Luce, R. (1959). *Individual choice behavior.* New York: Wiley.

MacDonald, M. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes, 9,* 157–201.

MacDonald, M., Pearlmutter, N., & Seidenberg, M. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review, 4,* 676–703.

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

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

Rumelhart, D., & Abrahamsen, A. (1973). A model for analogical reasoning. *Cognitive Psychology, 15,* 346–378.

Sadler, D., & Shoben, E. (1993). Context effects on semantic domains as seen in analogy solutions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 128–147.

Seidenberg, M. (1992). Connectionism without tears. In S. Davis (Ed.), *Connectionism: Theory and practice* (pp. 84–122). Cambridge, England: Cambridge University Press.

Shoben, E. (1991). Predicating and nonpredicating combinations. In P. J. Schwanenflugal (Ed.), *The psychology of word meanings* (pp. 117–135). Hillsdale, NJ: Erlbaum.

Shoben, E., Cech, C., & Schwanenflugel, P. (1983). The role of subtractions and comparisons in comparative judgments involving numerical reference points. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 226–241.

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

Trueswell, J., Tanenhaus, M., & Kello, C. (1993). Verb specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 528–553.

Whaley, C. (1978). Word–nonword classification time. *Journal of Verbal Learning and Verbal Behavior, 17,* 143–154.

Wisniewski, E., & 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, The Netherlands: Elsevier Science.

3

Conceptual Combination: Possibilities and Esthetics

Edward J. Wisniewski

Much of language understanding involves combining concepts into 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] For example, one might interpret *frog bowl* as a bowl for housing pet frogs. People create novel combinations in order to specify referents of discourse contexts and to extend the vocabulary of their language (Downing, 1977; Gerrig & Murphy, 1992). Besides their role in language, the study of novel combinations can inform theories of concepts (Hampton, 1987; Markman & Wisniewski, in press; Medin & Shoben, 1988). For example, studies of conceptual combination have identified

I thank Doug Medin and Tom Ward for trenchant comments on a previous draft and Bob Dylan for providing some of the inspiration.

[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 the 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.*

ways that 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, in press).

In general, the creation of new concepts by combining existing ones is a powerful and common way of expanding knowledge. For example, a recent newspaper article used the phrase *car boat* to refer to a new kind of boat that is also a car. Such a phrase provides access to a wealth of inferences that one can reasonably draw about this novel entity that do not necessarily need to be empirically derived. There is also a sense of creativity in bringing together disparate concepts to form new concepts that people readily understand (e.g., as in the use of *boomerang flu* to refer to a flu that goes away and comes back, or *bait car* to refer to a car that is used to catch carjackers). However, in psychology, formal models of concept learning primarily acquire new concepts by applying empirical methods to many examples of a category (Wisniewski, 1995). There has been considerably less emphasis on how existing concepts are combined and modified to produce new ones.

There are at least three levels at which one can understand conceptual combination. Marr (1982) referred 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, this volume; 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 chapter, 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 carrying out and representing this change. Constrained by this computational level account, I then provide an algorithmic level sketch of how concepts combine and change. This sketch does not constitute a full-fledged model, and more empirical work is needed to flesh out the details. However, I believe that my account is an advance over current views and that it captures essential ingredients that must go into a more comprehensive model. In particular, this account highlights the importance of two additional processes (comparison and construction) that have not been emphasized in the research on conceptual combination. Finally, I suggest that the interpretation of nominal metaphors also involves comparison and construction processes. However, psychological models of metaphor understanding emphasize one or the other of these processes, but not both. A central theme of this chapter is that conceptual change does not usually involve the straightforward transfer of knowledge from one concept or domain to another. Rather, knowledge in one concept guides the construction of new knowledge in another concept. In this sense, conceptual change is creative in going beyond simple addition. This view is the rule rather than the exception. A fundamental challenge to models of novel language understanding is specifying processes that use knowledge in one domain to construct knowledge in another domain.

A COMPUTATIONAL LEVEL ACCOUNT: WHY PEOPLE COMBINE CONCEPTS

Conceptual combination occurs in a communicative context and serves at least three goals. 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,

in contrast to fish bowls and cereal bowls, a frog bowl is typically used to contain frogs and may also have an island-like glass structure in its middle for this purpose. Second, combinations are often used to convey information in a concise and efficient way. For example, *football parking* designates an area for parking one's car while attending a football game. Even though this phrase is somewhat elliptical readers generally understand what it means. Third, combinations 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 establish cohesion in discourse contexts.

In achieving these goals, the speaker and 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 (rightmost) head noun: Frog bowls are different from other types of bowls. Second, the source of this difference comes from the (leftmost) modifier noun: In contrast to other bowls, frog bowls house frogs. Third, despite this difference, the referent of the combination shares important commonalities with the head noun category: Although frog bowls are different from other bowls, they are still containers and shaped like other bowls. In general, these constraints are consistent with basic goals of communication such as being informative and relevant, and avoiding ambiguity (Grice, 1975; Wilson & Sperber, 1981).

It would appear that speakers and listeners adhere to these constraints in producing and interpreting novel combinations. Imagine how bizarre a novel phrase that did not follow these constraints would sound. For example, suppose that *skunk squirrel* referred to any squirrel (violating the first constraint) or referred to a special kind of squirrel whose specialness had nothing to do with skunks (violating the second constraint), or did not refer to squirrels in any sort of way (violating the third constraint).

Why do people produce novel combinations that obey these constraints? Typically, people need to refer to changes in their environment: a newly discovered disease, a newly created product, an important change to something old. In general, people create change by modifying the familiar in small steps (leaving much constant), and they attempt to understand and refer to change by relating it to the familiar (Murphy, this volume; Ward, 1994; Weber & Dixon, 1989). Using a novel combination captures both constancy and change. For example, recently in a grocery store, *seafood sausage* was used to refer to patties of ground up seafood. In this case, people can use the concepts *seafood* and *sausage* and their interaction to infer properties of sausage that are likely to remain constant (shape, size, edibleness, method of cooking) and those that are likely to change (composition, calorie level, cooking time).

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 a model. Furthermore, knowing what people compute constrains a model's processing assumptions. For example, a psychologically plausible model should not compute things that people do not compute (cf. Pinker & Prince's, 1988, critique of Rumelhart & 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 individuals' interpretations of novel combinations. In many studies, I have asked research participants to assume that they have heard a novel combination in a conversation and to describe its most plausible meaning (Markman & Wisniewski, in press; Wisniewski, 1996a; Wisniewski & Gentner, 1991; Wisniewski & Love, 1996; Wisniewsk & 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 arbitrarily or randomly pairing their constituents, which generally produces novel combinations. Although the combinations are novel, their constituents name common, familiar things including human-made artifacts, plants or their parts, substances, abstract terms, and events. In these studies, participants typically proceed through the task at their own pace, writing down meanings for 15 to 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 to 20 interpretations per combination (Wisniewski, 1995).

Kinds of Combinations

Analyses of these interpretations suggest that people combine concepts in three basic ways (Wisniewski, 1996a). Two of these ways, relation linking and property interpretation, are very common among novel combinations. *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—for example, if one pictures a robin snake, interpreted by relation linking as "a snake that eats a robin," one sees a snake eating a robin. However, in picturing a robin snake, interpreted by property construction as "a snake with a red underbelly," one sees only a snake (whose underbelly has a color similar to that of a robin). There are also important conceptual distinctions between hybrids and relation-linking and property interpretation types—for example, a *robin canary*, interpreted with a property as a "canary with a red breast" or by relation linking as a "canary that preys on robins" is a canary, but a *robin canary* that is a cross between the two is not just a canary.

In these studies, I have focused on novel combinations in contrast to familiar ones (e.g., shoe box, bookstore) because it is more difficult to infer interpretation processes from the latter. Meanings of familiar combinations have often been modified or augmented after their initial creation (Levi, 1978). For example, eggplant once referred to a small, white, egg-shaped vegetable but now typically refers to a larger, purple, oblong vegetable (Elliot, 1988). Nevertheless, it is easy to find familar examples of these combinations, and linguistic classifications of familiar combinations are generally consistent with the claim of three basic types (see Wisniewski, 1996a, in press; Wisniewski & Love, 1996, for many familiar examples of these types).

Similarity Effects

The similarity between the modifier and head concepts strongly affects the kind of interpretation that is produced. In particular, a robust finding is that combinations with highly similar constituents (e.g., *robin canary, magazine newspaper*) almost always yield property interpretations or hybrids. My colleagues and I have obtained this result in many studies and have never failed to replicate this finding (Markman & Wisniewski, in press; Wisniewski, 1996a; Wisniewski & Love, 1996; Wisniewski & Markman, 1993). In one study (Wisniewski, 1996a; Experiment 2), 92% of 320 interpretations of 32 highly similar combinations were interpreted using these two strategies. For example, some property interpretations included "horse that produces milk," for *cow horse* and "newspaper with glossy ads" for *magazine newspaper.*

When the constituents of a combination are not similar, people are more likely to interpret the combination by relation linking (e.g., *apartment piano* is often interpreted as "a piano found in an apartment " and *cow cabbage* is often interpreted as "cabbage eaten by cows"). Also, when

people do interpret dissimilar combinations using properties, their meanings are more likely to involve emergent properties (i.e., properties not represented in either constituent, Wilkenfeld, 1995; Wisniewski, 1991a, 1991b). People can be quite creative in this regard. For example, one participant interpreted *ant canary* as "a canary with two feathers on its head which look like antennae." As another example, individuals sometimes interpret *rake pencil* as "a pencil with multiple lead points."

Referential Scope

The use of a novel combination indicates one or more new referents. In many combinations, these referents are subsets of the categories typically named by the constituents of a combination. 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 one would typically name with the terms *book* and *jar*. In many respects, this assumption 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.

Nevertheless, people are more creative and flexible in their interpretation of a noun's referent. The potential *referential scope* of a noun is broader. First, as implied by property interpretations, 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, a common interpretation of *skunk squirrel* is "a squirrel with black and white stripes." In this case, the referent of *skunk* is one of its properties and not the things that people call skunks. Second, people can treat a noun as referring to a representation of a thing rather than the thing itself. For example, a not infrequent interpretation of *moose pencil* is "a pencil with an eraser that looked like a moose." This interpretation refers to a representation of *moose* rather than to a moose as in "pencil stepped on by a moose." Finally, 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 *tiger chair* as "a chair made of tiger skin."

This flexibility is not a rare phenomenon and occurs in many familiar combinations (see Wisniewski, 1996a, for examples). Furthermore, this flexibility is evident in a major theory of metaphor understanding that assumes that nouns refer both to things and to prototypical properties of those things (Glucksberg & Keysar, 1990; Glucksberg, Manfredi, & McGlone, this volume). For example, in the literal statement "my job at the jail," *jail* refers to a jail, but in the nominal methaphor "my job is a jail," it refers to properties exemplified by jails such as "confining."

Adopting this broader referential scope for a noun helps achieve the communicative goals of conceptual combination: 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 efficiently 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." The use of these concise labels would not be possible if people could not flexibly expand the referential scope of a noun.

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 is limited to a property, representation, or thematic associate of the typical referent of the noun.

Conceptual Change

Conceptual combination is a process that involves conceptual change or rerepresentation. For example, consider *zebra clam*, interpreted as a "clam with stripes." Combining *zebra* and *clam* produces a new concept that is a modification of *clam*. What is the nature of that change? Of importance, the meaning of the new concept *zebra clam* is not "a copy of zebra's stripes," which has been added to the meaning of *clam*. Rather the stripes of a zebra clam are shorter and thinner than those of a zebra, and they run along

the outside of a clam shell and not a zebra body. In a sense, 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 also might be shorter and thinner than those of a zebra to be consistent with the size of a clam. This example illustrates that combining concepts can go beyond the simple copying of knowledge over from one concept to another. Instead, conceptual combination is creative in that new knowledge emerges out of interactions between different representations.

Of importance, this construction process is quite common. Many novel combinations are interpreted in this manner (Wisniewski, 1996b). It is also easy to find familiar combinations that refer to entities having a property resembling but not identical to that of the modifier. For example, the shapes of *butterfly chair*, *butterfly valve*, and *butterfly clip* are variations on the shape of a butterfly. *Zebra mussel*, *zebra fish*, *zebra wood*, and *zebra finch* refer to objects with different variants of a zebra's stripes. *Accordion bag* and *accordian pleat* refer to objects with different instantiations of an accordion's folds (see Wisniewski, in press, for other examples). Presumably, when people first coined these phrases they had to determine that a property of the "object to be named" was related to one of another object. For example, in wanting to create a new name for a mussel with stripes, they had to determine that the stripes were related to those of a zebra (as opposed to those of the American flag, a barber pole, or a tiger). To see such a relation, people may have had to invoke the construction process.

Summary

An extensive analysis of the meanings of novel combinations suggests that there are three basic kinds of interpretations. The similarity between the constituents strongly affects which kind of interpretation is selected: Combinations with highly similar constituents are almost always interpreted either by property construction or by hybridization. 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, 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.

ALGORITHMIC LEVEL: CURRENT MODELS OF HOW PEOPLE COMBINE CONCEPTS

There are two processing accounts of how people combine concepts, which are fairly broad in scope: the thematic relations approach and the schema approach.[2] Both accounts assume that the meaning of a combination involves a relation between the referents of the constituents. In the *thematic relations* view, nouns are combined by determining an abstract relation that holds between the nouns (Coolen, van Jaarsveld, & Schreuder, 1991; Shoben, 1993; Shoben & Gagne, this volume). These approaches assume that there are a relatively small set of such relations (perhaps one or two dozen). In evaluating this view, researchers have typically drawn on the relations that linguists have used to classify familiar combinations (e.g., Downing, 1977; Levi, 1978). For example, Levi (1978) 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, etc.).

The most developed view of the thematic relations approach is the work by Shoben (1993) and his colleagues (see Shoben & Gagne, this volume). They assume that a noun's combinatorial history influences inter-

[2] I have excluded from this discussion models that 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). I have also excluded a model of how certain adjectives are combined with nouns (Smith, Osherson, Rips, & Keane, 1988). In general, it is not possible to straightforwardly extend these models to account for other kinds of combinations (see Murphy, 1988; Wisniewski & Gentner, 1991, for arguments).

pretation of a novel phrase involving that noun. That is, people use the distributional knowledge of how nouns have been 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 been previously combined with other nouns in a similar manner. A second assumption is that the modifier's combinatorial history has more influence on interpretation (because it precedes the head noun). In a series of studies, Shoben and Gagne (this volume) show that the time to judge the sensibility of a novel combination is a function of the relative frequency of abstract 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 people have about a place, event, or object which has been acquired from 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. Schema-based models link concepts by relations through a process of slot filling (Brachman, 1978; Cohen & Murphy, 1984; Murphy, 1988; Shoben, 1993). 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. According to this account, 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. 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 po-

tential filler. For example, the filler of the *eats* slot of *snake* would have to be edible, whereas the filler of the *contains* slot of *box* would have to be smaller than a box. 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 through a relation, there are two important differences between the approaches. First, they emphasize different levels of abstraction in characterizing the meanings of combinations. Whereas the abstract relations view postulates that a relatively small set of very general relations exhausts the ways that 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 in determining meanings.

There are several reasons to believe that the schema approach better accounts for how people combine concepts. First, the use of more specific slots rather than abstract relations captures distinctions between meanings to which people are clearly sensitive. For example, in the abstract 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 people's 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 distinctions seem to be part of the basic, core meaning of such novel combinations that arise from the specific knowledge captured 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). 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 cannot be stirred.

Second, the thematic relations view overly emphasizes the role of a noun's combinatorial history in determining meaning. On the one hand, there is some evidence that how a noun combines with another is influenced by how it has been previously combined (Shoben & Gagne, this volume). 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 cannot be the frequency of these past meanings per se that determines whether people will interpret a novel *X box* in the same way. Some combinations are not interpreted in this manner (e.g., *skyscraper box*). Rather, as emphasized in the schema approach, interpretation involves figuring out whether a meaning is plausible (e.g., can *X* fit into a box?). 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 (as most things can be enclosed). The function of a noun's combinatorial history is probably to suggest candidate meanings that previously have been useful (and thus narrow down the search space), but most of the work in interpretation is determining a plausible meaning and creating a new representation.

In summary, two general proposals have been offered to explain how people combine concepts. The schema approach appears to be a more plausible psychological model. However, the computational level account suggests that both the thematic relations and schema views do not go far enough in explaining how people combine concepts. There are three major areas where such approaches are incomplete. First, both approaches focus on processes that derive relational interpretations of combinations and have not addressed the derivation of property and hybrid interpretations. Yet, property interpretations are common for many novel combinations (Wisniewski, 1996a) and not infrequent among familiar combinations (Wisniewski & Love, 1996). As I suggest later, property and hybrid interpretations require processes that differ from those that determine re-

lations between the modifier and head concepts. Second, current approaches assume that the referential scope of a noun is fairly narrow: It is taken to refer to the object or thing typically named by that noun. In contrast, a noun can refer to a property, representation, or thematic associate of the noun's typical referent.

Third, the models capture relatively small changes in representation that result when concepts combine. In the thematic relations approach, a new concept is created by linking one constituent to another through a relation. Being linked to another constituent is the only sense in which a constituent's representation undergoes modification. In schema approaches, a new concept is created by replacing the typical value of a slot with a new value (Cohen & Murphy, 1984; Murphy, 1988; Smith et al., 1988). One schema approach, the concept specialization model, goes further than slot value replacement by postulating an additional stage of processing that also involves conceptual change (Cohen & Murphy, 1984; Murphy, 1988). After slot value replacement, people use their world knowledge to elaborate or "fix up" this concept by adding attributes to a combination. For example, in interpreting *apartment dog* as "a dog that lives in an apartment," people may recall actual dogs that they know and use the memories to modify the combination by adding properties such as "yappy and neurotic" (Murphy, 1988; p. 540). Thus, this model's notion of conceptual change not only involves slot value replacement but also adds new attributes to the combination. However, as I have suggested, conceptual change does not generally result from adding attributes.

UPPING THE ANTE: TOWARD A COMPREHENSIVE MODEL

The task of algorithmically specifying how people combine concepts is a difficult one. In this section, I describe an approach to conceptual combination that attempts to capture a broader range of psychological phenomena. My description does not constitute a full-fledged model. However, it identifies what appear to be necessary ingredients that must go into constructing a model.

Representational Assumptions

On the basis of the phenomena that I have presented, it is clear that a model has no hope of succeeding if it uses simple representations of nouns. Much research suggests that concepts have complex representational structure (Barsalou, 1993; Barsalou & Hale, 1992; Cohen & Murphy, 1984; Gentner, 1989; Murphy, 1988; Ortony, 1979; Palmer, 1978; Wisniewski & Medin, 1994). In keeping with the strengths of the schema approach, I assume that a noun is a schema that consists of slots and fillers. 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 (see also Barsalou & Hale, 1992). A noun schema represents a variety of knowledge about its referent, including its typical properties (e.g., an object would have size, shape, color, parts, and so on) as well as the everyday interactions that the referent has with other entitites. I refer to these interactions as *scenarios*. In the latter case, for example, there could be a *function* slot in *soap* with the filler *cleaning*. This filler would be a scenario with the possible roles or slots of *agent*, *instrument*, and *recipient* and the typical fillers *person*, *soap*, and *solid object*, respectively. In this sense, the scenario representation of such everyday interactions is analogous to verb representations that describe actions, events, or states (cf. Fillmore, 1968; Levin, 1993; Schank, 1972). Finally, there are typically interconnections between components of a schema that capture dependencies between slots and fillers. For example, a complex property of *hammer*, such as "for pounding in nails," depends on the hammer having a handle, being grasped and lifted in a certain way, having a blunt, solid end, and so on.

At a general level, there are two basic steps in combining concepts, although these steps merge with each other and there is no clear border between them. First, people must determine what knowledge will initiate the formation of the new combination. That is, they must figure out which property of the modifier to integrate into the head noun or which relation will link the modifier and head concepts. People must then create a new representation. For example, the common interpretation of *fork spoon* as "a spoon with prongs" involves identifying the property "has prongs"

of the modifier *fork* as the candidate property for guiding the construction of a new property in *fork spoon*. On the basis of participants' interpretations, the new representation typically corresponds to a spoon with prongs on the opposite end from the spoon's bowl or a spoon with short prongs emanating from the tip of the bowl (as in a spork or runcible spoon). Below, I suggest that different processes are involved in property interpretation and relation linking.

Identifying Potential Properties

Deriving a property interpretation requires finding a *difference* between the modifier and head concepts that forms the basis of the property that will be constructed in the head concept. This claim follows from a computational level goal of conceptual combination: the need to create a new category differing in some significant way from the category referred to by the head noun and whose difference is related in some way to the modifier.

The need to find a difference raises three logical questions: How do people determine this difference, why does this difference (and not some other one) form the basis for the new combination, and where in the head concept is the difference integrated? For example, to interpret *zebra horse* as "a horse with stripes," people must determine that zebras have stripes but horses do not and therefore that this difference could form the basis of the interpretation. In general, people are not sensitive to the absence of information, and thus concepts do not explicitly represent the absence of properties (e.g., see Nisbett & Ross, 1980, for evidence). People also tend to select this particular difference for the interpretation of *zebra horse* and not other possible differences (e.g., lives versus does not live in Africa). Finally, people believe that a zebra horse is one with stripes that run down the sides of the zebra horse's body and neck just as they do in a zebra. They do not, for example, think that the stripes of a zebra horse are on its tail. Thus, the difference is integrated into certain parts of the representation and not others.

Finding Differences and Integrating Them in Concepts

One can begin to address these questions by assuming that people use a comparison process to align or put into correspondence the structure of

the modifier and head concepts (Wisniewski, 1996a, 1996b; Wisniewski & Gentner, 1991; Wisniewski & Markman, 1993). This assumption was inspired by some models of metaphor and analogy understanding that have emphasized the importance of comparison (Gentner, 1983, 1989; Holyoak & Thagard, 1989). By aligning the structure that is common to the head and modifier concepts, people find differences that then form the basis of an interpretation (cf. Markman & Gentner, 1993). The idea is that any pair of concepts will have both commonalities and differences. Finding commonalities leads to finding differences because commonalities are interconnected or related to differences. To illustrate, consider again *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, vertically 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 versus not having stripes).

Besides determining differences between concepts, the comparison process suggests where a property can be integrated into the representation of 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 versus not having stripes), the comparison process has determined where the stripes can be incorporated into *zebra horse*. (If a zebra only had stripes on its tail, the comparison process would suggest that a zebra horse only had stripes on its tail.)

There is much evidence that people use the commonalities between mental representations to find differences. Markman and Gentner (1993) asked research participants to list commonalities and differences of word pairs. Not surprisingly, participants listed more commonalities for pairs of similar words than for pairs of dissimilar words. However, the reverse did not hold: Participants did not list more differences for dissimilar pairs than for similar pairs. Instead, they (somewhat paradoxically) listed more differences for the similar pairs, and importantly, the differences were conceptually related to the commonalities (see also Gentner & Markman,

1994; Markman & Wisniewski, in press). For example, when participants listed "has wheels" as a commonality of car and motorcycle, they also tended to list "has four wheels versus two wheels" as a difference. Of importance, there is evidence that people use a comparison process in combining concepts (Wisniewski, 1996a).

There may be a variety of ways (not necessarily mutually exclusive) in which the comparison process is instantiated (Wisniewski, 1996a, discussed the possibilities). The major point to be made is that such a process is necessary whatever the actual mechanism is that carries it out. For example, I have described the comparison process as if it aligns one mental representation with another to find a difference. However, as an alternative, 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), or a prototypical property of the modifier might be rapidly activated on reading the modifier (cf. Barsalou, 1982; Glucksberg & Keysar, 1990). In these cases, comparison is still necessary to establish where the property is to be integrated into the head concept. For example, even though "having stripes" might be highly accessible on reading *zebra*, people still must determine how it fits into *horse*. It is the correspondences between the bodies and necks of zebras and horses that suggest that the stripes run alongside the neck and body of a zebra 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*, as horses have a very similar shape. The comparison process identifies this commonality and thus prevents it from being used in the interpretation.

Selecting a Difference

In general, there are multiple differences between concepts. At least several factors influence which difference is selected. One factor is the discourse setting that may suggest the relevant difference. Another factor is

a property's cue and category validity. 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. When people derive property interpretations, they tend to select properties that are high in cue and category validity (Wisniewski, in press).

This finding 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. A property with perfect cue and category validity unambiguously identifies the members of the category. Therefore, in cooperative communication, a listener is most likely to think of properties high in cue and category validity in understanding a novel combination, and speakers may construct combinations with these properties in mind.

Of course, cue and category validity cannot be the only factors that determine selection. Appealing solely to these factors predicts interpretations of combinations that people never give. For example, people would probably never literally interpret *fork tennis balls* as "tennis balls with prongs" or as "a package of four tennis balls" even though these interpretations are derived from a property of *fork* that is very high in cue and category validity (i.e., "has four prongs"). Again, the communicative goals of the computational level account constrain selection. The interpretation "tennis balls with prongs" compromises the functionality of *tennis balls* and may result in a referent that is outside the acceptable scope for *tennis balls*. On the flip side, the interpretation "package of four tennis balls" may stretch people's ability to see the connection between the "has four prongs" property of *fork* and the "package of four tennis balls" meaning of *fork tennis balls*. In this case, "has four prongs" is used to construct a new property in *fork tennis balls* that preserves so little of the "has four prongs" that it lies outside the acceptable referential scope of "has four prongs."

These constraints, coupled with the comparison process, provide the beginning of an algorithmic account of property construction. To illustrate, I return to *fork spoon* and provide a more detailed account of how it is interpreted as "a spoon with prongs" (see Wisniewski, in press, for more examples). 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 variations on the comparison process 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, which functionally are 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 (as in a runcible spoon or a *spork*). Alternatively, they could mentally attach the prongs to the top end of the spoon. 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.

Identifying Plausible Relations

In contrast to comparing the modifier and head noun representations in order to find an important difference that forms the basis of a new property, finding a relation between constituents amounts to constructing a plausible scenario involving the constituents. 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).

Clearly, a model of conceptual combination must assign and bind constituents to plausible roles in a scenario. However, analogous to property construction, the exact details of such a role assignment process are not well understood, and they remain an issue for future research. In assigning a constituent to a role, the process must determine the fit between the role and the constituent. Like the previously described schema approach (e.g., Murphy, 1988), fit is based on how well a constituent matches the preconditions of the role. If there is a sufficient match, the constituent is assigned to that role. However, examining how people combine concepts suggests that role fitting can be more interactive than suggested by the schema approach. That is, rather than finding a plausible role for a constituent, the constituent may be modified to plausibly fit a role. For example, a frequent interpretation of *robin termite* is "a termite that eats robin nests" (Wisniewski, 1996a). In this case, people construe *robin* in such a way that it plausibly fits the "what is eaten" role of the *eats* scenario (*robin nest* fits the constraints of what termites eat, whereas *robin* does not). As another example, *elephant box* is sometimes interpreted as "a very large, sturdy box for containing elephants" (Wisniewski, 1996a). In this case, a typical box is modified in such a way that *elephant* can plausibly fit the "what is contained" role of the *contains* scenario.

Plausible Properties Versus Plausible Relations

Up to this point, I have suggested that there are two different processes involved in combining concepts: a process that identifies and constructs properties and a process that links concepts by relations. One issue concerns the relationship between these processes. Do they operate concurrently or is there a preferred strategy for interpreting novel combinations? Some researchers have suggested that identifying and constructing properties is a last resort strategy—one that is used only when the constituents of a combination do not fit into a plausible scenario (Downing, 1977; Shoben & Gagne, this volume; Wisniewski & Gentner, 1991). Although

more research is needed to resolve this issue, recent work calls into question the "scenario first" hypothesis. In particular, which process dominates depends on factors such as the similarity between constituents and whether a process has succeeded in the previous context (Wisniewski & Love, 1996).

THE RELATION BETWEEN CONCEPTUAL COMBINATION AND NOMINAL METAPHORS

Recently, I have argued that property interpretations of noun–noun combinations are similar to interpretations of nominal metaphor interpretations (Wisniewski, in press). Nominal metaphors are language constructions of the form *X is a Y*, in which *X* and *Y* refer to nouns. A nominal metaphor can be rephrased as a noun–noun combination by reversing the order of the nouns and dropping the form of *be* to produce the phrase *YX*. 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. 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."

On the other hand, there are clear differences between noun–noun combinations and nominal metaphors. Noun–noun combinations allow for other kinds of interpretations (i.e., relation linking and hybridization). For example, people do not interpret "that job is a jail" by linking *job* and *jail* through a relation as in "that job is *at* a jail," whereas "that jail job" can be interpreted in this manner. There may be a variety of more subtle differences between these language constructions (see Wisniewski, in press, for details).

Nevertheless, 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 property 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 (see Wisniewski, in press, for more details). 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 using similar processes. Thus, a metaphor is understood by first aligning or putting into correspondence representational structure of the topic with that of the vehicle through a comparison process. These models include the structure mapping engine (SME; Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 1989) and the analogical constraint mapping engine (ACME; Holyoak & Thagard, 1989). Alignment is potentially a computationally intractable problem (Holyoak & Thagard, 1989) so much of the emphasis in the comparison approach has been on determining the psychological factors that constrain this process (see Gentner, 1983, 1989; Holyoak & Thagard, 1989, for details). 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 transfer is a copy and addition process (Falkenhainer et al., 1989; Gentner, 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1996).

However, my analysis of conceptual combination 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. Therefore, metaphors generally cannot 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, on understanding "the atom is like the solar system," people's sense in which electrons revolve around a nucleus may be differ-

ent from their 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 (Glucksberg & Keysar, 1990; Glucksberg et al., this volume; 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) described 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 but rather by a construction process similar to the one I have proposed in conceptual combination.

However, 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., this volume; McGlone, 1996). However, a comparison process is important for two reasons. First, the comparison process uses commonalities between the topic and vehicle to find differences which will form the basis of the

interpretation. Second, the comparison process finds correspondences between the topic and vehicle which indicate where new properties are to be integrated into the topic.

In summary, the interpretation of nominal metaphors may use the same processes as those involved in property interpretations of noun–noun combinations (i.e., comparison and construction). Current models of metaphor understanding have emphasized the importance of either comparison or construction, but not both. A complete model must employ both processes.

CONCLUSION

Providing a processing (algorithmic) account of how people combine concepts is a difficult task. I hope, however, that this chapter has 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. Of importance, I highlighted the necessity of conceptual change when concepts combine. This computational level explanation 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 constituents to find important differences between them. One or more of these differences provide the basis for modifying the head noun constituent. Of importance, I have shown why a comparison process is crucial to combining concepts in this manner. It provides a psychologically plausible way to determine differences and indicates where the

should be integrated into a concept. Of crucial importance, this integration involves 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 a 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 these comparison and property construction processes characterize how nominal metaphors are understood.

Although I have characterized the property construction process at a general level, an important problem to solve is specifying in more detail how this process operates and how it produces conceptual change. This problem is perhaps one of the most important ones to solve in language understanding.

REFERENCES

arsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10,* 82–93.

arsalou, 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.

salou, 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, CA: Academic Press.

hman, R. J. (1978). *A structural paradigm for representing knowledge* (BBN Report No. 3605). Cambridge, MA: Bolt, Beranck, & Newman.

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

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

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

, R. (1988). *The complete vegetarian cuisine.* New York: Pantheon Books.

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.), *Unconstraineds in linguistic theory* (pp. 1–88). New York: Holt, Rinehart & Winston.

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, England: Cambridge University Press.

Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science, 5*(3), 152–158.

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

Glucksberg, S., Gildea, P., & Bookin, H. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior, 21*, 85–98.

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

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics, 3, Speech Acts.* 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 G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Levi, J. N. (1978). *The syntax and semantics of complex nominals.* San Diego, CA: Academic Press.

Levin, B. (1993). *English verb classes and alterations: A preliminary investigation.* Chicago, IL: University of Chicago Press.

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

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

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 and Language, 35*, 544–565.

Medin, D. L., & Shoben, 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.

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.

Ortony, A. (1979). Beyond literal similarity. *Psychological Review, 86*, 161–180.

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, and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 216–271). *Psychological and biological models.* Cambridge, MA: Bradford Books/MIT Press.

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

Shoben, E. J. (1993). Comprehending nonpredicating conceptual combinations. In G. Nakamura, R. Taraban, & D. Medin (Eds.), *The psychology of learning and motivation* (Vol. 20, pp. 391–409). San Diego, CA: Academic Press.

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* (Vol. 1). East Lansing, MI: Philosophy of Science Association.

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

Ward, T. B. (1994). Structured imagination: The role of category structure in exemplar generation. *Cognitive Psychology, 27*, 1–40.

Weber, R. J., & Dixon, S. (1989). Invention and gain analysis. *Cognitive Psychology, 21*, 283–302.

Wilkenfeld, M. J. (1995). Conceptual combination: Does similarity predict emergence? Unpublished master's thesis, Texas A&M University, College Station.

Wilson, D., & Sperber, D. (1981). On Grice's theory of conversation. In P. Werth (Ed.), *Conversation and discourse.* London: Croom Helm.

Wisniewski, E. J. (1991a, May). Modeling conceptual combination. Paper presented at the Conference on Categorization and Category Learning in Humans and Machines, Department of Psychology, Texas Tech University, Lubbock.

Wisniewski, E. J. (1991b, May). *Relations versus properties: The effects of similarity on conceptual combination.* Paper presented at the Symposium on Conceptual Combination, Annual Meeting of the Midwestern Psychological Association, Chicago, IL.

Wisniewski, E. J. (1995). Prior knowledge and functionally related features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 449–468.

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

Wisniewski, E. J. (1996b). *Property instantiation in conceptual combination.* Manuscript submitted for publication.

Wisniewski, E. J. (in press). When concepts combine. *Psychonomic Bulletin & Review.*

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, The Netherlands: Elsevier Science.

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, *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 1083–1086). Boulder, CO: Erlbaum.

Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science, 18(2)*, 221–282.