** 431–467.

Markman, A. B., & Gentner, D. (1993b). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language,* **32,** 517–535.

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.

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.

Office stores program catalog. (1995). Evanston, IL: Northwestern University General Services Administration.

Reader's digest magic and medicine of plants. (1986). Pleasantville, NY: The Reader's Digest Association, Inc.

Reader's digest North American wildlife. (1992). Pleasantville, NY: The Reader's Digest Association, Inc.

Schank, R. (1972). *Conceptual information processing.* New York: Elsevier.

Schank, R. C., & Abelson, R. (1977). *Scripts, plans, goals, and understanding.* Hillsdale, NJ: Erlbaum.

Shoben, E. J., & Gagne, C. L. (1997). Thematic relations and the creation of combined concepts. In Ward, T. B., Smith, S. M., & Vaid, J. (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 31–50). Washington, DC: American Psychological Association.

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

Urdang, L., & Abate, F. R. (1983). *Idioms and phrases index.* Detroit, MI: Book Tower.

Webster's collegiate dictionary. (1995). New York: Random House.

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., & Bassok, M. (1997). *Stimulus compatibility with comparison and integration.* Manuscript submitted for publication.

Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 241–284). Amsterdam: North Holland.

Wisniewski, E. J., & Markman, A. B. (1993). The role of structural alignment in conceptual combination. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 1083–1086). Boulder, CO: Erlbaum.

(Received January 2, 1997)
(Revision received June 30, 1996)