

When concepts combine

EDWARD J. WISNIEWSKI

Northwestern University, Evanston, Illinois

I present a computational level account of how people combine concepts, and I use this account to evaluate current models of conceptual combination. Constrained by this account, I then provide an algorithmic level description of how people combine concepts. The algorithmic level account highlights the importance of two additional processes (comparison and construction) in explaining how some concepts combine and change. I then show that the interpretation of nominal metaphors involves these processes as well. Current approaches to metaphor understanding emphasize the importance of one or the other of these processes, but not both.

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.¹ Recently, for example, a newspaper article used the phrase *car boat* to refer to a new kind of boat that is also a car, and in a grocery store, *seafood sausage* was used to refer to patties of ground seafood. Like analogies and metaphors, such novel combinations are an important language construction. People create them in order to specify referents of discourse contexts and to extend the vocabulary of their language (E. V. Clark, personal communication, 1995; Downing, 1977). Besides revealing their role in language, the study of novel combinations can also inform theories of concepts (Barsalou, 1991; Hampton, 1987; Markman & Wisniewski, 1997; Medin & Shoben, 1988). Studies of conceptual combination, for example, have identified ways in which prototype theories need to be extended (Medin & Shoben, 1988) as well as differences between the conceptual structure of superordinates and basic level concepts (Markman & Wisniewski, 1997).

Achieving a good understanding of how people combine concepts is a challenging task. There are few straightforward rules for how one concept combines with others, especially when the concepts are nouns (Downing, 1977; Murphy, 1988; Wisniewski & Gentner, 1991). For example, the same noun can be combined with other nouns to produce a variety of meanings: *corn oil* means "oil made from corn," but *baby oil* means "oil rubbed on babies," and *lamp oil* means "oil for lighting lamps." As this example illustrates, there is no simple rule for how *oil* combines with other concepts to produce these meanings.

However, the meanings of nouns constrain the ways in which they can be sensibly combined. It would be very implausible, for example, to interpret *diamond oil* as "oil for lighting diamonds." The context in which a novel combination occurs also constrains its interpretation (Gerrig & Murphy, 1992; Murphy, 1990).

There are at least three levels on which one can understand conceptual combination. Marr (1982) refers to these levels as the *computational*, *algorithmic*, and *implementational* levels. The computational level addresses both the goal or purpose of conceptual combination (why people combine concepts) and what results from combining concepts (the output of conceptual combination). The algorithmic level describes procedures for how concepts are combined to produce the output. Most psychological approaches to understanding conceptual combination have been formulated at this level (Cohen & Murphy, 1984; Heit & Barsalou, 1996; Martin & Billman, 1994; Shoben, 1993; Shoben & Gagne, 1997; Smith, Osherson, Rips, & Keane, 1988; Thagard, 1984). Finally, the implementational level specifies how the algorithm is realized in a physical device, such as a human brain or a digital computer. These levels are not independent, and an explanation at one level can constrain that at another level. For example, an algorithmic account of conceptual combination depends on knowing what the output of that algorithm should be (an aspect of the computational level account).

In this paper, I provide a computational level account of conceptual combination and relate this account to present algorithmic (processing) level descriptions. One conclusion is that models fall short in explaining the basic phenomena that characterize conceptual combination. Perhaps of greatest importance, concepts *change* when they combine, and current models do not go far enough in accounting for this change. Constrained by this computational level account, I then provide an algorithmic level account of how concepts combine and change. This account highlights the importance of two additional processes (comparison and construction).

Finally, I show that the interpretation of some noun-noun combinations mirrors the interpretation of nominal

I thank Miriam Bassok, James Hampton, Doug Medin, Tom Ward, and two anonymous reviewers for trenchant comments on previous drafts, Kristen Carpenter and Beth Lynch for drawing the pictures, and Bob Dylan for providing some of the inspiration. Correspondence should be addressed to E. J. Wisniewski, Northwestern University, Department of Psychology, Swift Hall, 2029 Sheridan Rd., Evanston, IL 60208-2710 (e-mail: edw@nwu.edu).

metaphors, implying that similar processes operate in both domains. In particular, models of metaphor understanding require processes of both comparison and construction. Some models, however, focus on the comparison process and do not emphasize construction of new knowledge. They view conceptual change as a straightforward transfer of knowledge from one concept to another. Other models note the importance of construction but do not emphasize comparison. I will argue, however, that both processes are crucial to deriving interpretations of metaphors and accounting for conceptual change.

A Computational Level Account: Why People Combine Concepts

Conceptual combination occurs in a communicative context and serves at least three goals (see Downing, 1977, for a related discussion). First, people create novel combinations in order to designate significantly new categories: ones that have important, enduring characteristics that distinguish them from similar categories. For example, *seafood sausage* refers to a new category that differs from other types of sausage in its composition, cooking time, taste, and so on. These differences are stable characteristics that importantly affect how we interact with this category. The use of *seafood sausage* helps people to pick out this category and to infer its important properties. Phrases created with this goal in mind are apt to become part of the vocabulary of a language. In contrast, the use of some novel phrases imply more temporary relationships. For example, a person was once told to sit in the apple-juice seat (i.e., a chair in front of which a glass of apple juice had been placed; Downing, 1977, p. 818). This phrase was created under special circumstances to achieve a very specific goal rather than to designate a significant new category. It is less apt to become part of the vocabulary of a language. Second, combinations are used to convey information in a concise and efficient way. For example, *football parking* designates an area for parking one's car while one attends a football game. Even though this phrase is somewhat elliptical, readers generally understand what it means. Third, combinations may function as anaphora, in that they are used to refer back to a previous referent in a discourse context. For example, after describing a man who received the first artificial heart (Barney Clark), a speaker may use the phrase *heart man* to assert new information about this man, and listeners can determine its referent and construct its interpretation (Gerrig & Murphy, 1992). The use of anaphora helps to link information to the appropriate referent and thus establishes cohesion in discourse contexts (Garrod & Sanford, 1994).

In achieving these three goals, both the speaker and the listener implicitly assume certain constraints on the production and interpretation of a novel combination. First, the combination refers to a category that differs in some way from other categories named by the head noun (i.e., the rightmost noun). For example, *seafood sausage* refers to a type of sausage that differs from other kinds

of sausage. Second, the source of this difference can be determined from the modifier noun (i.e., the leftmost noun): *seafood sausage* is sausage *made of seafood*. Third, despite the differences, the referent of the combination still shares significant commonalities with the head noun category. So, even though *seafood sausage* is different from other kinds of sausage, it has enough in common with other types of sausage so that it is still called sausage. In general, these constraints are consistent with basic goals of communication, such as informative and relevant, and to avoid ambiguity (Grice, 1975; Wilson & Sperber, 1981).

In addition to explaining why people combine concepts, a computational level account specifies the output or result of conceptual combination. Determining what people compute when they combine concepts defines the generality to be attained by the model. Furthermore, knowing what people compute constrains the model's processing assumptions. For example, a psychologically plausible model should not compute things that people do not compute (cf. Pinker and Prince's (1988) critique of Rumelhart and McClelland's (1986) connectionist model of learning the past tense of English verbs).

A Computational Level Account: The Output of Combining Concepts

One way to determine the output that results from combining concepts is to examine interpretations of novel combinations. In a number of studies, I have asked subjects to assume that they have heard a novel combination in a conversation and to describe its most plausible meaning (Markman & Wisniewski, 1997; Wisniewski, 1996a; Wisniewski & Gentner, 1991; Wisniewski & Love, 1996; Wisniewski & Markman, 1993). The combinations are novel in the sense that they do not occur as lexicalized entries in dictionaries and people have generally not heard or seen them before. In many cases they were created by the arbitrary or random pairing of their constituents, which generally produces novel combinations. Although the combinations are novel, their constituents name common, familiar things such as human-made artifacts (e.g., *box, car, trumpet*), animals (e.g., *elephant, squirrel, fish*), plants (e.g., *cactus, tree, daffodil*), and substances (e.g., *clay, copper, chocolate*). In some experiments, I have used abstract terms (e.g., *promise, moral*) and events (e.g., *talking, writing*) (Wisniewski & Markman, 1993). Subjects typically proceed through the task at their own pace, writing down meanings for 15–20 novel combinations. From these studies, I have collected a very large and fairly representative set of interpretations involving approximately 500 novel combinations with 10–20 interpretations per combination (Wisniewski, 1994).

Kinds of combinations. Analyses of these interpretations suggest that people combine concepts in three basic ways (Wisniewski, 1996b). *Relation-linking* interpretations involve a relation between the referents of the modifier and head concepts. For example, people sometimes interpret *robin snake* as "a snake that *eats* robins."

In *property* interpretations, people assert that one or more properties of the modifier concept apply in some way to the head concept, as in “snake with a *red underbelly*,” for *robin snake*. A third, less frequent type of interpretation is *hybridization*. These interpretations refer to a combination of the constituents (e.g., a *robin canary* is “a bird that is a cross between the two—half robin and half canary”) or to a conjunction of the constituents (e.g., a *musician painter* could refer to someone who is both a musician and a painter). Note that logically, there could be other kinds of interpretations. For example, a combination could refer to a disjunction of its constituents. However, this combination type and others do not appear (at least in English).

There are important conceptual distinctions between these combination types. If, for example, one pictures a robin snake, interpreted by means of relation linking as “a snake that eats a robin,” one sees a snake eating a robin. However, if one pictures a robin snake, interpreted as having a property, as “a snake with a red underbelly,” one sees only a snake (whose underbelly has a color similar to that of a robin). As another example, *box clock* interpreted by means of relation linking as a “clock inside a box” involves a scene with two objects, one contained in the other. On the other hand, *box clock* interpreted as having a property and hence as “a square box,” refers to a single object with a box-like shape. There are also important conceptual distinctions between hybrids and relation-linking and property interpretations—for example, a *robin canary* interpreted as having a property and therefore as a “canary with a red breast” and a *robin canary* interpreted by means of relation linking as a “canary that preys on robins” are both canaries, but a *robin canary* that is a cross between the two is not.

I have focused on novel combinations in contrast to familiar ones because it is more difficult to infer interpretive processes from the latter. Meanings of familiar combinations have often been modified or augmented after their initial creation (Downing, 1977; Levi, 1978). For example, eggplant once referred to a small, white, egg-shaped vegetable, but it now typically refers to a larger, purple, oblong vegetable (Elliot, 1988). Nevertheless, it is easy to find familiar combinations that belong to these categories. Interpretations involving relations include *dish washer*, *tent peg*, *elephant gun*, *horse fly*, *paper weight*, and *stapler gun*. Those involving properties include *cat fish*, *soap stone*, *banana pepper*, *squirrel monkey*, *string tie*, *mule deer*, and *box fan*. Hybrids include *apartment hotel*, *toaster oven*, *clock radio*, *house boat*, *singer songwriter*, and *rose pink*.

Referential scope. The use of a novel combination indicates one or more new referents. In many combinations, these referents are subsets of the things typically named by the constituents. For example, people might refer to a kind of jar that is used to hold very small books as a *book jar*. The books and jars referred to by *book jar* are subsets of the things that we would typically name with the terms *book* and *jar*. In many respects, this as-

sumption is consistent with the goals of the computational level account. A speaker who uses the typical referent of a noun may be more likely to succeed in conveying what he or she intends the listener to understand. However, people are also flexible in their interpretation of a noun’s referent. The potential referential scope of a noun is broader. Property interpretations are a good example of this flexibility: people construe the modifier noun as referring to a property or characteristic of the thing named by the noun rather than the thing itself. For example, in interpreting *skunk squirrel* as “a squirrel with black and white stripes,” the referent of *skunk* is one of its properties.

There are other ways in which people interpret a noun’s referent. In what I have called *representation construal*, people treat a noun as referring to a representation of a thing rather than the thing itself. For example, *moose pencil* was interpreted as “a pencil with an eraser that looked like a moose,” *car box* was interpreted as “a box that contained a toy car,” *stone squirrel* was interpreted as “an ornamental squirrel made of stone,” and so on (Wisniewski, 1996a). These interpretations refer to representations of *moose*, *car*, and *stone*, respectively, rather than to mooses, cars, or stones, as in “pencil stepped on by a moose,” “box found inside a car,” and “squirrel that lives among stones.” People also interpret nouns as referring to things associated with a constituent, such as something that is thematically related to the constituent. For example, people often interpret *artist collector* as a “collector of the *works of an artist*,” *tiger chair* as “a chair made of *tiger skin*,” and *robin termite* as “a termite that eats *robin’s nests*.”

People are quite willing to flexibly construe the referent of a noun, at least when interpreting novel combinations presented outside of discourse contexts. Wisniewski (1996a) found that among 3,003 meanings of 224 novel combinations given by 189 subjects, an astonishing 70% involved the construal of a noun’s referent as something other than the typical category named by the noun. This flexibility is also suggested by a major theory of metaphor understanding which assumes that nouns have dual reference (Glucksberg & Keysar, 1990; Glucksberg, McGlone, & Manfredi, in press). When used literally, a noun refers to its typical category but when used metaphorically it refers to a property that 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 a sense of being confined.

Like the basic combination types, these kinds of construals are seen among familiar combinations. *Animal crackers*, *alphabet soup*, *piggy bank*, and *chocolate easter bunny* show representation construal. In our culture, the function of many human-made artifacts is to represent (e.g., *replica*, *mannequin*, *painting*, *tattoo*, *photograph*, *carving*, *television*, *diagram*, *book*, *model*, etc.), and so in hindsight it is not surprising that representation construal is a common way of combining concepts. Familiar combinations that involve construing a referent to be a

thematic associate include *football parking* (described previously), *pepper pot* ("a West Indian stew made of highly seasoned vegetables and meat or fish"), *safflower oil* ("oil made from the seed of the safflower herb"), and *xerox key* ("key to the room that contains the xerox machine").

Of importance, adopting this broader referential scope for a noun helps us achieve the goals of conceptual combination. From a communicative perspective, the use of novel combinations of familiar terms efficiently captures new situations. For example, using *artist* to refer to the works of an artist allows someone to use the novel phrase *artist collector* to refer to a collector of the works of an artist (see also Nunberg, 1979). People almost always derive this meaning in the absence of any context. Without this flexibility, people might have to use a longer and perhaps more awkward sounding phrase, such as "the collector of the works of the artist."

At the same time, this extended referential power has limits and is constrained. In general, people construe a noun's referent in only a few basic ways, and the referents are conceptually related to the noun's typical category. That is, the extended referential scope of a noun appears limited to a property, representation, or thematic associate of the category.

Conceptual change. Concepts change when they are combined. Consider the novel phrase *zebra clam*, interpreted as a "clam with stripes." Combining *zebra* and *clam* produces a new concept by modifying *clam*. What is the nature of the change? One possibility is that the property *stripes* is represented hierarchically at multiple levels of abstraction in *zebra* (Barsalou, 1993; Marr, 1982; Wisniewski & Medin, 1994) and that people transfer a fairly abstract representation of this property after doing some minimal processing for plausibility. For example, in interpreting a phrase like *zebra clam*, people may assess the plausibility of a clam's having stripes by noting that stripes are a type of texture, that objects have textures, and that clams are objects with textures that do not already have stripes. Thus, they may conclude that *zebra clam* could plausibly mean "clam with stripes." Therefore, the concept *zebra clam* would be the concept *clam* plus an abstract representation of *stripes*. In effect, this abstract representation of *stripes* has been copied from *zebra* and added to *clam*.

Alternatively, a property in the modifier acts as a source of information for instantiating a *new version* of that property in the head concept. The new version is constructed to fit constraints specified by both the modifier

and head concepts. For example, the stripes of a zebra clam might show an alternating dark and light pattern (as they do in a zebra), but they might also be shorter and thinner than those of a zebra, in order for them to be consistent with the size of a clam. Thus, the concept *zebra clam* would be the concept *clam* plus a new version of stripes that has been instantiated through an interactive construction process that is sensitive to constraints specified by the modifier and head concepts. The new version of *stripes* resembles but is not identical to the *stripes* in *zebra*.

Wisniewski (1996b) provides evidence that people carry out this construction process rather than a copy and addition process. If people actually construct properties, the instantiation of a property in a combination (e.g., the instantiation of *stripes* in *zebra horse*) should match its corresponding instantiation in the modifier concept (e.g., the instantiation of *stripes* in *zebra*) to the extent that constituents of the combination are similar. Thus, the stripes of a zebra horse should resemble those of a zebra more than the stripes of a zebra clam should resemble those of a zebra. To examine whether subjects would be sensitive to these differences, I constructed quadruplets of combinations whose interpretations attributed a property of the modifier to the combination. Table 1 presents some examples. Notice that the interpretations for each quadruplet involve an identical property but that the constituents of a combination vary in their similarity to each other. Furthermore, each modifier and head noun appear in both a similar and a dissimilar combination.

In one study, subjects saw these combinations and their interpretations presented in pairs (e.g., they saw *cactus dandelion* and *cactus pig* and their corresponding interpretations). For each pair, subjects selected the interpretation that sounded more natural to them. The idea is that when subjects read an interpretation they will instantiate the property in the combination and notice the degree to which it matches the instantiation of this property in the modifier. Thus, for the similar combinations, the modifier will better capture the sense of the property in the combination. Subjects overwhelmingly preferred the interpretations of the similar combinations.

In a second study, subjects were more likely under time pressure to judge that interpretations of similar combinations were more plausible than those of dissimilar combinations. For example, subjects were more likely to judge as plausible that *snow sand* is white sand than that *snow soda* is white soda. This result cannot be explained

Table 1
Quadruplets of Novel Combinations and Possible Interpretations

Similar Combinations	Dissimilar Combinations
snow soda: white soda	milk soda: white soda
milk sand: white sand	snow sand: white sand
cactus pig: prickly pig	porcupine pig: prickly pig
porcupine dandelion: prickly dandelion	cactus dandelion: prickly dandelion
alarm clock sparrow: sparrow that wakes you up	alarm clock tv: tv that wakes you up
rooster tv: tv that wakes you up	rooster sparrow: sparrow that wakes you up

Note—Similar combinations are presented on the right.

by suggesting that it is more plausible that sand rather than soda can be white, because subjects were also more likely to judge as plausible that *milk soda* is white soda than that *milk sand* is white sand. For similar combinations, the instantiation of a property in the modifier better matched its instantiation in the combination. Thus, subjects considered these interpretations more plausible when making speeded judgments.

In a final study, subjects provided interpretations of the similar and dissimilar combinations. The degree to which the instantiation of a property in the modifier matches its instantiation in the combination makes predictions about how subjects will phrase their interpretations. The idea is that subjects follow cooperative principles in which they attempt to accurately phrase their interpretation (Grice, 1975). Thus, for similar combinations, subjects should be more likely to explicitly describe a property of the combination as (also) belonging to the modifier or as resembling one of the modifier. For example, subjects should be more likely to say that *zebra horse* is a "horse with the stripes of a zebra" or "horse with stripes that look like those of a zebra's." However, for *zebra clam*, they should say "striped clam" or "clam with stripes" rather than "clam with the stripes of a zebra." This difference in phrasing was found between the interpretations of similar and dissimilar combinations.

Figure 1 shows artistic renditions of the output of this construction process. These illustrations correspond to not infrequent interpretations that subjects gave to the novel combinations *fork spoon* and *rake pencil* (Wisniewski, 1996a). The two illustrations of *fork spoon* highlight different ways in which the same property of *fork* ("has prongs") is instantiated or realized in this combination.²

One objection to these findings is that they involve very unusual or bizarre combinations (e.g., *rake pencil*, *alarm clock sparrow*), contrived by the experimenter, which do not occur in everyday language use. Thus, according to this argument, people rarely employ these extreme property construction processes. Again, it is easy to find familiar combinations that correspond to these forms of property construction. Figure 2 illustrates the referents of a variety of familiar combinations. The first column highlights the property as manifested in the modifier. The next two columns illustrate how the property is differentially instantiated in two combinations.

Summary. An extensive analysis of the meanings of novel combinations suggests that there are three basic kinds of interpretations. People are also flexible in how they interpret the referent of a noun. Although people interpret a noun as referring to the object or thing typically named by that noun, the referent can also be a property, a representation, or a thematic associate of the object. Of importance, combining concepts can involve property construction in which a property of the modifier is used to guide the creation of a new property in the combination. Concepts are not combined simply by adding a copy of a property of the modifier to the new combination.




Figure 1. Illustrations of the referents of people's interpretations of some novel combinations.

The Algorithmic Level: Current Models of How People Combine Concepts

There have been a number of processing accounts of how people combine concepts. Some approaches have had a fairly narrow focus. In particular, several models describe how people interpret hybrid combinations—those that refer to a conjunction of their constituents, such as *pet fish* (Hampton, 1987; Martin & Billman, 1994; Thagard, 1984). There is also a model of how certain adjectives are combined with nouns (Smith et al., 1988). These models provide accounts of only a small range of conceptual combination phenomena. They cannot be straightforwardly extended to account for the majority of noun–noun combinations—that is, to those which are interpreted with the use of relations or properties (see Murphy, 1988; Wisniewski, 1996a; Wisniewski & Gentner, 1991, for arguments). Therefore, I will not discuss them further.

There are two processing accounts which are broader in scope: the thematic relations approach and the schema approach. Both accounts assume that the meaning of a noun–noun combination involves a *relation* between the



Figure 2. Illustrations of the referents of some familiar combinations.

referents of the constituents. In the *thematic relations* view, nouns are combined by determining a fairly general relation that holds between the nouns (Coolen, van Jaarsveld, & Schreuder, 1991; Gagne & Shoben, 1993; Shoben & Gagne, 1997). These approaches assume that there is a relatively small set of such relations (perhaps one or two dozen). In developing this view, researchers draw on the relations that linguists have used to classify familiar combinations (e.g., Downing, 1977; Levi, 1978). For example, Levi suggested that 16 relations can be used to classify the meanings of most familiar combinations (e.g., the FOR relation, *plant food*; the CAUSE relation, *electric shock*; the MAKE relation, *honey bee*; the FROM relation, *peanut butter*, etc.).

The most developed view of the thematic relations approach is the work by Shoben and his colleagues. They assume that a noun's combinatorial history influences the interpretation of a novel phrase involving that noun. That is, people use the distributional knowledge of how nouns have previously combined to interpret a novel

combination. For example, when *mountain* is used as a modifier, it typically instantiates a locative relation (e.g., *mountain stream*, *mountain resort*, *mountain goat*) and is only rarely involved in other types of relations (e.g., *mountain range*). Therefore, people may 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. A second assumption is that the *modifier's* combinatorial history has more influence on interpretation (perhaps because it precedes the head noun). Shoben and Gagne (1997) show that the time to judge the sensibility of a novel combination is a function of the relative frequency of thematic relations associated with the modifier noun rather than with the head noun.

In the *schema approach*, a concept is viewed as a schema or frame. A schema represents the basic knowledge that we have about a place, event, or object which has been acquired from our everyday interactions with these things or from other sources of knowledge such as

written materials. Schemas represent this knowledge in the form of slots and fillers that refer to the dimensions of the entity, along with their typical or default values (Minsky, 1975; Rumelhart, 1980). For example, a schema for *elephant* might include the slots *color* and *habitat* and the typical fillers *gray* and *zoo*, respectively. Of importance, the fillers of the schema can themselves be representationally complex, also consisting of schemas with slots and fillers (cf. Ortony, 1979). As a result, processes that operate on a schema could be recursively applied to its fillers (Barsalou & Hale, 1992).

Schema models link concepts by relations through a process of slot filling (Brachman, 1978; Cohen & Murphy, 1984; Gerrig & Murphy, 1992; Murphy, 1988). For example, according to the concept specialization model (Cohen & Murphy, 1984; Murphy, 1988), one interprets a noun–noun combination by filling a slot of the head noun with the modifier noun. Thus, one might interpret *robin snake* by filling a slot in *snake* (e.g., the slot *eats*) with the modifier concept *robin* to produce the meaning “a snake that eats robins.” As a result, the filled slot captures a relation between the objects denoted by the modifier and head concepts. Of course, not all slots represent relations between objects (e.g., the color slot in *elephant* with the filler *gray* denotes a default property of *elephant*).

To explain which slot is selected to be filled by the modifier, these models emphasize the importance of constraints on the fillers of a slot. That is, a slot specifies preconditions that must be met by a potential filler. For example, the filler of the *eats* slot of *snake* would have to be edible (Brachman, 1978). These constraints are derived from people’s typical interactions with the referents of concepts.

Although the thematic relations and schema views are similar in assuming that interpretation involves linking one referent to another via a relation, they have two important differences. First, they emphasize different levels of abstraction in characterizing the meanings of combinations. Whereas the thematic relations view postulates that a relatively small set of general relations determines the ways in which nouns can be combined, the schema approach does not make an analogous claim for slots. Rather, slots represent more specific relations. Second, the thematic relations view emphasizes a noun’s combinatorial history in determining the interpretation of a novel combination, whereas the schema view emphasizes constraints on slots.

The schema approach better accounts for how people combine concepts. First, the use of slots rather than more general relations captures distinctions between meanings to which people are clearly sensitive. For example, in the thematic relations view, the novel combinations *paint spoon* and *blueberry spoon* would share the FOR relation between their constituents. However, the FOR relation fails to capture crucial differences between the interpretations of these combinations: a paint spoon is used to stir paint (and not to eat paint), but a blueberry spoon is used to eat blueberries (and not to stir them). These dis-

tinctions are part of the basic, core meaning of such novel combinations; they arise from the specific knowledge in the constituent concepts.

The schema approach could capture these differences by assuming that the concept *spoon* represents the different functions of a spoon (*stirring* and *eating*) and the slots associated with these functions. For example, both functions would have *agent*, *instrument*, and *object* slots along with appropriate constraints (e.g., the filler of the *object* slot in *eating* must be edible). This knowledge is derived from the usual interactions that people have with spoons. In *paint spoon*, *paint* would fill the *object* slot of *stirring* but not the *object* slot of *eating*, because paint can be stirred but is not edible. In *blueberry spoon*, *blueberry* would fill the *object* slot of *eating* but not the *object* slot of *stirring*, because blueberries are edible and perhaps because they cannot be stirred without destroying their integrity. People’s verbal descriptions of the meanings of novel combinations more closely correspond with this level of abstraction (Downing, 1977; Wisniewski, 1994, 1996a).

Second, the thematic relations view overly emphasizes the role of a noun’s combinatorial history in determining meaning. On the one hand, how one noun combines with another is influenced by how it has previously combined with other nouns (Shoben & Gagne, 1997). However, combinatorial history may be more of an effect rather than a cause. For example, most familiar combinations of the form *X box* (in which *X* stands for another noun) have the meaning “box that contains *X*” (Urdang & Abate, 1983). However, it is not the frequency of these past meanings per se that determines whether people interpret a novel *X box* in a similar way. There are combinations of this form that people do not interpret in this manner (e.g., *skyscraper box*). Rather, as emphasized in the schema approach, interpretation involves determining whether a noun fits the constraints associated with a relation. That most familiar “*X boxes*” have the interpretation “box that contains *X*” more likely reflects the intuition that there are few constraints on what can be contained in a box (since most things can be enclosed). A noun’s combinatorial history may suggest candidate meanings (and thus narrow down the search space), but most of the work in interpretation consists of evaluating the plausibility of a meaning and creating a new representation.

Although the schema approach is a more plausible psychological model, the computational level account suggests that both the thematic relations and schema views do not go far enough in explaining how people combine concepts. They have two important limitations. First, these models have focused on processes that derive relational interpretations of combinations and have not addressed the derivation of property and hybrid interpretations. As I will suggest later, property and hybrid interpretations require processes that differ from those that determine relations between the modifier and head concepts.

Second, these approaches capture relatively small changes in representation that result when concepts com-

bine. In the thematic relations approach, a new concept is created by linking one constituent to another via a relation. Being linked to another constituent is the only sense in which a constituent's representation undergoes modification. In the schema approach, a new concept is created by filling a slot of the head concept with the modifier concept (Cohen & Murphy, 1984; Gerrig & Murphy, 1992; Murphy, 1988). There is an additional stage of processing that also involves conceptual change. After slot filling, people use their world knowledge to elaborate or "fix up" this concept. World knowledge lies outside of the constituent concepts. According to Murphy (1988):

It isn't possible to form the complete concept if one knows only the two concepts involved; instead, one must have access to a very large data base—one's world knowledge—that can influence the process. (p. 533)

An important source of world knowledge is *extensional feedback* (Hampton, 1987), which refers to people's actual knowledge about objects in the world. Murphy (1988) describes how this knowledge is used to elaborate the meaning of *apartment dog* (p. 540). Once an initial interpretation is constructed by slot filling, people may recall actual dogs they have met and use these memories to modify the concept. Modification involves adding attributes (e.g., "yappy and neurotic") to the combination.

Thus, the model's notion of conceptual change involves not only filling a slot in the head concept with the modifier to produce a new concept but also adding new attributes to this concept. However, as illustrated by the examples in Figures 1 and 2, conceptual change goes beyond the addition of attributes. Also, the model presents a somewhat bleak and unconstrained view of how this change comes about: people must rely on knowledge outside of the constituent concepts to elaborate their interpretations, and there is a vast amount of such knowledge that potentially can be used. However, much of conceptual change may be a function of the knowledge in the constituent concepts themselves coupled with appropriate processing assumptions. For example, in interpreting *zebra clam*, the discrepancy in size between a zebra and a clam importantly determines how the stripes are instantiated in *clam*. It is likely that knowledge about size is captured in the concepts of *zebra* and *clam*.

TOWARD A COMPREHENSIVE MODEL

In this section, I describe an algorithmic approach to conceptual combination which is intended to capture a broader range of psychological phenomena. The approach is an extension of the schema model of conceptual combination. It retains the representational assumptions of that model and uses a modified version of its slot-filling process to account for relation-linking interpretations. However, it augments the schema model with comparison and construction processes that are required for property and hybrid interpretations. Table 2 illustrates this relationship between the approaches.

My approach does not constitute a full-fledged model. For one thing, novel combinations occur in discourse contexts and such influences will need to be integrated into a model. For example, it is highly likely that discourse contexts will guide the operation of the hypothesized processes to varying degrees. Also, the model is formulated at a general level and has not yet been implemented in a computer program that simulates and predicts experimental results. Nevertheless, my approach identifies necessary processes that must go into constructing a more detailed model.

Relation-Linking Interpretations

In the augmented schema approach, schemas contain *scenarios*, which correspond to verbs describing actions, events, or states, with various roles. For example, *soap* would contain a *cleaning* scenario with roles such as *recipient*, *agent*, and *instrument*, respectively corresponding to what is cleaned, who did the cleaning, what was used to assist in the cleaning. Interpretations of novel combinations are frequently structured around verbs that implicitly relate the constituents to roles associated with the verb (e.g., *motorcycle screwdriver* is described as a "screwdriver for repairing motorcycles") (Wisniewski, 1994). Many researchers suggest further that verbs are decomposed into such roles (e.g., Fillmore, 1968; Gentner, 1981; Levin, 1993; Schank, 1972). Also, many familiar combinations have the pattern noun-nominalized verb (e.g., blood donor, book request), in which the constituents fill roles associated with the verb that has been nominalized (Levi, 1978).

Thus, generating a relation-linking interpretation amounts to creating a plausible scenario involving the constituents of the combination. In this case, a plausible scenario is one in which each constituent is bound to a different role within the same scenario. For example, a plausible interpretation of *truck soap* is "soap for cleaning trucks," because *truck* can be bound to the *recipient* role of *cleaning* (i.e., the thing being cleaned) and *soap* to the *instrument* role (what is used to do the cleaning). The idea is that a constituent is matched to the preconditions of the role. If there is a sufficient match, the constituent is assigned to that role.

The use of scenarios allows the augmented model to explicitly indicate the different roles that the modifier and head noun are playing in an action, event, or state. In contrast, the slot-filling process of schema models derives a relation-linking interpretation by finding a role or slot for the modifier to fill, without explicitly indicating the role that the head noun fills.

Table 2
Processing Differences Between the Schema and Augmented Schema Models

Phenomenon	Process	
	Schema	Augmented Schema
Relational interpretations	slot filling	scenario creation
Property interpretations	—	comparison & construction
Hybrid interpretations	—	comparison & construction

In general, this approach constructs relation-linking interpretations in a manner similar to the schema model. The slot-filling process of the schema model operates only on the *constituents* of the combination to produce relational interpretations. However, in the augmented schema model, I assume that the scenario creation process can be recursively applied to *associates* in a constituent's schema. This extension allows the model to capture relational interpretations such as "someone who collects *works* of an artist" (for *artist collector*) and "termite that eats *robin's nests*" (for *robin termite*) (Wisniewski, 1996a).

Property Interpretations

In order to derive a property interpretation, there must be an important difference between the modifier and head concepts that forms the basis of the interpretation. For example, interpreting *zebra clam* as "a clam with stripes" implies that clams typically differ from zebras in not having stripes. This (almost self-evident) observation is a consequence of the computational level goals of conceptual combination: the need to create a new category that differs from the category referred to by the head noun and whose difference is related to the modifier. Furthermore, the difference must then be used to construct a new property in the head concept. In the example above, "having stripes" of *zebra* must be used to create a new version of "having stripes" which is incorporated into *clam*. Thus, an algorithm that derives property interpretations must determine a difference between the modifier and head noun concepts and create new knowledge in the head concept.

I will argue that comparison and construction processes are essential for carrying out property interpretation. The *comparison* process determines commonalities and differences between the modifier and head concepts that form the basis of the interpretation. In doing so, the comparison process also determines "where" in the head concept new knowledge is to be integrated. The output of the comparison process is used by a *construction* process to create new knowledge in the combination.

Comparison. The importance of comparison processes in language understanding was originally emphasized in models of metaphor and analogy (Gentner, 1983, 1989; Holyoak & Thagard, 1989). Applied to conceptual combination, the idea is that people align or put into correspondence the structure of the modifier and head concepts. By aligning the structure that is *common* to the head and modifier concepts, people find differences that potentially form the basis for an interpretation. Any pair of concepts has both commonalities and differences. Finding commonalities leads to finding differences, because commonalities are interconnected or related to differences (Markman & Gentner, 1993a, 1993b). To illustrate, consider *zebra horse*. Roughly speaking, one puts the representation of a horse's body into correspondence with the representation of a zebra's body because they are similarly shaped and because they have similar conceptual relations to similar components (similar, verti-

cally oriented legs are connected to a similar underside of the body at similar places). By aligning this common structure, one finds an important difference between zebras and horses (having vs. not having stripes).

Many studies suggest that when people compare mental representations, they use the commonalities between these representations to find differences (Gentner & Markman, 1994; Markman & Gentner, 1993a, 1993b; Markman & Wisniewski, 1997). A general result is that subjects not only list more commonalities for similar than for dissimilar concepts but also list *more differences* for similar concepts. Of importance, 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 list "has four wheels versus two wheels" as a difference (Markman & Gentner, 1993b).

Besides determining differences between concepts, the comparison process suggests where a property can be integrated into the combination. For example, in placing the body and neck of *horse* into correspondence with the body and neck of *zebra* and noting a difference (having vs. not having stripes), the comparison process has determined where the stripes can be incorporated into *zebra horse*. In particular, they should run alongside the body and neck of a horse just as they do in a zebra. (If a zebra only had stripes on its tail, the comparison process would suggest that a zebra horse only had stripes on its tail.)

Of importance, there is evidence that a comparison process occurs in conceptual combination. The frequency of hybrids among combinations increases as the similarity between the modifier and head noun increases (Wisniewski, 1996a). This result strongly implies that subjects compare the modifier and head nouns. Because hybrids refer to a combination or conjunction of the constituents, they must have many important properties of *each* constituent. This constraint is difficult to meet unless both constituents are highly similar, since otherwise, it is likely that important properties of the constituents would conflict. For example, it is very difficult to interpret a *drill pamphlet* as a hybrid, since it would require both the function of a drill and a pamphlet and these properties conflict. If subjects compare representations, they should detect increasingly more commonalities as the similarity between constituents increases and thus be more likely to interpret a combination as a hybrid.

There are a variety of more specific ways in which this comparison process could be instantiated. A number of models of metaphor, analogy, and similarity use a local-to-global algorithm coupled with certain constraints that increase the efficiency of finding correspondences (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983; Goldstone, 1994; Holyoak & Thagard, 1989). Another approach would be to use a global-to-local algorithm (cf. Schyns & Oliva, 1994). For example, suppose that *zebra* and *horse* included representations of perceptual components with spatial extent (cf. Barsalou, 1993). Comparison could occur by first putting similarly located large

areas into correspondence (e.g., aligning the topmost part of one representation with the topmost part of the other) and then focusing the comparison on smaller corresponding spatial locations within these larger areas.

As a third possibility, a property of the modifier might be accessible prior to the comparison process. For example, a linguistic context might indicate the property of the modifier that is relevant to the novel combination (cf. Gerrig & Murphy, 1992). Alternatively, a prototypical property of the modifier might be rapidly activated upon reading of the modifier (cf. Barsalou, 1982; Glucksberg & Keysar, 1990). In these cases, comparison is still necessary in order to establish where the property is to be integrated into the head concept. For example, even though “having stripes” might be highly accessible upon reading *zebra*, people must still determine how it fits into *horse*. It is the correspondences between *zebra* and *horse* such as *body* and *neck* which suggest that the stripes run alongside the neck and body of the horse. These correspondences are determined by the comparison process. Furthermore, the comparison process may also be necessary to establish that the property of the modifier is not already present in the head noun concept. For example, a salient property of zebras is their shape, yet this property would not form the basis of an interpretation of *zebra horse*, because horses have a very similar shape. Presumably, the comparison process identifies this commonality and thus prevents it from being used in the interpretation.

To appreciate the advantages of a commonality-driven alignment process to find differences, consider a fairly simple alternative mechanism. In the example of *zebra horse*, one could begin by attempting to match a property of *zebra* with each property of *horse*. If the property of *zebra* failed to match all properties of *horse*, it would be considered a difference that could form the basis of the new combination’s meaning. This process could be made efficient by assuming that it operates in parallel. However, unlike the commonality-driven comparison process, identifying a difference by this method provides no information about *where* it should be integrated into combination. So, this method can determine that “has stripes” does not match any property in *horse* but cannot determine where this property should be integrated into *horse*.

Selecting a difference. In general, there are multiple differences between concepts. One issue is which difference is selected to form the interpretation. At least several factors influence selection. First, context may indicate the relevant property. Second, nouns are sometimes used to refer to salient properties (Glucksberg & Keysar, 1990). Thus people may consider these properties in forming their interpretations. For example, *elephant* and *whale* are sometimes used to refer to large things, as in the familiar combinations *elephant garlic* and *whale shark*.

A third factor may be the cue and category of a property associated with the noun. The *cue validity* of a property with respect to a category is the conditional probability that something belongs to a category, given that it has the property. For example, given that something has

wings, it is very probable that it is a bird (although it could be a bat or an insect). The *category validity* of a property with respect to a category is the conditional probability that something has the property, given that it is in the category. For example, given that something is a bird, it is very probable that it has wings. People may have a bias to construe nouns as referring to properties that are high in cue and category validity. This assumption makes sense, given the communicative goals of the computational level account. Properties high in cue and category validity are representative of the category and distinguish it from other categories. In the limit, a property with perfect cue and category validity applies to every member of the category and only to the category’s members. Therefore, in cooperative communication, a listener may likely think of properties high in cue and category validity in understanding a novel combination, and speakers may produce combinations with these properties in mind.

Selection is also a function of whether a property can be plausibly applied to the head noun concept. For example, people would probably never literally interpret *fork tennis ball* as “a tennis ball with prongs,” even though this interpretation is derived from a prototypical property of *fork* that is high in cue and category validity. Again, the communicative goals of the computational level account constrain selection. The interpretation “tennis ball with prongs” compromises the functionality of *tennis ball* and would result in a referent that is outside the acceptable scope for *tennis ball*.

Construction. As I previously showed, people generally do not interpret a combination by copying over a property of the modifier to the head concept. Rather, a property in the modifier acts as a source of information for constructing a new version of that property in the head noun concept. Construction is an interactive process, in which the new property is a function of constraints specified by both the modifier and head noun concepts. More specifically, in keeping with the computational level goals of conceptual combination, the new property must bear enough resemblance to its source in the modifier so that people can determine how the modifier contributes to the meaning of the combination (cf. Nunberg, 1979, for a related discussion on the creation of new word meanings). At the same time, the construction of the new property must not alter the head noun concept in such a way that it destroys its integrity (as in the example of *fork tennis ball* interpreted as “a tennis ball with prongs”).

To illustrate the construction process, consider how it might operate (in conjunction with comparison) to produce the interpretations illustrated in Figure 1. In interpreting *fork spoon*, people could begin by aligning the *handle* of *fork* with the *handle* of *spoon*, and the *end* of *fork* with the *end* of *spoon* and note an important difference: forks have prongs on their ends but spoons have “little bowls” on their ends (of course, other comparison algorithms are possible, as noted earlier). The comparison process identifies where in the representation of *spoon* the property “has prongs” can be incorporated (on

the *end of spoon*). However, there is a conflict between mentally connecting this property to the *end of spoon* and staying within the referential scope of *spoon*. In particular, one cannot mentally replace *little bowl* with *prongs* or mentally add the full length prongs to the *end of the little bowl*, because these manipulations would create fork spoons that functionally were forks but not spoons (thus violating communicative goals of the computational level account). People can resolve this conflict by mentally attaching the prongs to the end of the little bowl and shortening them (Figure 1a) or by mentally attaching the prongs to the top end of the spoon (Figure 1b). In these cases, the prongs of a fork spoon are similar but not identical to those of a fork: they are either attached to an end that is opposite to the little bowl or they are shorter and attached to the end of the little bowl. Furthermore, the attachment to the opposite end or the shortening preserves the function of a spoon while still yielding a property that resembles the one from which it originated. Thus, if someone points to the artifact depicted in Figure 1a or 1b and says “that’s a fork spoon,” a listener will understand why it is specifically called a fork spoon.

Analogous reasoning explains why the “teeth” of a rake are not simply added to *rake pencil*. A rake pencil with teeth exactly like those of a rake could not be a *pencil*. On the other hand, a rake pencil with multiple lead points (Figure 1c) would have a somewhat different but (apparently) acceptable variation on the typical function of *pencil* (one subject suggested that such a pencil was used to draw the lines on music paper). To construct this new property, a comparison process may use commonalities between *rake* and *pencil* to identify the difference “has teeth versus has a lead point.” In particular, the teeth of a rake and the lead point of a pencil have similar spatial positions relative to the rest of *rake* and *pencil*, respectively. They are also connected to similar things which may also be aligned (i.e., parts which are long and thin relative to the rest of *rake* and *pencil*, respectively). Further, the teeth and the lead point are similar in that both are pointy.

Once aligned, a new property is constructed in *rake pencil* which respects constraints specified by both *pencil* and *rake*. In particular, a pencil cannot have the teeth of a rake without compromising its functionality, so it has multiple lead points (which are at least shaped in a similar way to the teeth). Further, it cannot have as many lead points as there are teeth (but it does have more than one and thus preserves the multiplicity aspect of *teeth*). The rake pencil also preserves the perpendicular relation in *rake* between the handle and the base of the rake, which holds the teeth. Thus, if someone points to the artifact depicted in Figure 1c and says “that’s a rake pencil,” a listener will understand why it is specifically called a rake pencil.

In summary, some novel combinations are interpreted by using comparison and construction processes. In comparison, the representational elements of the modifier are aligned with ones in the head noun. The com-

parison process uses commonalities to find a difference between the modifier and the head noun which will form the basis of a new property to be constructed in the combination. Further, by identifying correspondences between the modifier and head concept, the comparison process helps to indicate where in the head noun the property should be integrated. Constrained by the computational level account, the construction process must create a new property in the combination which preserves enough of the sense of the modifier property from which it originated. Further, it must respect constraints provided by the head noun such that the resulting meaning maps onto a referent which is within the acceptable scope of the head noun.

Scenario creation versus comparison/construction. I have discussed the scenario creation process separately from the comparison/construction process. Another issue concerns the relationship between these processes. Do they operate concurrently, either as competing or as independent processes? Or, do the processes operate sequentially? A number of researchers have argued that people first attempt to derive a relation linking interpretation and only derive a property interpretation if there is no plausible relation between the modifier and head noun concepts (Downing, 1977; Shoben & Gagne, 1997; Wisniewski & Gentner, 1991). Thus, this assumption implies the sequential operation of processes with scenario creation operating initially and comparison/construction only applied if scenario creation fails to produce a plausible interpretation.

On first glance, this view appears to explain why combinations with highly similar constituents are very rarely interpreted by relation linking and almost always by property mapping and hybridization (Wisniewski, 1996a; Wisniewski & Markman, 1993). Recall that a plausible scenario is one in which each constituent fits or matches a *different* role within the same scenario (e.g., a plausible interpretation of *truck soap* would involve *soap* playing the *instrument* role and *truck* the *recipient* role in a *cleaning* scenario). However, highly similar constituents are highly likely to play the *same* roles in a scenario and therefore cannot be plausibly incorporated into the same scenario. For example, *magazine newspaper* is likely to activate a *reading* scenario with (at least) *agent* and *object* roles. However, both constituents compete for the object role (and neither constituent matches the *agent* role). An alternative scenario involving “newspapers which describe magazines” is implausible, given that newspapers typically describe the news.

On the other hand, some highly similar nouns nevertheless can play different roles in the same scenario (they are difficult to find). For example, a musician could play music for dancers. Recently, Wisniewski and Love (1996) used these nouns to test the “scenario creation first” hypothesis. Specifically, they created pairs of novel combinations that met the following constraints: one combination’s constituents were highly similar in comparison with those of its counterpart, the constituents in both combinations could involve the same scenario, and it was at

least as plausible that the highly similar constituents could be involved in the scenario as it was that the less similar constituents could be involved in it. For example, the constituents of *dancer musician* are more similar than those of *mourner musician*, both combinations can be interpreted using the *plays for* scenario, and subjects believe that “musician who plays for dancers” is more plausible than “musician who plays for mourners.” Other subjects were given these combinations to define. In general, the pattern of findings contradicted the scenario first hypothesis. For example, people tended to generate property or hybrid interpretations for *dancer musician*, *spear chisel*, and *goose vulture*. In contrast, they interpreted *mourner musician* as a musician who plays music for mourners, *barrel chisel* as a chisel used on barrels, and *fish vulture* as a vulture that preys on dead fish, even though the scenarios involving these nouns were judged to be less plausible than the same ones involving the highly similar constituents. This finding suggests that for highly similar combinations, scenario construction and comparison/construction processes compete. Since it is easier to align representations with similar structure and to find their commonalities and differences, the comparison/construction process “wins” over scenario creation even though there is a plausible interpretation involving a scenario.

THE RELATION BETWEEN CONCEPTUAL COMBINATION AND NOMINAL METAPHORS

Nominal metaphors are language constructions of the form *X is a Y*, in which *X* and *Y* refer to nouns. It is possible to rephrase a nominal metaphor as a noun–noun combination by reversing the order of the nouns and dropping the form of *be* to produce the phrase *YX*. For example, *that job is a jail* can be rephrased as *that jail job*. Both constructions can be interpreted as “a job that is confining.” However, this equivalence may only hold between nominal metaphors and noun–noun combinations that are interpreted by property construction. Nominal metaphors that correspond to hybrid and relation-linking interpretations of noun–noun combinations apparently do not exist. For example, it seems that people would not interpret *that job is a jail* by linking *job* and *jail* via a relation as in “that job is *at* a jail,” whereas *that jail job* can be interpreted in this manner.

This hypothesized relationship between noun–noun combinations and nominal metaphors is consistent with a recent theory of metaphor understanding. According to the interactive property attribution model, people understand a metaphor as a class inclusion statement whose vehicle has two referents—the literal referent and a (covert, unnamed) category described by one or more properties that the literal referent exemplifies (Glucksberg & Keysar, 1990; Glucksberg, Manfredi, & McGlone, in press; Glucksberg, McGlone, & Manfredi, in press). For example, in *that job is a jail*, the vehicle *jail* refers to both a literal jail (i.e., a building used to confine prisoners)

and to properties that are prototypical of *jail* (e.g., confining, unpleasant, punishing, etc.). People understand a metaphor by applying one or more of these properties to the topic. In this example, they apply properties such as *confining*, *unpleasant*, and *punishing* to the topic *job*.³

On this view, property interpretations of noun–noun combinations are similar to interpretations of nominal metaphor interpretations. Just as the modifier of a noun–noun combination is construed as a property that is then applied to the head noun, so is the vehicle of a nominal metaphor construed as a property that is then applied to the topic. However, noun–noun combinations are different from nominal metaphors in that they allow for other kinds of interpretations (i.e., relation linking and hybridization).

I recently conducted a study that provides evidence for this view. One group of subjects defined 26 novel noun–noun combinations of the form *that YX*. A second group defined phrases involving the same nouns but that were presented in the form *that X is a Y*. Among other things, the results showed a fairly striking correspondence between the interpretations of the noun–noun combinations and those of the nominal metaphors. The most common interpretation of every noun–noun combination was also the most common interpretation of the corresponding nominal metaphor. In addition, none of the metaphors were interpreted by relation linking or by hybridization. In contrast, relation-linking interpretations occasionally occurred among noun–noun combinations: “a thief who steals art” for *that artist thief*, “a teacher who trains students to be butchers” for *that butcher teacher*, “ocean on an earthquake” for *that earthquake ocean*, and so on.

Despite these commonalities and differences, there remain outstanding issues in regard to the relationship between noun–noun combinations and nominal metaphors. For one thing, the order of the nouns in a combination is the reverse of the order in a nominal metaphor. There are well-known order effects on the processing of information (Nisbett & Ross, 1980; Tversky, 1977). Also, there may be differences in the relative ease and speed with which combinations and metaphors are understood. The syntactic form of a nominal metaphor provides more constraints on its interpretation which may facilitate processing. Finally, an interesting issue concerns the conditions under which one phrasing is used as opposed to the other.⁴ People may prefer a nominal metaphor to describe a current situation but a combination to describe a habitual situation. For example, “that teenager is a volcano” might refer to a teenager presently revealing a bad temper whereas “that volcano teenager” may refer to a teenager who displays a bad temper with some regularity.

COMMON PROCESSES IN CONCEPTUAL COMBINATION AND METAPHOR

Given the correspondence between property interpretations of combinations and nominal metaphor interpretations, it is a short step to suggest that both phenomena involve the same processes. Previously, I argued that prop-

erty interpretation involves two processes: a comparison process that aligns the modifier and head nouns to determine differences between them, and a process that constructs a new property that is integrated into the head noun. In like fashion, nominal metaphors would be understood by comparing and aligning the vehicle with the topic to determine differences between them and constructing a new property that is integrated into the topic. Of theoretical interest, current models of metaphor understanding embody one or the other process (to varying degrees), but not both.

In general, there have been two distinct approaches to metaphor understanding in cognitive psychology. The *comparison approach* has focused more on analogies than on metaphors but generally assumes that both are interpreted through the use of similar processes. Thus, a metaphor is understood by first aligning or putting into correspondence representational structure of the topic with that of the vehicle via a comparison process. These models include the “structure mapping engine” (Falkenhainer, et al., 1989; Gentner, 1983, 1989) and the “analogical constraint mapping engine” (Holyoak & Thagard, 1989). Usually, knowledge of the vehicle is then transferred to the topic. Figure 3 schematically depicts how the metaphor “heat flow is water flow” would be under-

stood according to these models. It shows (partial) representations of the topic and vehicle domains, the mapping between them (indicated by the approximately horizontal lines connecting elements of the topic to those of the vehicle), and the new knowledge about the topic which is carried over from the vehicle (boldfaced in the figure). In this example, I assume that a person hearing the metaphor “heat flow is water flow” does not know what causes water flow and learns the cause as a consequence of understanding the metaphor. In particular, the person concludes that heat flow is water flow because just as the greater water pressure in the beaker than in the vial causes water to flow from the larger beaker through the pipe to the smaller beaker, so does the greater temperature of the coffee than that of the ice cube cause heat to flow from the coffee through the bar to the ice cube. In this case, the person has acquired new knowledge about what causes heat flow.

As Holyoak and Thagard (1989) have noted, finding correspondences between the topic and vehicle is in principle a computationally intractable problem. If the topic and vehicle are each represented by m predicates and n objects and the mapping is one to one, then there are $m!n!$ possible mappings. This number grows exponentially with a linear increase in m and n . Therefore,

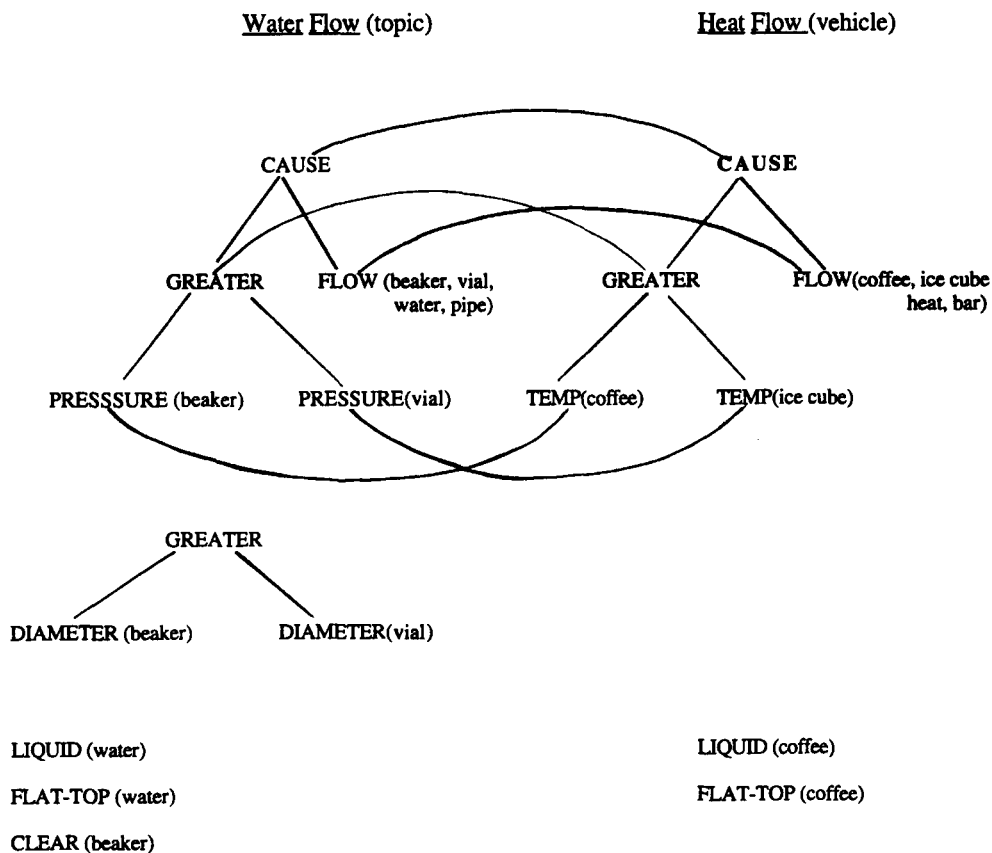


Figure 3. The comparison approach to understanding “heat flow is water flow.”

much emphasis in the comparison approach has been placed on determining the psychological factors that constrain the alignment process. These factors include structural consistency, semantic similarity between predicates of the domains, and pragmatic factors such as the goal of the understander (Holyoak & Thagard, 1989). In brief, structural consistency specifies that if a relation in the vehicle corresponds to a relation in the topic, the constituent predicates and arguments of the paired relations should also correspond. Another factor (related to structural consistency) is called *systematicity* (Gentner, 1983, 1989). In brief, systematicity states that correspondences between higher order relations (i.e., relations that take relations as arguments) should constrain first-order relations (i.e., those that take objects as arguments), which in turn constrain correspondences between objects. More recent work has focused on developing a comparison process that is sensitive to working memory limitations (Hummel & Holyoak, in press).

Given that comparison is involved in the interpretation of noun–noun combinations and metaphors, it is crucial to determine the psychological constraints on this process. However in the comparison approach, there has been considerably less emphasis on constructing new properties in the topic on the basis of properties of the vehicle. In fact, the typical view is that knowledge is transferred from one domain to another by a copy and addition process (Falkenhainer et al., 1989; Gentner, 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1996). For example, Gentner (1989) states that “the learner simply carries across a known system of predicates from the base to the target” (p. 201).

Holyoak and Thagard (1989) make a similar claim:

the general form of analogical transfer is to find correspondences among elements of the source and of the target, and then construct candidate inferences about the target by *essentially copying over propositions from the source* after substituting the appropriate corresponding elements from the target domain. (p. 303; italics added)

This view is exemplified in the metaphor depicted in Figure 3. Understanding the metaphor involves the copying over of the CAUSE predicate from the vehicle to the topic and connecting this relation to its arguments. As another example taken from this approach, understanding the analogy “the atom is like the solar system” is assumed to involve copying the predicate REVOLVE from *solar system* to *atom* and replacing arguments (planets, sun) in the *solar system* domain with corresponding arguments (electrons, nucleus) in the *atom* domain.

However, my analysis suggests that this copy and addition process is generally not used to understand metaphors or analogies. Furthermore, metaphors and analogies prototypically involve comparisons of *conceptually different* domains, with different predicates applying to each domain. Metaphors, therefore, cannot generally be understood by straightforward transfer of predicates from one domain to another, because the predicates in one domain will often not apply in the other domain. Even

when they do apply, they may not apply in the same way. For example, once one has understood “the atom is like the solar system,” one’s sense in which electrons REVOLVE around a nucleus may be different from one’s sense in which planets REVOLVE around the sun (the frequency and size of revolutions are different in these domains). As previously noted, people are sensitive to such differences in conceptual combination (Wisniewski, 1996b).

A second view of metaphor understanding is the interactive property attribution approach that I have described above (Glucksberg & Keysar, 1990; Glucksberg, Manfredi, & McGlone, 1997; Glucksberg, McGlone, & Manfredi, in press; McGlone, 1996). The strengths and limitations of this view are the reverse of those of the comparison approach: it acknowledges the importance of construction but does not specify a role for comparison. For example, in earlier work on which this approach is based, Glucksberg, Gildea, and Bookin (1982) describe how the same properties of the vehicle *butcher* are differentially instantiated in the topics *surgeon* and *pianist*:

The statement *X is a butcher* can always be taken to mean that *X* is negatively evaluated, and that *X* is grossly and characteristically incompetent as well. The particular way that *X*’s incompetence is instantiated will depend on who or what *X* might be. If *X* is a surgeon, the incompetence takes the form of botched operations, with bleeding, disfigurement, and death the likely consequences. If a pianist is a butcher, then the competence is not merely the forgetting of certain parts of the piano pieces or the lack in the repertoire, but that the music is plowed through insensitively, too loudly, without any hint of subtlety or beauty. (p. 95)

In these examples, the authors imply that the vehicle provides information that acts as a source for the construction of new properties in topic. Metaphors are not understood by copying over predicates and replacing arguments.

The property attribution approach does not specify a role for comparison in metaphor understanding. In fact, advocates of this approach have often criticized the comparison view (Glucksberg et al., 1997; Glucksberg, McGlone, & Manfredi, in press; McGlone, 1996). However, a comparison process is important for two reasons. First, the comparison process finds commonalities and differences between the topic and vehicle that will form the basis of the interpretation. Second, the comparison process finds correspondences between the topic and vehicle that indicate where new properties are to be integrated into the topic.

To illustrate the need for both comparison and construction, consider the metaphor *that dinner is a roller coaster*, which subjects defined in the study noted above. Most subjects interpreted this phrase as “a dinner with courses that alternate between tasting good and bad.” (The corresponding noun–noun combinations also yielded these meanings.)

How is this metaphor understood? It appears to be derived from knowledge in *roller coaster* that describes an event divided into repeating subevents (“going up and going down” or “building up excitement and releasing

excitement"). These subevents are characterized by dimensions that alternate between extreme values (a height dimension in which a person is relatively high to the ground or low to the ground, an emotion dimension in which a person feels either very excited or very relieved). This knowledge guides the construction of an interpretation. However, a comparison process is required to determine the aspects of *dinner* that are relevant to deriving the interpretation. This process identifies several important correspondences between the domains. Like *roller coaster*, *dinner* can refer to an event that consists of repeating subevents ("eating a course"). Like the subevents in *roller coaster*, the subevent in *dinner* has a dimension that can be characterized by extreme values: a course can taste good or bad. In identifying these correspondences, an important difference is found between these domains: whereas roller coasters have a dimension with extreme values that *alternate*, *dinner* has a dimension that can take on extreme values that *do not typically alternate*. With this difference as a guide, the metaphor is understood by modifying *dinner* so that it refers to a series of courses that alternate between tasting good and bad. This interpretation is based on both commonalities and a difference between the domains that were found by the comparison process.

Of further importance, the mechanism that creates new knowledge in the topic is not a copy and addition process. For example, prior to one's understanding that *dinner is a roller coaster*, *dinner* referred to a series of courses with their taste unspecified. However, now that the metaphor is understood, *course* must be modified to *explicitly* indicate that it tasted good or tasted bad. Clearly, this new knowledge is not copied from the vehicle to the topic. Furthermore, the sense of ALTERNATE as applied to *dinner* is different from its sense in *roller coaster*. For example, the frequency with which courses alternate from tasting good to tasting bad is slower than the frequency of going up and down in a roller coaster. Thus, ALTERNATE is not simply copied from *roller coaster* to *dinner*. Rather the sense of ALTERNATE in *dinner* is sensitive to constraints specified by *dinner* while still preserving aspects of its meaning in *roller coaster*.

The interpretation of this metaphor also highlights another important characteristic of metaphor understanding: because metaphors typically involve conceptually different domains, it follows that the comparison process must often align nonidentical predicates. In this example, "a course that *tastes good*" might be aligned with "a roller coaster that is *high* off the ground," and "a course that *tastes bad*" might be aligned with "a roller coaster that is *low* off the ground." (Alternatively, the taste predicates might be aligned with predicates describing emotions.)

There may be a variety of factors that influence the alignment of nonidentical predicates. First, people may align nonidentical predicates by recognizing common semantic components between them. For example, people often interpret *snake vase* as "a taller, thin vase" (Wisniewski, 1994). In this case, they apparently align the horizontal length of *snake* with the vertical height of

vase, perhaps because *length* specifies the principal axis of *snake* just as *height* specifies the principal axis of *vase* (see Marr, 1982, for evidence that people are sensitive to this property). Second, certain alignments may be hard-wired into the cognitive system, especially those involving different sensory modalities. For example, loudness and brightness are both coded by temporal firing rates of neurons, and through this common code people may align these different dimensions (Marks & Bornstein, 1987). Third, structural isomorphism between domains may lead to the alignment of nonidentical predicates (Holyoak & Thagard, 1989). Finally, different dimensions may become alignable because their values are correlated (Lakoff & Johnson, 1980; Marks & Bornstein, 1987). For example, Lakoff and Johnson speculate that *happy* aligns with the spatial orientation *up* and *sad* aligns with *down* because "drooping posture typically goes along with sadness and depression, erect posture with a positive emotional state" (p. 15).

CONCLUSION AND SUMMARY

Providing a processing (algorithmic) account of how people combine concepts is a difficult task. I hope, however, that in the present paper I have taken some important steps toward achieving this objective. As a first step, I outlined the goals of conceptual combination and described the important phenomena that result when people combine concepts. In the latter regard, I identified three conceptually distinct kinds of combinations that cover the combinatorial space. I especially highlighted the necessity of conceptual change when concepts combine. This computational level analysis is an essential prerequisite for developing a processing account. It delineates the range of output that the algorithm should produce. In addition, why people combine concepts also constrains how they combine them.

As a second step, I described several processes that operate when people combine concepts. These processes correspond to the conceptually distinct types of output that I identified in the computational level account. One process incorporates the constituents into a plausible scenario in which the constituents play different functional roles. A second process compares the modifier and head noun concepts to find important differences that can form the basis for modifying the head noun. By also finding correspondences between the modifier and head noun, comparison also indicates where these differences should be integrated into the combination. Using these correspondences and differences, a construction process produces conceptual change of a form that goes beyond the addition of relatively unmodified knowledge to a concept. Properties are constructed in a concept rather than added to a concept. The modifier concept provides a property that acts as source of information for the construction of a new property. The head concept provides constraints on how this property is created and instantiated. I also suggested that similar comparison and construction processes are necessary in order to understand nominal metaphors.

However, current approaches to metaphor understanding emphasize one or the other process, not both.

I characterized the construction process at a general level and provided a number of post hoc descriptions of its operation. An important problem to solve consists in specifying in more detail how construction produces conceptual change. This problem is perhaps one of the most important ones to solve in language understanding.

REFERENCES

- BARSALOU, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, **10**, 82-93.
- BARSALOU, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, pp. 1-64). New York: Academic Press.
- BARSALOU, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, M. A. Conway, & P. E. M. Morris (Eds.), *Theories of memory* (pp. 29-101). Hillsdale, NJ: Erlbaum.
- BARSALOU, L. W., & HALE, C. R. (1992). Components of conceptual representation: From feature lists to recursive frames. In I. Van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 97-144). San Diego: Academic Press.
- BRACHMAN, R. J. (1978). *A structural paradigm for representing knowledge* (BBN Report No. 3605). Cambridge, MA: Bolt, Beranck, & Newman.
- 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**, 341-352.
- DOWNING, P. (1977). On the creation and use of English compound nouns. *Language*, **53**, 810-842.
- ELLIOT, R. (1988). *The complete vegetarian cuisine*. New York: Pantheon.
- FALKENHAINER, B., FORBUS, K. D., & GENTNER, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, **41**, 1-63.
- FILLMORE, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Unconstrained in linguistic theory* (pp. 1-88). New York: Holt, Rinehart & Winston.
- GAGNE, C., & SHOBE, E. (1993, April). *The influence of relational information on interpreting nonpredicating combinations*. Paper presented at the 65th annual meeting of the Midwestern Psychological Association, Chicago.
- 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* (pp. 675-698). New York: Academic Press.
- GENTNER, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, **13**, 56-83.
- GENTNER, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, **7**, 155-170.
- GENTNER, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity, analogy, and thought* (pp. 199-241). Cambridge: Cambridge University 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 & Cognitive Processes*, **7**, 205-230.
- GLUCKSBERG, S., GILDEA, P., & BOOKIN, M. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning & Verbal Behavior*, **21**, 85-98.
- GLUCKSBERG, S., & KEYSAR, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, **97**, 3-18.
- GLUCKSBERG, S., MANFREDI, D., & MCGLONE, M. S. (1997). How metaphors create new categories. In T. B. Ward, S. M. Smith, & J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 327-350). Washington, DC: American Psychological Association.
- GLUCKSBERG, S., MCGLONE, M. S., & MANFREDI, D. (in press). Property attribution in metaphor comprehension. *Journal of Memory & Language*.
- GOLDSTONE, R. L. (1994). Similarity, interactive-activation and mapping. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 3-28.
- GRICE, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol 3. Speech acts* (pp. 41-58). New York: Academic Press.
- HAMPTON, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, **15**, 55-71.
- HEIT, E., & BARSALOU, L. W. (1996). The instantiation principle in natural categories. *Memory*, **4**, 413-451.
- HOLYOAK, K. J., & THAGARD, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, **13**, 295-355.
- HUMMEL, J. E., & HOLYOAK, K. J. (1996). LISA: A computational model of analogical inference and schema induction. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- HUMMEL, J. E., & HOLYOAK, K. J. (in press). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*.
- LAKOFF, G., & JOHNSON, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- LEVI, J. N. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.
- LEVIN, B. (1993). *English verb classes and alterations: A preliminary investigation*. Chicago: University of Chicago Press.
- MARKMAN, A. B., & GENTNER, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory & Language*, **32**, 517-535.
- MARKMAN, A. B., & GENTNER, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology*, **23**, 431-467.
- MARKMAN, A. B., & WISNIEWSKI, E. J. (1997). Similar and different: The differentiation of basic level categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 54-70.
- MARKS, L. E., & BORNSTEIN, M. H. (1987). Sensory similarities: Classes, characteristics, and cognitive consequences. In R. E. Haskel (Ed.), *Cognition and symbolic structures: The psychology of metaphoric transformation* (pp. 49-65). Norwood, NJ: Ablex.
- MARR, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- MARTIN, J. D., & BILLMAN, D. O. (1994). Acquiring and combining overlapping concepts. *Machine Learning*, **16**, 121-155.
- MCGLONE, M. S. (1996). Conceptual metaphors and figurative language interpretation: Food for thought? *Journal of Memory & Language*, **35**, 544-565.
- MEDIN, D. L., & SHOBE, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, **20**, 158-190.
- MINSKY, M. (1975). A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 211-277). New York: McGraw-Hill.
- 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 & Language*, **29**, 259-288.
- NISBETT, R. E., & ROSS, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- NUNBERG, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics & Philosophy*, **3**, 143-184.
- PALMER, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259-303). Hillsdale, NJ: Erlbaum.
- PINKER, S., & PRINCE, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, **28**, 73-193.

- RUMELHART, D. E. (1980). Schemata: The building blocks of cognition. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 33-58). Hillsdale, NJ: Erlbaum.
- RUMELHART D. E., & MCCLELLAND, J. L. (1986). On learning past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 216-271). Cambridge, MA: MIT Press, Bradford Books.
- SCHANK, R. (1972). *Conceptual information processing*. New York: Elsevier.
- SCHYNS, P. G., & OLIVA, A. (1994). From blobs to boundary edges: Evidence for a time- and spatial-scale-dependent scene recognition. *Psychological Sciences*, *5*, 195-200.
- SHOEN, E. J. (1993). Comprehending nonpredicating conceptual combinations. In G. Nakamura, R. Taraban, & D. Medin (Eds.), *The psychology of learning and motivation* (Vol. 20, pp. 391-401). San Diego: Academic Press.
- SHOEN, E. J., & GAGNE, C. L. (1997). Thematic relations and the creation of combined concepts. In T. B. Ward, S. M. Smith, & J. Vaid (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.
- THAGARD, P. (1984). Conceptual combination and scientific discovery. In P. Asquith & P. Kitcher (Eds.), *PSA: Proceedings* (Vol. 1). East Lansing, MI: Philosophy of Science Association.
- TVERSKY, A. (1977). Features of similarity. *Psychological Review*, *84*, 327-352.
- URDANG, L., & ABATE, F. R. (1983). *Idioms and phrases index*. Detroit, MI: Book Tower.
- WILSON, D., & SPERBER, D. (1981). On Grice's theory of conversation. In P. Werth (Ed.), *Conversation and discourse*. London: Croom Helm.
- WISNIEWSKI, E. J. (1994). *Interpretations of novel noun-noun combinations* (Technical Report). Northwestern University, Department of Psychology.
- WISNIEWSKI, E. J. (1996a). Construal and similarity in conceptual combination. *Journal of Memory & Language*, *35*, 434-453.
- WISNIEWSKI, E. J. (1996b). *Property instantiation in conceptual combination*. 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., & LOVE, B. (1996). *Properties versus relations in conceptual combination*. Manuscript in preparation.
- WISNIEWSKI, E. J., & MARKMAN, A. B. (1993). The role of structural alignment in conceptual combination. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 1083-1086). Hillsdale, NJ: Erlbaum.
- WISNIEWSKI, E. J., & MEDIN, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, *18*, 221-282.

NOTES

1. Many combinations involve more than two nouns (e.g., *FBI shoe print expert*). Although researchers typically have not examined these combinations, the motivations for their use and processes involved in understanding them are probably similar to those involved in understanding noun-noun combinations. On the other hand, multiple noun combinations introduce other ambiguities into understanding such as determining which noun modifies another. Whereas the leftmost noun typically modifies the rightmost (head) noun in a noun-noun combination, this generalization does not hold for multiple noun-noun combinations. For example, although *shoe* modifies *print* in *FBI shoe print expert*, *FBI* modifies *expert*.
2. When interpreting this novel combination, subjects may have been influenced by knowledge of runcible spoons or sporks, which bear resemblance to the fork spoon depicted in Figure 1a.
3. In the study of metaphor, a number of different terms have been used to refer to the constituents of a metaphor. In particular, for *that X is a Y*, *X* has been called the *topic*, *source*, and *base*, and *Y* the *vehicle* and *target*. I use the more traditional terms *topic* and *vehicle*.
4. I thank an anonymous reviewer for bringing this issue to my attention.

(Manuscript received March 26, 1996;
revision accepted for publication September 12, 1996.)