Understanding Word and Sentence
G.B. Simpson (Editor)

Chapter 10

On the Combinatorial Semantics of Noun Pairs: Minor and Major Adjustments to Meaning

Edward J. Wisniewski
University of Michigan
Ann Arbor, Michigan U.S.A.

Dedre Gentner
University of Illinois
Champaign, Illinois U.S.A.

The process of conceptual combination involves accessing two or more concepts and determining how they fit together to form a new concept. In a sense, conceptual combination is very broad in scope, involved in many situations in natural language understanding. For example, understanding a story probably includes the combining of concepts of the individual sentences. Understanding a sentence, in turn, probably includes combining the meanings of noun, verb, and prepositional phrases. To understand a noun, verb, or prepositional phrase, we combine the meanings of individual words. In this chapter, we will focus on how people combine concepts when they attempt to understand complex noun phrases (i.e., noun phrases other than those consisting of a noun or a determiner and a noun). For example, to understand a phrase like "elephant tie," one might combine the concepts *elephant* and *tie* in such a way to mean, "a tie worn by circus elephants" or "a tie with a picture of an elephant on it."[1] These are possible interpretations of the phrase "elephant tie." Recently, there has been a fair amount of psychological research on this kind of conceptual combination (e.g., Osherson & Smith, 1982; Smith & Osherson, 1984; Medin & Shoben, 1988; Murphy, 1990). Several models have been proposed to account for this process (Hampton, 1987; Smith, Osherson, Rips, & Keane, 1988; Cohen & Murphy, 1984; Murphy, 1988). Besides models of understanding complex noun phrases, there has also been research on how people combine the meanings of nouns and verbs in understanding sentences (Gentner & France, 1988).

This chapter is organized into three parts. In the first part, we introduce the process of conceptual combination and discuss its importance as a topic of investigation. We will outline some empirical results of studies on conceptual combination and general challenges that any theory of conceptual combination must address. In the second part, we present three psychological models of conceptual combination and evaluate them in terms of the specific psychological findings and the general challenges outlined in the first part of the chapter. The models are the attribute inheritance model (Hampton, 1987), the selective modification model (Smith et al., 1988) and the concept specialization model (Cohen & Murphy, 1984; Murphy, 1988).

In the third part of the chapter, we examine one of the major assumptions of two of the conceptual combination models. Both the selective modification and concept specialization models propose a type of *slot filling* as the primary mechanism for combining concepts. (A process called elaboration is also very important in the concept specialization model.) In these models, concepts are composed of slots and fillers (as in frames and schemata). One combines a pair of concepts by filling a slot in one concept (which we will call the *head* concept) with that of the other concept (which we will call the *predicate* concept).[2] A slot in the head concept is restricted to having the predicate concept as its filler or value. For example, to interpret a combination like *red box*, one finds a slot in *box* (i.e., the *color* slot) that can be filled by the concept *red*. The concept *red box* is thereby restricted to having *red* as the filler of its *color* slot.

In this section, we will suggest that slot filling may be a common default strategy for combining concepts. To address this hypothesis, we discuss some results from a preliminary study that examines the kinds of descriptions that people give for noun-noun concepts. In this study, people defined novel combinations of count and mass nouns that were either artifacts or natural kinds. While we found evidence for slot filling, a number of examples from this study appear to be exceptions to the slot-filling view of conceptual combination. In order to account for the full range of results, we will propose some other mechanisms that may be involved in conceptual combination. In addition, we will argue that these mechanisms operate on much richer conceptual representations than those typically emphasized in the literature on conceptual combination. In particular, the relational structure in concepts plays an important role in how they are combined.

The Conceptual Combination Problem

Studying how people combine concepts is important for several reasons. First, the use of novel complex noun phrases is a very common, natural, and creative way to fill a vocabulary gap. People often introduce new terms into a language by combining existing words rather than inventing new words. For example, to denote a particular kind of table that supports computers, one might introduce the phrase "computer table," instead of inventing a new word. Novel complex noun phrases are common in newspaper headlines, where they concisely convey important information (e.g., "Record Pentagon Procurement Overcharges Cited," which appeared in the Washington Post on New Year's day, 1989). The coining of novel noun phrases is evident in children at an early age (E. Clark, Gelman, & Lane, 1985). Indeed, many researchers have speculated that conceptual combination provides a route to language change and growth (Downing, 1977; H. Clark, 1983; Gentner & France, 1988). Downing (1977, page 823) suggests that the creation of novel concept combinations, "serves as the back door into the lexicon."

Second, any general-purpose natural language processing system will have to interpret complex noun phrases. Building systems that can understand such combinations is very difficult (e.g., Brachman, 1978; Cottrell, 1988; Finin, 1980; Hirst, 1983). Knowledge of how people combine concepts might assist in the development of such systems.

Third, studying how concepts combine can provide a way to constrain theories about how concepts are represented. In fact, work by Edward Smith and Daniel Osherson on conceptual combination presented a serious challenge to theories of concepts that were derived from research on single concepts (Osherson & Smith, 1981, 1982; Smith & Osherson, 1984). In particular, they demonstrated that prototype theories augmented with fuzzy set theory accounts of conceptual combination could not predict a number of findings on how people combine concepts. This research provided clues to conceptual structure that one may not have been able to discover by just studying single concepts. Based on our own studies of how concepts combine, we will also suggest how concepts should be structured.

For these reasons and others, there has been increasing interest in conceptual combination. Below, we describe some recent studies, as well as a number of general characteristics of conceptual combination that make it a challenging problem.

Representational Assumptions

Crucial to any discussion of how concepts are combined is some notion of how they are represented. Researchers in the field have used different representations for concepts as well as different terminology for the same representation. To keep our discussion of representational issues explicit and clear, we will briefly define some terms. The term "attribute" or "feature" will refer to any property of an object that is represented in the concept of that object. So, for example, "has a pair of wings," "is colored red," and "flies" might be attributes or features that are represented in the concept of *robin*. In describing his model, Hampton (1987) uses attribute in this manner. Many researchers, however, distinguish between slots and fillers when discussing properties of objects that are represented in concepts. In slot and filler notation, the attributes of *robin* would be represented as the tuples (robin part wings), (robin color red), (robin locomotion flies). The first term in each tuple is the concept name, the second term is a slot of the concept, and the third term is a filler or value of the slot. One can view slots and their fillers as dimension-value pairs of a concept. Researchers often discuss slots and fillers in the context of frames. A *frame* is a knowledge structure that represents one's concept of a stereotypical situation or object (Minsky, 1975). It consists of a list of tuples that are generally true of the stereotypical situation or object. (A *frame instance* would represent a specific example of that situation or event—e.g., a specific robin.)[3] So, the frame for *robin* would include the tuples above (as well as others). With some slight modifications (discussed later), Smith et al. (1988) use frames to represent concepts in their model. One can view a concept as being more than a list of slots and fillers, however. A *structured* frame consists of a structured list of slots and fillers. The list is structured in the sense that it captures various relationships between a slot and its filler and between different slots and fillers. For example, in the concept *pie*, the slot *made-of* might indicate that its filler must be something edible, capturing the fact that pies are made of edible things. As a second example, in the concept *rectangle*, the slot *area* might indicate that its filler is the product of two other fillers (namely, the fillers of the *height* and *width* slots). With some slight modifications, Murphy (1988) uses structured frames to represent concepts in his model.

Ambiguity

Combining concepts involves three kinds of ambiguity—syntactic, lexical, and relational ambiguity. In *syntactic* ambiguity, the concept that a constituent modifies is ambigous. When the number of constituents is more than two, there is the "who modifies whom" problem of combining concepts. In a concept pair in English, the first concept almost always modifies the second concept. However, when there are more than two concepts, determining who modifies whom is not straightforward. The combination is syntactically ambiguous. Often, combinations are nested within other combinations. So, in *solid state RCA color television*, a system or person must recognize that *solid* modifies *state* and that this combination modifies *television*. In *water meter cover adjustment screw*, *water* modifies *meter* and the combination modifies *cover*, which in turn forms a new combination *water meter cover* that modifies the final concept. These examples also suggest that a model of conceptual combination must have a mechanism for the recursive processing of nested combinations.

In *lexical* ambiguity, one or more of the meanings of the constituent words of the combination is ambiguous. Many words of English have more than one meaning and many common words have a very large number of meanings (Hirst, 1983). As an example of a phrase that has ambiguous constituents, consider "ball bat." The word "ball" could mean (among other things) any of the balls used in playing sports or it could mean the kind of ball that one dances at. The constituent "bat" could refer to a type of animal or sports equipment.

In *relational* ambiguity, the relation between the constituents is ambiguous. In this case, it is neither the individual constituents nor their syntactic relationship that is ambiguous but rather, how they fit together—their conceptual relationship. Consider the example of elephant tie above. Assume that the meanings of the constituents are unambiguous (e.g., that *elephant* means a type of animal and that *tie* means a type of clothing). Given an appropriate context, *elephant* and *tie* could be related to each other in many possible ways, as in, "a tie worn by circus elephants," "a tie that is large like an elephant," "a tie that has pictures of elephants on it," and so on. In fact, constituents can be related to each other in an arbitrary number of ways that can be very plausible, given the appropriate context (Kay & Zimmer, 1976; H. Clark, 1983).

Concept-dependent Combination Processes

How a particular predicate concept is combined with the head concept often depends on the head concept. To see why, consider a simple, straightforward system in which concepts are combined in a way that is *independent* of the head concept. A particular predicate concept would be combined with any head concept in the same way. For example, the concept *red* might be combined with any head concept *X*, to mean "an X that has the color red." Therefore, for each concept, a straightforward rule would describe how the concept combines with all others when it functions as the predicate.

However, there is increasing evidence that, in human language, predicate concepts do not combine with all concepts in the same way (e.g., Halff, Ortony, & R. C. Anderson, 1976; Rips & Turnbull, 1980; Murphy, 1988; Medin & Shoben, 1988). As an example, a *flower man* is a man who sells flowers, a *flower garden* is a garden that contains flowers, a *flower painting* is a painting that depicts a flower, a *flower necklace* is a necklace made out of flowers, and so on. In these cases, how the predicate concept *flower* combines with the head concept varies as the head concept varies. Medin and Shoben (1988) provide evidence from typicality ratings for this claim. Results of their second experiment suggest that, for example, *gold* is combined with *coin* to mean "a coin made out of gold" but that it is combined with *railing* to mean "a railing with the color of gold." Murphy (1988) also showed that the meaning subjects gave for a simple adjective varied with the noun that it was combined with. These findings suggest that specifying the combinatorial rules of conceptual combination will not be straightforward.

Typicality effects

Smith and Osherson (1984) describe several findings involving the typicality of adjective-noun concepts. A complete theory of conceptual combination needs to specify what it is about the structure of adjective-noun concepts that accounts for these effects. The first finding, called the conjunction effect, is that the typicality of an instance to an adjective-noun concept (i.e., a conjunction) exceeds its typicality to the noun concept. So, a particular red apple is more typical of the concept *red apple* than it is of *apple*. The second finding is the compatible-incompatible conjunction effect—that is, the conjunction effect is greater for incompatible than compatible conjunctions. Here, an incompatible conjunction is one in which the adjective denotes an unlikely filler for a slot of a noun (e.g., as in *blue apple*) and a compatible conjunction is one where the adjective denotes a likely filler for a slot of a noun (e.g., as in *red apple*). So, the extent to which a blue apple is judged more typical of *blue apple* than *apple* is greater than the extent to which a red apple is judged more typical of *red apple* than *apple*. The third finding, called the reverse conjunction effect, is that the typicality of a noninstance to an adjective-noun concept is less than its typicality to the noun concept. So a blue apple is less typical of *red apple* than it is of *apple*.

Relative Ease of Combining Concepts

Some concepts are easier to combine than others. Reaction time studies suggest that several factors affect how easy it is to understand complex noun phrases. One factor is the *form class* of the predicate concept—in particular, whether the predicate is an adjective or noun. Murphy (1990) found that subjects understood adjective-noun pairs faster than noun-noun pairs. For example, people understood "pleasant punishment" more quickly than "bear punishment." It is unlikely that this result could be explained by differences in the familiarity of the objects named by the phrases or in word frequency. Noun-noun pairs were used that had been previously judged as interpretable. Furthermore, the predicate nouns actually had higher word frequencies than the adjectives. One possible reason for the difference is related to the different roles that form classes play in language. In general, adjectives function as operators whose role in language is to pick out a particular slot of a noun to fill. Often, an adjective picks out the same slot of many different nouns. For example, "green" picks out the color slot of the nouns in green apple, green table, green grass, and so on. In contrast, a noun primarily serves to establish reference to individual objects or categories (Gentner & France, 1988). Using a noun as an operator *violates* its preferred use as a referent. An adjective-noun phrase might be easier to understand than a noun-noun phrase because both of its constituents are playing their primary roles whereas in a noun-noun phrase, the predicate noun is playing a role that violates its preferred role.

A second factor that may affect ease of understanding is conceptual complexity. Murphy (1990) prefers this explanation for why noun-noun pairs are more difficult to understand than adjective-noun pairs. Nouns are more conceptually complex than adjectives. In general, compared to adjectives, there is more knowledge represented in nouns. To understand a noun-noun phrase, people must combine two relatively complex representations whereas to understand an adjective-noun phrase, they must combine one simple and one complex representation. Therefore, combining a pair of nouns should involve more computation.

Within the class of adjective-noun concepts, Murphy (1990) has identified three factors that affect comprehension: typicality, relevance, and predication. People more quickly understand a combination containing an adjective that is *typical* of the noun than one containing an adjective that is *atypical* (e.g., "red apple" is easier to understand than "blue apple"). (Familiarity of adjectives was controlled for by having each adjective serve both as a typical and atypical adjective, when paired with different nouns.)

A combination containing an adjective that is *relevant* to the noun is easier to understand than one containing an adjective that is *irrelevant*. According to Murphy, an adjective is relevant if it picks out a slot that is present in the noun (i.e., is part of the representation of the noun). So, the adjective "cold" is relevent to *beer* because it picks out *temperature*—a slot that is part of the concept of *beer*. In contrast, "cold" is irrelevant to *garbage* because *temperature* is not part of the concept of *garbage*. One must infer the fact that garbage has a temperature, presumably by inheritance from its superordinate. Therefore, the phrase "cold beer" is easier to understand than "cold garbage." Relevance is independent of typicality. In this experiment, typicality was measured by the proportion of objects in a noun category that had the adjective property. In the example above, Murphy found that the proportion of objects in the category *beer* that had the attribute *cold* was judged to be about the same as the proportion in the category *garbage* that had this attribute.

Finally, people more quickly understand combinations containing an adjective that has a *predicating* relationship to a noun than one having a *nonpredicating* relationship. An adjective is predicative if the combination can be mapped onto a sentence of the form *noun be adjective* that makes sense and reflects the meaning of the combination (e.g., Levi, 1978). So, "ugly" is a predicating adjective in "ugly painting" because the phrase can be mapped onto "The painting is ugly" — a sentence that makes sense and reflects the meaning of "ugly painting." In contrast, "rural" is a nonpredicating adjective in the phrase "rural policeman." The sentence "The policeman is rural" does not make sense.

Within the class of noun-noun compounds, several studies have investigated differences in ease of understanding. Murphy (1990) found that context can speed the interpretation of noun-noun phrases. In one study, novel noun-noun phrases were preceded by either helpful or neutral contexts. A helpful context was one that plausibly indicated how the predicate noun was related to the head noun. The neutral context mentioned the predicate noun and the head noun but did not indicate a plausible relation. When subjects actually read sentences containing the noun-noun phrases, they understood those embedded in the helpful contexts more quickly than those embedded in the neutral contexts. This finding suggested that when the context specified the relation between the predicate noun and the head noun, the combination process was faster.

Wisniewski (1990) identified another factor that predicts how easy it will be to combine concrete, artifact nouns—namely, the functional scope of the object named by the head noun. The *functional scope* of an artifact refers to the range of objects that can enter into the artifact's function. The range can be relatively unconstrained such that many objects can participate in the object's function. For example, the functional scope of *soap* is unconstrained because many objects can enter into its function (i.e., many objects can be cleaned). The range can also be relatively constrained such that few objects can participate in the function. For example, the functional scope of *comb* is constrained because few objects can enter into its function (i.e., few objects can be straightened or styled with a comb). Typically, functions whose scopes are unconstrained are those for which achieving the function depends on a nearly universally-applicable characteristic of objects, i.e., a characteristic that almost all concrete objects can possess. For example, the functional scope of a box is relatively unconstrained (i.e., a box can be used to contain many things). This is because the function of box depends for its achievement on a nearly universally applicable characteristic—i.e., the characteristic of occupying finite volume so that being contained is possible.

Wisniewski found that people more quickly understood noun-noun phrases involving head nouns with unconstrained scopes (e.g., "jacket box") than those involving head nouns with constrained scopes (e.g., "jacket fork"). The results suggested that for artifact nouns, people often interpret compounds by trying to relate the predicate noun to the function of the object named by the head noun. In the case of a head noun with unconstrained functional scope, people can easily relate the predicate noun to the head noun's function and thus interpret the noun-noun pair. But, in the case of a head noun with constrained functional scope, people must seek other ways to meaningfully relate the constituents, thus increasing comprehension time.

Differential Mutability

Gentner (1981) suggested that words vary in terms of their *mutability*. Specifically, she formulated the verb mutability hypothesis: the meanings of verbs and other predicate terms are more likely to be altered to fit the context than are the meanings of object-reference terms. In support of this hypothesis, Gentner and France (1988) examined how people paraphrased noun-verb combinations that were either semantically natural or semantically strained. An example of a combination that is semantically natural is "The lizard limped." The combination meets the requirement that an animate object (i.e., lizard) be the agent of the verb (i.e., limped). In contrast, "The lantern limped" is semantically strained—the agent of "limped" is not animate. Gentner and France found that when a combination was semantically strained, people paraphrased the combination by altering the meaning of the verb more than the noun. So, people might paraphrase "The lantern limped" as "The lighting device gave off a flickering light" (which alters the meaning of "limped" and preserves the meaning of "lantern"). They seldom paraphrased such combinations by altering the meaning of the noun and preserving the meaning of the verb (e.g., as in "The bright person walked lamely").

Gentner and France (1988) suggested that the nouns in a *noun is a noun* sentence show a similar pattern of differential mutability (though less extreme). In these sentences, the predicate noun (which functions as operator) typically adapts its meaning to the subject noun (which functions to establish object reference). Thus, "The acrobat is a hippopotamus" conveys a clumsy acrobat, whereas "The hippopatomus is an acrobat" conveys an agile hippopotamus. The interpretation of noun-noun concepts should parallel this finding. That is, the first noun should function as an operator whereas the second noun should function to establish object reference (we will discuss some exceptions to this rule later).

Emergent and Interacting Features

Features often emerge in concept combinations that are not present (or at least not salient) in the constituents of those combinations. Murphy (1988) found that subjects judged certain features to be typical of adjective-noun concepts but atypical of the noun or adjective concept alone. For example, people believe "lose money" is a typical feature of *empty store* but an atypical feature of *store* or *empty*. Murphy (1988) argued that such a feature was not a conceptual part of either the constituent *empty* or *store* but rather emerged through an interaction of the constituents and people's general world knowledge. Gentner and France (1988) also suggested that features emerge in concept combinations that are not present in the constituents. They found that when a noun-verb combination was semantically strained, people often altered the verb's meaning by invoking a novel meaning of the verb. For example, in one case, the sentence "The lizard worshipped," was paraphrased as "The small gray reptile lay on a hot rock and stared unblinklingly at the sun." In this example, the feature "stared unblinkingly at the sun" is not highly typical of either concept. Gentner and France (1988) argued that when paraphrasing the meaning of verb in such combinations, people often go beyond simply selecting from a range of prestored aspects of verb meanings. Instead, they adapted the meaning of the verb to fit the noun's meaning. The implication of all of these findings is that conceptual combination is not a closed operation and involves knowledge other than the lexical entries of the constituents.

Features in a concept combination also interact. For example, people believe that wooden spoons are large spoons whereas metal spoons are small spoons (Medin & Shoben, 1988). In this example, the *made of* dimension of *spoon* interacts with the *size* dimension. One interpretation of this finding is that a combination inherits correlations contained in the head concept that are made relevant by the predicate concept. So, the representation of *spoon* might contain information about the correlation "spoons that are made of wood are also large." This correlation would influence the interpretation of *wooden spoon*.[4] The implication of this finding is that concepts cannot be represented simply as lists of features, as implied by many past theories (e.g., Rosch & Mervis, 1975; Smith, Shoben, & Rips, 1974). Instead, concepts also capture dependencies and relations between features. For example, features may be statistically correlated (as in the spoon example above), causally connected, (e.g., *has wings* and *flies* in the concept *bird*), functionally related (e.g., the legs of a table typically support its top), mathematically related (e.g., the volume of a cube is the product of its height, width, and length), and so on. A model of conceptual combination must take into account such dependencies.

Concept Combinations as a Heterogeneous Class

We began this chapter by suggesting that conceptual combination is very broad in scope. Restricting conceptual combination only to noun-noun and adjective-noun pairs, there are still a number of psychologically important dimensions along which such combinations can vary. First, as previously noted, there can be form class differences—the predicate term of a combination can be either a noun or an adjective. In addition, an adjective can have a predicating or nonpredicating relationship to the head noun (as described above).

Combinations can be *conjunctive* or *nonconjunctive*. A conjunctive concept designates a category whose members belong to both constituent categories (Hampton, 1987). For example, the members of *pet iguana* are both pets and iguanas. The members of *red truck* are both red things and trucks. In contrast, the members of a nonconjunctive category are members of only one constituent category (that named by the head noun). So, apartment dogs are dogs but not apartments.

Combinations also vary in their degree of familiarity—from well-known, lexicalized terms (e.g., "apple pie") to novel phrases coined by eccentric writers. (For example, the counterculture author Richard Brautigan, 1967, titled the last chapter of *Trout Fishing in America*, "The Mayonnaise Chapter," a reference to the fact that the chapter ended with the word "mayonnaise.") It is assumed that novel terms are interpreted by combining the meanings of the individual constituents to form a new meaning (Murphy, 1988). A lexicalized term, on the other hand, is assumed to be interpreted by directly accessing its meaning, rather than by deriving it from the meanings of its constituents. As support for this claim, Lees (1968) suggested that as a combination is frequently used in a language, it may lose components of its initial meaning or gain aspects of meaning that are not derivable from its constituents. For example, the lexicalized term "marshmallow" originally named a type of plant that lived in marshes, then came to mean a confection made from the root of this plant and today means a soft, spongy confection made of sugar and corn syrup, roasted over camp fires. Here, the term marshmallow has lost its initial meaning (based on the constituents "marsh" and "mallow") and gained a meaning that is not even derivable from its constituents.

Summary

Taken together, these phenomena suggest that an adequate theory of conceptual combination will be rather complex. In noun-noun pairs, the predicate noun can be related to the head noun in arbitrary ways. Therefore, it does not look promising that a theory can be constructed out of a set of rules that maps constituents onto a small set of relations between them. The theory must also postulate some way of choosing the appropriate meanings of constituents (i.e., resolving lexical ambiguity) as well as the appropriate relations between them (resolving relational ambiguity). For combinations that contain more than two constituents, the theory must have a mechanism for selecting which constituents to combine (i.e., resolving syntactic ambiguity). Importantly, these mechanisms will have to interact with the context surrounding a concept combination. Finally, such a theory must also be able to represent not only the rich, internal structure of concepts but the general, world knowledge that lies outside those concepts, since this knowledge is often used to combine those concepts.

Models of Conceptual Combination

In this section, we describe the representational and processing assumptions of three models of conceptual combination. All of these models are *intensional:* A combination *XY* is formed by using representations of *X* and *Y*. These theories can be contrasted with extensional theories of conceptual combination (e.g., Osherson & Smith, 1981; Zadeh, 1965) in which a combination *XY* is formed by intersecting the sets of the members corresponding to *X* and *Y*. Strong evidence suggests that the psychological validity of extensional theories is untenable (e.g., Osherson & Smith, 1982; Murphy, 1989).

These intensional models have focused chiefly on the interpretation of adjective-noun compounds and/or noun-noun compounds (although Smith et al. have extended their model to account for adverb adjective-noun compounds like "very red apple"). We will restrict our evaluation of these models to how well how they account for the range of psychological findings on adjective-noun and noun-noun combinations that were reviewed in the first section. To describe each model, we first note the scope of combination phenomena that it has explicitly addressed, then discuss the model's representation of concepts and the combinatorial processes that it postulates. Finally, we summarize evidence for and against the model.

The Attribute Inheritance Model

Hampton (1987) proposed a model for how people interpret novel conjunctive concepts, such as *machine vehicle*, *tool weapon*, and *sport game*. These concepts are a subset of noun-noun combinations. As mentioned, the members of a conjunctive category are members of both constituent categories. So, a member of the category *pet shark* is both a pet and a shark. Specifically, Hampton's model has been applied to conjunctive concepts of the form "X that is a Y" (i.e., "pet that is a shark") rather than conjunctive noun-noun concepts. (It is not clear whether this syntactic difference would significantly affect how people process the two types of conjunctions.)

In Hampton's model, concepts are represented as lists of independent attributes, weighted by importance. As a default, a conjunctive concept is formed by taking the union of the attributes belonging to its constituents, and reweighting them in the resulting conjunctive concept. The weighted importance of an attribute for the conjunctive concept is a rising monotonic function of the attribute's importance weights associated with its constituents. There are several cases in which an attribute of a constituent concept will fail to be inherited by the conjunctive concept. First, attributes that are true of one constituent but impossible or highly implausible for the other will not be included. For example, the attribute "is warm and cuddly" which is generally true of pets is a highly implausible attribute for sharks. Therefore, the conjunctive concept pet shark would not contain this attribute. Second, if the average importance of an attribute for the constituent concepts is low then it may be correspondingly low for the conjunctive concept and fail to be inherited. Third, attributes from each constituent may be incompatible or conflict with each other such that the conjunctive concept may contain one but not both. So, "lives in a domestic environment" (for the constituent *pet*) and "lives in the ocean" (for the constituent *shark*) are incompatible and only one of them would be contained in *pet shark*. Hampton argues that the attribute that is chosen is the one that is most compatible with the other attributes of the conjunctive concept.

In addition to these attribute inheritance failures, there are also situations that strongly predict that a attribute will be inherited. Specifically, an attribute that is necessary or highly probable for either constituent will also be included in the conjunctive concept. So, the attribute "has gills" would be included in *pet* *shark* since it is a necessary or highly probable attribute of sharks.

Hampton also hypothesized that a constituent of a conjunctive concept would often *dominate* the other constituent in the sense that it would have more salient and important features. As a consequence, combining the constituents would result in a conjunctive concept that included more important attributes of the dominant concept. This hypothesis was based on previous studies in which Hampton (1988) noted that constituent concepts often contributed unequally to the determination of the typicality of items in their conjunctive concept. That is, for some pairs of constituents, *X* and *Y*, Hampton (1988) found that the typicality of items to the constituent *X* carried more weight in predicting the typicality of those items to *X that is a Y* than the typicality of those items to *Y* did. On this basis, Hampton categorized a number of concepts as dominant.

Evidence for the Model

Hampton obtained evidence supporting his model from subjects' listings of attributes and ratings of their importance In general, attributes that belonged to (i.e., were rated as important for) either or both constituents were inherited by the conjunctive concept, supporting the model's assumption that a conjunctive concept includes the important attributes of its constituents. Not all attributes belonging to the constituents were inherited by the conjunctive concept, however. But, in more than half of these inheritance failures, the attribute had a low average importance rating for the constituent concepts. Therefore, as predicted by the model, these attributes should not be inherited.

Hampton also showed that attributes that were necessary for defining a constituent were inherited by the conjunctive concept whereas those that were impossible for a constituent were not inherited. To show this, Hampton classified an attribute as necessary for a constituent if subjects had rated it as "necessarily true of all possible examples of the constituent" and impossible for a constituent if subjects had rated it as "necessarily false of all possible examples of the constituent." In virtually all cases, necessary attributes for one or both of the constituents were also necessary attributes for the conjunctive concept. Impossible attributes for one or both of the constituents were also impossible attributes for the conjunctive concept.

To investigate whether the importance of an attribute in the conjunctive concept was a rising function of its importance in each constituent, Hampton performed regression analyses. In general, a weighted average of the constituent scores best predicted the importance of an attribute in the conjunct.

Finally, Hampton showed that dominant concepts did in fact have more attributes that were important than nondominant concepts. Furthermore, in regression analyses, the importance of an attribute for the dominant concept carried more weight in predicting the importance of that attribute in the conjunctive concept.

Evaluation of the Model

Hampton's work has shown that attributes defining a conjunctive concept are, to a large extent, reasonably predicted by a model that specifies the inheritance of attributes from the individual constituents of the conjunctive concept. Hampton has identified a number of possible factors (incorporated into the model) that determine which attributes are inherited and what their importance will be in the conjunctive concept, as well as which attributes will be excluded in the conjunctive concept. These factors include the necessity, impossibility, and importance of the features in the constituents and the dominance of a concept. Given the predictive success of the model, future work might explicitly specify the processes that underly these factors. Currently, such factors are primarily determined by subjective ratings. One example of an unspecified process in the model concerns inheritance failures of impossible attributes. Given an important attribute for one constituent, exactly what process determines that the attribute is impossible for the other constituent, and therefore impossible for the conjunctive concept? As another example, subjective ratings determine dominance of one concept with respect to another concept. What are the underlying reasons for why a concept dominates over another?

The most important limitation of the attribute inheritance model is its lack of generality. Currently, it only applies to a very small set of conceptual combination phenomena of the form, "X that is also a Y," where "X" and "Y" are nouns. Many — perhaps most — noun-noun combinations are not conjunctive (Murphy, 1988). For example, dog sleds are sleds but they are not dogs, apartment dogs are dogs but not apartments, and so on. As formulated, Hampton's inheritance rules will not correctly apply to these noun-noun concepts. For example, in the case of *dog sled*, inheritance of necessary attributes predicts that "breathes" should be an attribute of *dog sled* (because dogs must breathe) whereas noninheritance of impossible attributes predicts that it should not be (because sleds cannot breathe). Inheritance of necessary attributes predicts that "has walls" should be an attribute of *apartment dog* whereas noninheritance of impossible attributes predicts that it should not be, and so on. In general, nonconjunctive combinations are characterized by *attribute inheritance asymmetry* between the predicate noun and head noun. Almost all of the attributes of the combination come from the head noun. For example, the attributes of *apartment dog* are almost all taken from *dog* just as the attributes from *dog apartment* are almost all taken from *apartment*. Therefore, feature necessity and impossibility depends primarily on the head noun.

Other research suggests that the focus of combining nonconjunctive concepts is on determining a *plausible relation between the constituents*, rather than on selecting attributes from the constituents. For example, slot filling models (see next section) would predict that a plausible meaning of *apartment dog* is "a dog that lives in apartments." This meaning captures a relation between the constituents, rather than an attribute that is inherited from one or the other constituent. It does not appear possible to extend the attribute inheritance approach to nonconjunctive combinations without considerable alteration.

It is also unclear whether the model could be extended to encompass adjective-noun concepts. One could view many adjective-noun concepts (e.g., *blue apple*) as conjunctive concepts of the form "X that is a Y" (e.g., "apple that is blue colored"). However, at least in some cases, it appears that Hampton's model incorrectly predicts that the conjunct will fail to inherit important attributes. In the blue apple example, it seems obvious that the attribute "has the color blue" should be a salient feature of *blue apple*. However, this attribute will be unimportant in the *apple* constituent (since almost all apples are red, yellow, and green) and important in the concept *blue*. If *apple* is viewed as the dominant constituent, then a weighted average of the importance of this feature for the two constituents (giving more emphasis to the importance of the attribute for *apple*) would predict that "has a blue color" will not be inherited by *blue apple*.

Also, certain adjective-noun concepts are not conjunctive. As with noun-noun combinations that are not conjunctive, the focus of conceptual combination is not on selecting which features from the constituents become inherited. Consider the combination *square bicycle*. One might interpret this combination as a conjunction of the form "bicycle that is also square shaped." However, our intution suggests that this interpretation does not make sense (one could not ride a bicycle that was literally square-shaped). In contrast, compare this combination to *square box* which probably does make sense when interpreted as box that is also square shaped. A more plausible interpretation of *square bicycle* is "bicycle with a square frame." In this case, the adjective applies to *part* of the object named by the head concept. The feature has a square frame appears to be an emergent rather than inherited property of *square bicycle* (see discussion in last section).

Selective Modification Model

Smith, Osherson, Rips and Keane (1988) proposed a model for constructing adjective-noun concepts from individual adjective and noun concepts. More specifically, it was developed to account for typicality judgments involving adjective-noun concepts. The model has also been applied to adverb adjective-noun concepts (e.g., *very red apple*). In general, the model postulates that the adjective directs the formation of the adjective-noun concept by restricting the filler of a noun slot to the adjective concept and by increasing the diagnosticity of this slot. To take a simple example, the adjective "green" would direct the formation of the concept green apple by restricting the color *slot* of *green apple* to the filler *green* and by increasing the diagnosticity of the *color* slot.[5]

The model has three major characteristics. First, nouns and adjectives are viewed as simple frames which consist of the typical properties associated with their instances. For example, a frame for *apple* would contain typical properties such as "has a round shape," "has a smooth texture," and "has the color red." More specifically, the frame consists of a list of slots and a set of fillers associated with each slot. Each slot has a diagnosticity value and each filler of the slot has a salience score. The diagnosticity of a slot measures how useful the slot is in discriminating examples of the category from examples of contrast categories. The salience score of a filler of a slot reflects its subjective frequency among examples of the category as well as its perceptibility. For ease of explaining the model, the salience score is viewed as the "number of votes" for a particular slot filler. Figure 1 (adapted from Smith et al., 1988) shows examples of the *apple*, *brown*, and *brown apple* frames. Second, the model proposes a mechanism for operating on adjective and noun frames to produce an adjective-noun frame. The mechanism operates as follows. The slots of the adjective frame select the corresponding slots of the noun frame. For each of these slots in the noun frame, there is an increase in the salience of the filler indicated by the adjective and a decrease in the salience of other fillers (i.e., votes get shifted to the filler from the other fillers). In addition, there is an increase in the diagnosticity of the slot.

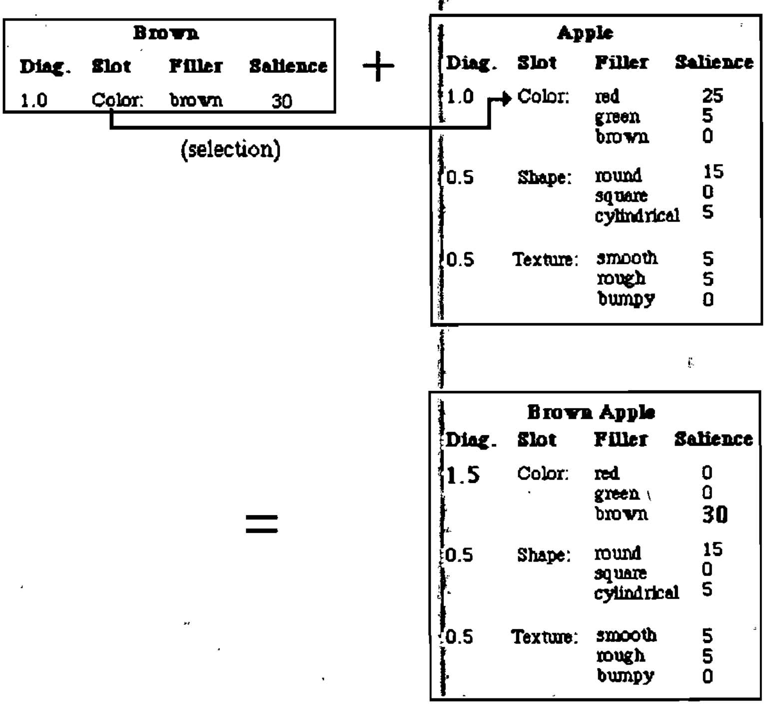

Figure 1. Using slot selection (or slot filling) to combine brown and apple to form brown apple, in the selective modification model.

Figure 1 also illustrates the operation of this mechanism for the frames *brown* and *apple*. Here, the *brown* frame contains a single slot (*color*) and it selects that slot in the *apple* frame, increasing the salience of the filler *brown* and the diagnosticity of the slot *color*, in *brown apple*. The votes for the other fillers of *color* are shifted to the filler *brown*. Although the model is not stated in these terms, one can view selection as slot filling. By selecting a slot in the noun frame, the adjective frame essentially restricts or limits the slot to having the adjective concept as its filler.

The third characteristic of the model is a process for computing the typicality (or similarity) of an instance with respect to its frame, using a slightly modified version of Tversky's (1977) contrast rule. According to this rule, typicality is an increasing function of the votes for an slot filler that are common to the instance and frame and a decreasing function of the votes for a slot filler that are distinct to the frame and distinct to the instance:

$$\text{Typicality}(I, F) = \Sigma_i\, [af_i\,(I + F) - bf_i\,(F - I) - cf_i\,(I - F)$$

where I is the instance, F is the frame, i indexes the slots, $I + F$ is the set of votes of the slot fillers of slot i that overlap in the instance and frame, $F - I$ designates the set of votes of the fillers of slot i that are distinct to the frame, and $I - F$ designates the set of votes of the fillers of slot i that are distinct to the instance. The parameters a, b, c determine the relative contributions of these sets. Figure 2 shows an example of how the typicality of an apple instance to the apple frame has been computed (a, b, and c all have the value 1). In this example, for the slot *color*, the instance and frame have 25 overlapping votes for the filler *red*. The frame, in turn, has 5 votes for the filler *green* that are distinct, and the instance has 5 votes of the filler *red* that are distinct. These sets are then multiplied by the diagnosticity value of the *color* slot. The equation is applied to the other slots in a similar manner.

Apple Instance

Slot	Filler	Salience
Color:	red	30
	green	0
	brown	0
Shape:	round	20
	square	0
	cylindrical	0
Texture:	smooth	30
	rough	0
	bumpy	0

Apple Frame

Diag.	Slot	Filler	Salience
1.0	Color:	red	25
		green	5
		brown	0
0.50	Shape:	round	15
		square	0
		cylindrical	5
0.25	Texture:	smooth	25
		rough	5
		bumpy	0

$$\text{Typicality}(\text{Apple}_I, \text{Apple}_F) = 1 * (25 - 5 - 5) + 0.50 * (15 - 5 - 5) + 0.25 * (25 - 5 - 5) = 21.15$$

Figure 2. Computing the typicality of instance of apple to apple in the selective modification model.

The selective modification model was developed primarily to account for typicality findings involving adjective-noun concepts. As discussed in the first section of this chapter, there were three major findings. The first finding was the conjunction effect. In some cases, the typicality of an instance to an adjective-noun concept exceeds its typicality to the noun concept (e.g., a brown apple is more typical of *brown apple* than of *apple*). The second finding is the compatible-incompatible conjunction effect. That is, the conjunction effect is greater for incompatible than compatible conjunctions. So, the extent to which a blue apple is judged more typical of *blue apple* than *apple* is greater than the extent to which a red apple is judged more typical of *red apple* than *apple*. The third finding was the reverse conjunction effect. In some cases, the typicality of a noninstance to an adjective-noun concept is less than its typicality to the noun concept. So, a blue apple is less typical of *red apple* than it is typical of *apple*.

The model accounts for these findings in a straightforward way. For example, to account for the conjunction effect, consider the *apple* and *brown apple* frames (shown in Figure 1). Notice that both frames are identical except for the *color* slot. This slot has been modified in *brown apple* (during conceptual combination) such that the number of votes for the filler *brown* (i.e., its salience) and the diagnosticity of the slot *color* are greater in *brown apple* than in *apple*. Also, the votes for the other *color* fillers are less in *brown apple* than in *apple*. Note that more votes of the *color* fillers for a particular brown apple will match those of *brown apple* than *apple*, and the *color* slot will have greater diagnosticity in *brown apple* than in *apple*. Furthermore, more votes of the *color* slot for the particular brown apple will mismatch those of *apple* than *brown apple*. Therefore, using the equation above, a particular brown apple will be more typical of *brown apple* than *apple*. As a second example, to account for the reverse conjunction effect, note that a particular brown apple will mismatch a highly salient filler (i.e., *red*) in *red apple* and that this mismatch will be increased (relative to *apple*) by the increased diagnosticity of *color* for *red apple*. Therefore, the particular brown apple will be less typical of *red apple* than it will be of *apple*.

The selective modification model has also proposed a mechanism for combining adverbs with adjectives and nouns. In particular, the model has examined adverbs that intensify or diminish aspects of concepts. For example, an adverb like "very" appears to increase the filler of a slot. In "very red fruit," "very" increases the redness of *red fruit*. Other adverbs like "slightly" and "non" appear to decrease a slot's filler. In "slightly red fruit" and "non red fruit," the adverbs decrease the redness in *red fruit*. In the model, these adverbs function as scalars that multiply the salience scores of slot fillers. The adverb "very" is a scalar greater than 1, "slightly" is a scalar between 0 and 1, and "non" is a scalar less than or equal to 0. For example, "very" would multiply the votes for *red* in *red apple* by some scalar greater than 1. Thus, *very red apple* would have more votes on *red* than *red apple* would.

Evidence for the Model

Smith et al. empirically tested their model by collecting data that allowed them to derive the frames of various concepts and to predict typicality ratings of instances to those frames. They then compared the predicted typicality ratings of the model with actual typicality ratings of subjects.

Specifically, Smith et al. had subjects list slot and filler pairs (e.g., *color-red*) for various examples of vegetables and fruits (e.g., onion, carrot, peach, apple). Using these data, they constructed frames for *vegetable* and *fruit*. The number of subjects who listed a particular slot-filler pair was taken as the salience or number of votes for that slot-filler pair in the concept. Diagnosticity values for each slot were estimated by measuring the extent to which the fillers of that attribute were associated with vegetable but not fruit (or vice versa). Smith et al. then calculated the typicality of the examples to fruit, vegetable and to the eight adjective-noun combinations of the four adjectives "red," "white," "round," and "long" with "fruit" and "vegetable." Next, they compared these typicality ratings predicted by the model with those given by another group of subjects. The predicted values were highly correlated with the subject ratings for most of the adjective-noun concepts that were tested.

Smith et al. also calculated the model's typicality ratings of examples to adverb adjective-noun concepts that paired the adverbs, "very," "slightly," and "non" with the adjective-noun concepts above, and compared them to subjects' ratings. They generally found reasonable correlations between obtained typicality ratings and ratings predicted by the model.

Evaluation of the Model

The selective modification model has a number of strengths. First, it postulates a clear, well-specified combinatorial mechanism. Second, Smith and his colleagues have carefully tested the model's predictions by comparing them to psychological studies of how adjectives and nouns are combined, as well as how some adverbs, adjectives, and nouns are combined.

One limitation of the current model is that in general, it cannot be applied to noun-noun concepts. To understand why the model would have difficulty with constructing a noun-noun frame from two noun frames, consider one of the major processes in the model: the slots of the adjective select the corresponding slots of the noun and modify them in the combination. To use one of Smith et al.'s examples, in forming the combination *shriveled apple*, the slots of the adjective *shriveled* (*texture* and *shape*), would select the corresponding *texture* and *shape* slots in the noun *apple*. These slots would (and probably should) be modified in the *shriveled apple* frame. If we apply an analogous process for combining two noun frames, the slots of the predicate noun should select the corresponding slots of the head noun and modify them in the combination. Clearly, however, many of the slots of the head noun frame that correspond to those in the predicate noun should not be modified. For example, in *dog house*, likely corresponding slots might be *texture*, *shape*, and *size* (to name a few). We would not want the *dog house* frame to contain *dog's* fillers for *texture*, *shape*, and *size*. Dog houses are not furry and shaped like dogs, and they are certainly larger than dogs. Thus, it does not appear possible to extend the selective modification model to noun-noun combinations without considerable alteration.

This limitation may not be not overly important in evaluating the model, however. One might argue that there is form-class specificity in how concepts are combined and that the selective modification model addresses how adjective-noun concepts are formed. Other processes are involved in noun-noun combinations. This view is reasonable, given that the primary roles of adjectives and nouns are different. Adjectives inherently are predicate terms whereas nouns primarily are used as object referents and only secondarily used as predicate terms.

More seriously, the generality of the model's assumptions about adjective-noun combination has been called into question. First, the model assumes that adjectives and nouns are combined using a closed operation—an adjective modifies slots within the noun or adds slots to the noun that are within the adjective. Such an assumption is contradicted by Murphy's (1988) finding that subjects judged certain attributes to be considerably more typical of adjective-noun concepts than of either the noun or adjective concept alone. These results imply that such attributes may not be present in either the adjective or the noun concept but rather emerge through an interaction of the noun and adjective concepts and general, world knowledge. (Gentner and France (1988) found similar results with noun-verb combinations.)

Second, the selective modification model treats the attributes of a concept as independent. Hence, it predicts that modifying one attribute of a concept should not affect other attributes. However, recall Medin and Shoben's finding that attributes in concept combinations can be correlated (e.g., wooden spoons tend to be large whereas metal spoons tend to be small). This finding suggests that an adjective may affect more than one slot in a noun. For example, in *large spoon*, "large" not only determines the filler of the *size* slot of *spoon* but also determines the filler of its *made-of* slot. (Some of the Medin and Shoben results are open to an alternative interpretation, provided by Smith and Gray (1990)—see footnote 3). Besides this specific finding, there is substantial evidence that relations between attributes are very important and that in general, many concepts are best represented as having complex, relational structure (Gentner, 1975, 1981, 1983, 1989; Gentner & Clement, 1988; Gentner & France, 1988; Goldstone, Gentner, & Medin, 1989; Malt & Smith, 1984; Medin, Altom, Edelson, & Freko, 1982; Medin & Shoben, 1988; Medin & Wattenmaker, 1987; Murphy & Medin, 1985; Murphy & Wisniewski, 1989a, 1989b; Palmer, 1978;

Norman & Rumelhart, 1975).

Smith et al. also note that the model does not apply to several kinds of adjective-noun concepts. First, in some combinations, the adjective indicates a slot that is normally not a slot of the noun. For example, "upside-down" cues the slot *orientation* which would not normally be associated with *fruit*. Thus, *upside-down fruit* could not be formed by selecting the appropriate slot in fruit. Presumably, some additional mechanism would add the slot to the noun (along with its diagnosticity, values, and the values' saliences). Second, they note that some adjectives can have complex effects on the nouns that they combine with. For example, "fake" in "fake apple" leaves some slots of *apple* intact (e.g., *shape*, *color*) but negates others (e.g., *taste*, *edibleness*). In the next section, we will make a related claim that mass nouns (e.g., glass, chocolate) have complex effects on their head nouns when they play the role of the predicate term in a combination.

For all its admirable explicitness, the selective modification model does not present a complete picture of adjective-noun combination. It may be more accurate to say that Smith and his colleagues have identified one important process that operates when adjectives and nouns combine. Namely, when an adjective combines with a noun, it may select one or more slots of the noun, changing their diagnosticity and the saliences of their possible fillers. It appears that other mechanisms or representational assumptions are needed to specify how general knowledge affects the combination process and how filling a slot affects other slots in the concept. There is some suggestive evidence that selective modification may be a first stage in adjective and noun combination, with other processes operating later (Smith & Gray, 1990).

Concept Specialization Model

The concept specialization model was explicitly formulated to account for both the interpretation of adjective-noun and noun-noun concepts. As with the selective modification model, slot filling is an important mechanism in the concept specialization model (Cohen & Murphy, 1984; Murphy, 1988). Importantly however, the model hypothesizes a second process that operates in conceptual combination. This process, called elaboration, is driven by people's general background knowledge that lies outside the concepts being combined. The model hypothesizes a richer representation for concepts than the first two models.

The formulation of the concept specialization model was influenced by an AI model called KL-ONE (Brachman, 1977; 1978; 1979). This model represented concepts as structured sets of slots and fillers and it proposed slot filling as one of the primary mechanisms for combining concepts. (Brachman called the mechanism slot restriction rather than slot filling.) The concept specialization model also incorporates these characteristics but differs in that it is being proposed as a psychological model of conceptual combination. Therefore, some characteristics of the model reflect findings in the psychological literature on concepts.

In this model, a concept is a structured set of slots. A slot specifies a default filler and a list of other possible fillers that may fill the slot, weighted by typicality. The slots are structured in the sense that relations between slots and between slots and their possible fillers are also represented. Although this aspect is not spelled out in detail, fillers of different slots may be statistically correlated, and may have causal, numerical, functional, or logical dependencies between them. Concepts are organized into hierarchies and can inherit slots from concepts higher in the hierarchy.

Several of the representational aspects of the model are designed to capture concept typicality. In particular, the possible fillers of a slot and the subconcepts of a concept are ordered by typicality. The typicality of a subconcept to its parent concept (e.g., the typicality of *robin* to *bird*) is measured by computing the degree of family resemblance, after Rosch and Mervis (1975). According to Cohen and Murphy, family resemblance is calculated by counting the number of slots and slot fillers that a subconcept shares with the other subconcepts of the parent and subtracting the number it shares with non-subconcepts of the parent. Thus, subconcepts that have many slots and fillers in common with each other and few in common with non-subconcepts will be more typical of the parent concept.

Concepts are combined using a two-stage process. The first process is similar to that outlined in the selective modification model. Here, a combination is created by filling one of the slots of the head concept with the predicate concept. For example, to interpret elephant box, one would fill a slot in box (e.g., a slot like *contains*) with the predicate concept *elephant*. So, *elephant box* might be interpreted as "a box that contains elephants." Figure 3 illustrates the slot filling process for *elephant box*. There are several ways of determining which slot to fill. First, a predicate term may be listed in the head concept as one of the possible fillers for a slot. If the predicate term is a possible filler of more than one slot, then presumably the slot of which it would be more typical is selected. Second, context (e.g., a discourse setting) can drive the slot-selection process, by activating a slot in the head concept. So, during a discussion of washing and the mention of a phrase like "finger cup," a slot will be activated that reflects a "cup used for washing fingers" interpretation of *finger cup*. Third, one may use general knowledge to determine the best slot. Especially for novel combinations, the predicate term may not be listed as a possible filler for a slot and discourse may be insufficient for selecting an appropriate slot.

The second process is called *elaboration* and it involves refining and augmenting the combination, using world knowledge (Murphy, 1988). This knowledge is used to infer other likely characteristics of the combination. To continue the example above, one might reasonably conclude that an elephant box is larger than the usual box and augment the combination with this fact (probably by filling the *size* slot of *elephant box* with *large*). One also might conclude that an elephant box is sturdier than the usual box and therefore is made of wood rather than cardboard. In Figure 3, *elephant box* has been elaborated to reflect these conclusions. Elaboration may be based on some type of plausible reasoning process (Collins, 1978; Collins & Michalski, 1989). It may also involve recalling examples of the head concept. So, one might recall examples of boxes that contained elephants and use them to refine and augment the meaning of *elephant box*. Hampton (1985) calls this process extensional feedback (see also Cohen & Murphy, 1984). It involves accessing knowledge of actual objects in the world.

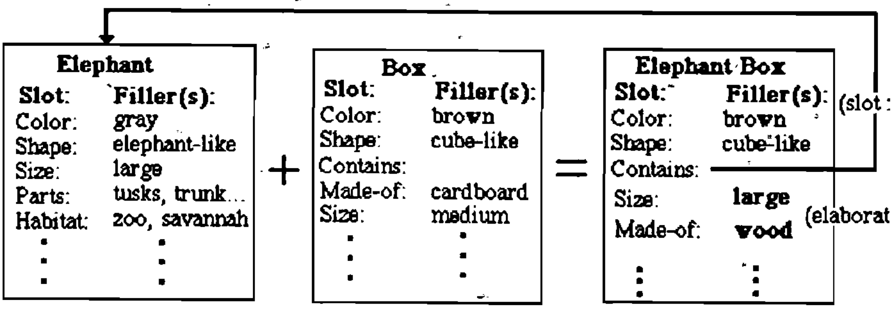

Figure 3. Using slot filling and elaboration to combine ***elephant*** and ***box*** to form ***elephant box***, in the concept specialization model.

Besides the processing assumptions of slot filling and elaboration, the model also assumes that conceptual combination is typically asymmetrical: a combination of the form *XY* is not at all the same as one of the form *YX*. So, for example, an apartment dog is not the same as a dog apartment. This assumption has been emphasized in order to contrast the model with extensional models of conceptual combination, in which an *XY* combination would be formed by intersecting the sets corresponding to *X* and *Y*. By commutativity of set intersection, this view would predict that an apartment dog is the same as a dog apartment. The asymmetry of conceptual combination may be due to the different roles that the predicate and head nouns play (Gentner & France, 1988). In an *XY* combination, the meaning of *X* is more mutable (because *X* functions as an operator) and the meaning of *Y* is more stable (since it serves to designate the referent of an object). The reverse is true for a *YX* combination. Therefore, the meaning of *XY* will be different from *YX* .

Evidence for the Model

Murphy (1988, 1990) details several studies that provide support for the general assumptions of the model. One line of support for the use of general knowledge comes from a previously mentioned study suggesting that features appear in a combination which are not present in either the adjective or noun concepts but which emerge through an interaction of these concepts, mediated by general knowledge. In addition, Murphy's finding that irrelevant-adjective noun concepts were more difficult to understand than relevant-adjective noun concepts suggests that one needs to access concepts outside of the constituents (specifically, superordinate concepts) to understand the former. So, in trying to understand *cold garbage* one will not find an appropriate slot in *garbage* that *cold* can fill (because a slot like *temperature* is not relevant to *garbage*). Instead, one must determine an appropriate slot via inheritance from one of the superconcepts of *garbage*. In contrast, understanding *cold beer* should be easier because such a slot is represented in the concept *beer*.

Murphy's (1988) finding that a helpful context speeds up the interpretation of noun-noun concepts suggests that context may activate slots in a concept, thus speeding the combination process. Context may suggest a plausible slot for the predicate term, making the combination process easier.

Evaluation of the Model

The concept specialization model provides a unifying account of how adjective-noun and noun-noun concepts are interpreted. In this respect, it is more general than the selective modification model (which applies only to adjective-noun concepts) and the attribute inheritance model (which applies only to the small subset of noun-noun concepts that are conjunctive). It is also the only model that has attempted to account for the important role of context in conceptual combination.

On the other hand, while the model's notion of world knowledge (used in the elaboration process) seems necessary to capture emergent features and to determine which slots to fill, it is a vague principle. Murphy (1988) has noted that the model refers to people's knowledge in a rather unconstrained manner and that its use of knowledge is not spelled out to any degree. Moreover, the model has not been empirically evaluated as carefully as either the attribute inheritance model or the selective modification model. Further development of the model will need to take these issues into account.

Nevertheless, the concept specialization model is extremely plausible. Indeed, we suspect that it is the default model for combining nouns. However, in the next section, we will argue that in some cases, it is necessary to go beyond the model's processing and representational assumptions. In particular, while the model assumes structured representations for nouns, the importance of such structure for combining concepts has not been demonstrated. We will suggest that in a number of noun-noun combinations this structure (which includes relations between slots) plays a very important role in the combination process. Besides an emphasis on structured representations, we will also suggest that a complete model of noun-noun combination must employ other processes besides slot filling.

Summary

We have described a number of models of conceptual combination. It is clear that no model is complete. There are two major reasons why this is the case. First, psychological research on conceptual combination is relatively recent compared to work in other areas (memory retrieval, attention, structure and processing of single concepts, etc). There is not yet a large base of empirical studies on how people combine concepts. We need to learn more about this cognitive process that people do naturally and easily.

Second, as argued in the last section, any complete model of conceptual combination will have to be complicated and extensive. The approach taken in all of these models is to carve out a piece of the problem and first attempt to understand that well. In doing so, these models have made a number of implicit, simplifying assumptions which make them incomplete at this point in their development. They avoid the problem of "who modifies whom" by assuming that combinations are composed of only two constituents. The models also implicitly assume that it is clear which meanings of the constituents are being combined (thus avoiding the problem of lexical ambiguity). The models also limit the types of combinations that they address. Smith et al.'s selective modification model has been applied to a subset of adjective-noun concepts. Hampton's attribute inheritance model has been applied to the subset of noun-noun concepts that are conjunctive. Murphy's concept specialization model is the most general model—specifying how people interpret both noun-noun and adjective-noun concepts. However, the model has not explicitly addressed the important role of conceptual structure (e.g., relations between slots) in combining nouns. In the next section, we will suggest that this structure sometimes is involved in combining noun meanings. We will also describe other processes besides slot filling that operate noun-noun combinations.

How do People Define Novel Combinations —What is a Pony Chair?

In this section, we address the generality of slot filling in conceptual combination. Our goal is to examine people's descriptions of novel combinations to see how well they fit this view. We are especially interested in determining other strategies that people use to combine concepts as well as the kinds of noun representations that would be needed to accommodate these strategies. As mentioned, slot filling is a major component of both the selective modification model and the concept specialization model. The authors of these models imply that slot filling typically occurs when people combine concepts. The concept specialization model also postulates a second process (elaboration) that follows slot filling. We will focus here on noun-noun concepts and will not examine the generality of the process for adjective-noun concepts. In terms of evaluating the two models, our data are more directly applicable to the concept specialization model (since it explicitly addresses how noun-noun concepts are combined).

Assumptions and Plausibility of the Slot Filling Process

There are three underlying assumptions involved in slot filling. First, the process is applied to the head concept and to a slot that the head concept contains or can inherit from its superordinate concepts. Second, the resulting combination *XY* is basically a *Y* with an additional restriction on one of its slots. Third, the process involves restricting the filler of the slot to the predicate concept (and not other concepts). So, when forming a combination *XY*, people restrict the filler of a slot in the head concept *Y* to the predicate concept *X*. For example, consider a very plausible meaning of *book box*: "box that contains or holds books." Assume that the concept *box* has a number of slots that can be filled by other concepts. When people interpret a phrase like "book box" they search for a slot in the head concept box that can be filled by the predicate concept *book*. In this case, people interpret *book box* by filling a slot of *box* (that corresponds to "contains" or "holds") with *book*. This slot is restricted to having *book* as its filler. A book box is a box except that it contains books and not other things. (Of course, such a representation does not rule out the possibility that a book box could contain other things. We will ignore this subtle distinction.)

Intuitively, it seems that slot filling is a very natural strategy for combining concepts. There may be several reasons for why people prefer this strategy. First, it allows one to use the predicate noun as a predicate while preserving the integrity or cohesiveness of its meaning, as well as the meaning of the head noun. That is, slot filling may involve *minor adjustments* to noun meanings. Gentner (1981, 1982) has suggested that concrete nouns, relative to other parts of speech, have highly coherent, internally constrained meanings and that people prefer to preserve those meanings whenever they can. Simple nouns typically refer to objects in the world and their meanings incorporate a large amount of perceptual information that is determined by those objects. Other parts of speech, especially verbs, are less tightly constrained by the perceptual world. As previously noted, verbs are more likely to change their meanings than nouns (Gentner & France, 1988). Also, compared to nouns, languages vary more in terms of which meaning components they conflate into verbs (Gentner, 1981; 1982; cf. Talmy, 1978).

Slot filling amounts to asserting a relation between the head noun and the predicate noun (e.g., "box that contains books") and does not disrupt the basic meanings of the nouns. Of course, slot filling may alter the meaning of either noun to some extent. For example, assume that "pear pie" means "a pie made of pears" and that *pear* fills the *made-of* slot of *pie*. In this case, the predicate noun "pear" probably refers to sliced and peeled pears rather than the typical pear. As another example, the head concept *soap* in *tank soap* is probably different than the typical soap (e.g., more abrasive, more concentrated). It is likely that the meanings of these nouns have been altered. However, they probably retain enough of their original meanings so that people would agree that they still refer to pears and soap. Later, we will present examples suggesting that this is not always the case.

Second, in terms of computation, slot filling may be an easier strategy than others. In general, one only has to check the meaning of the predicate noun rather than to alter its structure. Specifically, slot filling may require that one check whether the predicate noun fits certain constraints on the slot. In the example of *pear pie*, filling the *made-of* slot of *pie* with *pear* might involve checking whether *pear* fits a constraint on the *made-of* slot such as *being edible*. In contrast, we will suggest that other strategies require one to dismantle and significantly alter the meaning of one or both nouns. That is, some strategies involve *major adjustments* to noun meanings. Presumably, these adjustments are more computationally complex than those involved in slot filling.

An Experiment

The study that we will describe was largely exploratory in nature. We were interested in assessing the generality of slot filling as well as discovering other combinatorial strategies and the corresponding representations that they operate upon. One way to examine such strategies is to collect a large number of descriptions of many novel combinations. The obvious problem with this approach is that one needs a way to meaningfully sample from the huge number of possible noun combinations. To introduce some constraints, we varied nouns along three conceptually important dimensions: predicate versus head noun position, artifact versus natural kind, and count noun versus mass noun.

Intuitively, we also believed that nouns varying along these dimensions might interact in interesting ways when they were combined. These interactions might result in situations where slot filling was more or less preferred as a combinatorial strategy. For example, intuitively, the "predicate versus head noun position" probably cues whether a noun is an operator or a referent. On the other hand, count nouns and mass nouns may differ in terms of how natural it is to use them as referents and operators. Objects (particularly artifacts) are often composed of mass quantities (e.g., windows made of glass, vases made of clay, etc). One uses count nouns to refer to such objects rather than the mass terms of which they are composed of. Therefore, in a novel combination, one might prefer to use a count noun as an object referent and a mass noun as an operator. A mass/artifact-count combination preserves the preferred roles of the constituents whereas an artifact-count/mass combination violates those roles. As a result, it might be more straightforward to interpret a mass/artifact-count combination by slot filling than an artifact-count/mass combination. In fact, in the predicate position, mass nouns may function like adjectives, picking out a particular slot (i.e., *composition*) to fill. On the other hand, subjects might use some other strategy to interpret count-mass terms (e.g., use the predicate noun as the referent and the head noun as an operator).

We used three groups of nouns to create the noun-noun phrases. One group consisted of 10 count nouns and a second group consisted of 10 mass nouns. A third group of nouns consisted of 10 count nouns, as in the first group. Half of the nouns in each group were artifacts and half were natural kinds. The three groups of nouns are shown in Figure 4. To form noun phrases, we first paired each noun from group 1 with each noun from group 3 and paired each noun from group 2 with each noun from group 3. This procedure resulted in 200 pairings. It also resulted in a hierarchy of combination types, shown in Figure 5. For each pairing, we then formed the two noun-noun phrases that were possible (e.g., for the pairing of "robin" and "clock," the two phrases "robin clock" and "clock robin" were possible). This procedure resulted in 400 noun-noun phrases. Each of 20 subjects defined 20 of these noun-noun phrases.

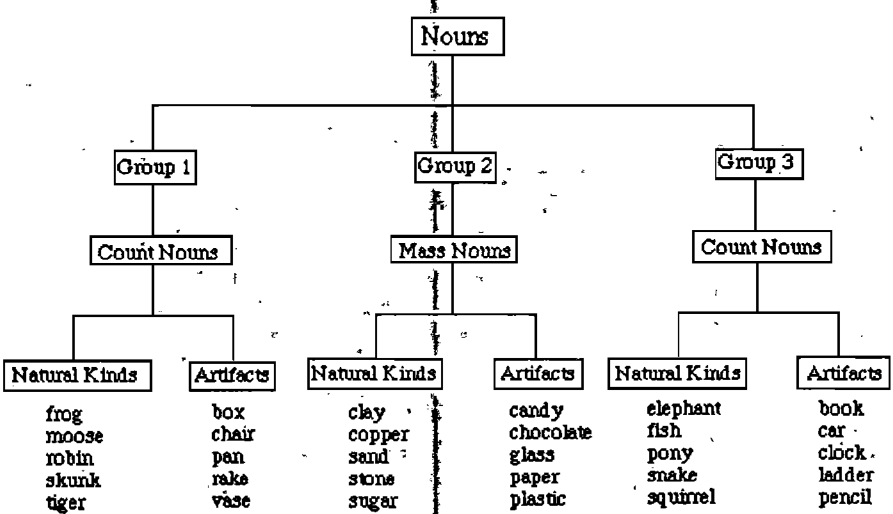

Figure 4. The three groups of nouns used in the experiment.

Subjects read the novel noun-noun phrases and were asked to write down descriptions of their most likely meanings. They were told to pretend that they

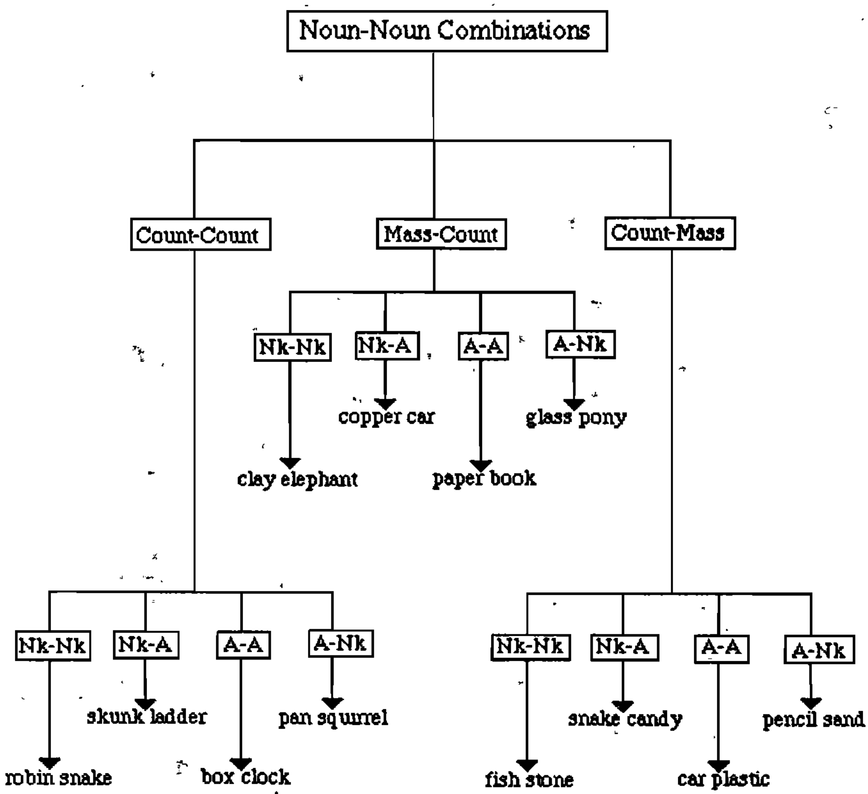

Figure 5. A hierarchy of combination types (with examples of each type) used in the experiment.

had just heard a phrase during a conversation and that they should think of a meaning that seemed most natural to them. Subjects were instructed to try to arrive at meanings that were specific and clear and to define every phrase.

Generality of Slot Filling

To look at the generality of slot filling, we asked two questions about a description that would provide evidence for the slot filling view. The first question was what is the referent of a subjects' description: Slot filling predicts that the referent will be a type of the head concept. For example, if a subject described *book box* as above, "a box that contains or holds book," then the referent would clearly be a type of *box*. In most of the subjects' descriptions, the referent could be determined syntactically. Typically, as in this example, it is the first noun that is mentioned in the subject's description. However, this is not always the case. For example, consider the description for *chocolate pony*, "a pony made of chocolate." Syntactically, a pony is being described. However, conceptually, the referent is not really a pony, but rather chocolate in the *shape* of a pony (according to our intuitions).

The second question is whether all or part of a description can be characterized as *Y slot-relation X*, where *Y* is the head concept, *X* is the predicate concept, and *slot-relation* is some relation being asserted between *Y* and *X* that corresponds to a slot contained in the head concept. The description of *book box* can be characterized as a *Y slot-relation X* ("box contains books"). In contrast, a description such as "squirrel with a black stripe down its back" for *squirrel skunk*, cannot be characterized in this manner, as no relation is being asserted between *skunk* and *squirrel*. Rather, it appears that a property of *squirrel* is being asserted of *skunk*.

To address the first question (i.e., to determine the referents of the subjects' descriptions), we gave the descriptions to a group of undergraduate judges. Each of 20 judges read half (200) of the definitions. For each description, they determined whether the referent was a type of the predicate concept, a type of the head concept, both, or some other object. Specifically, subjects were asked to answer the following question about each description:

> What is the object that is being described? That is, what would be the best name for the object that would let someone know what it really is.

For a given description, this procedure resulted in 10 judgements about the identity of the referents.

The referent of a description was determined by the consensus of the judges. In general, the head noun was the referent, as predicted by slot filling. The judges believed that a majority of the descriptions described types or kinds of the head concept. Interestingly, however, for 151 (38%) of the 400 descriptions, the head concept was not the referent. Two examples of this violation were *chair ladder*, which was described as "a chair that for necessity is used as a ladder," and *paper elephant*, which was described as "paper in the shape of an elephant." (Notice that in the both cases, the predicate noun functions as the referent and head noun as the operator.) Two examples of descriptions in which the head concept was judged as the referent were "a tiger that preys on horses/ponies, etc" for *pony tiger*, and "glass for holding pencils" for *pencil glass*.

To address the second question (i.e., to determine whether a description could be characterized as a *Y slot-relation X*), two graduate students from the University of Michigan read each definition and decided which of two categories it belonged to. If a description included a relation between the two objects named in the phrase, it was placed in the *relation* category. Otherwise, the description was categorized as *other*. We gave the raters several examples of the *relation* and *other* categories, using descriptions that were not from the experiment. For example, they were told that "a factory that is smelly and processes fish" belonged to the *relation* category, as it asserts a "processes" relation between *factory* and *fish*. As another example, they were told that "a dangerous man" belongs to the *other* category (no relation is being asserted between a pair of nouns). Also, some descriptions specified a relation between the head and predicate noun even though one of the nouns was not explicitly mentioned. An example of such a description was "a tiger that likes to read a lot" for the phrase *book tiger*. In this description, a relation between *book* and *tiger* is strongly implied even though book has been omitted from the description. Judges were instructed to place these descriptions in the *relation* category.

Note that this procedure provides a liberal test of the generality of slot filling. The raters only determined that a relation was being asserted between the head concept and predicate concept. They did not have to judge which relations corresponded to slots contained in the head noun. This leads to a generous count, for it includes relations that may not be part of the head noun's frame. For example, one might argue that in the description for *ladder skunk*, "a skunk that climbs ladders, " the slot being filled in *skunk* (i.e., *climb*) originates in *ladder*. (Certainly, climbing is much more typical of ladders than of skunks.) The procedure also does not distinguish those descriptions based solely on slot filling from those that included other strategies in addition to slot filling.

The two raters initially agreed on 87% of their judgments about the descriptions. Differences in scoring were discussed and resolved. The raters judged only 40% of the descriptions as stating a relation. Two examples of descriptions that were categorized as *relation* were "a pan for frying fish" for *fish pan*, and "car made out of copper" for *copper car*. Two examples of descriptions that were judged as other were "a square box" for *box clock* and "a ladder whose rungs are far apart" for *frog ladder*.

The results of these analyses suggest that nouns are not always combined by slot filling. Indeed, the majority of the descriptions in our corpus were not classified as *Y slot-relation X*. An examination of those descriptions that did not conform to slot filling suggests two general conclusions. First, there are other important processes besides slot filling that are used to combine nouns. Some of these processes involve major adjustments to meaning (relative to slot filling). Second, some of these processes operate on noun representations that are more complex than those currently proposed in the literature. We will argue that noun representations must include more than slots and fillers. Importantly, they must include relations between slots within a noun (i.e., internal relations) as well as relations between slots of different nouns (i.e., external relations). In addition, fillers of slots can themselves be complex structures (i.e., nested structures). Below, we describe some of these conjectured processes and the representations that they operate upon. At this point, we will make no claim about their generality, except to say that the noun-noun descriptions that suggest these processes were not rare occurrences in our data.

Other Strategies for Combining Noun Meanings

As we have suggested, slot filling is an example of a process that preserves the basic meanings of the predicate and head nouns. The combination that results from slot filling is a type of the head noun with some relation (designated by the slot) to the predicate noun. In general, this process does not significantly alter the meaning of either noun. We will now suggest that many of the descriptions that do not conform to slot filling reflect processes in which only *part of* the meaning of the predicate noun is involved in the resulting combination. We will also suggest that under some conditions (for example, often when mass and count nouns combine) only part of the meaning of the head noun is involved in the resulting combination.

A large number of the subjects' descriptions (approximately 30%) had the form *property Y* or *Y with property*, where *Y* is the head concept (e.g., "a large frog" was the definition for *elephant frog*). In these descriptions, a property is being asserted of the head concept, rather than a relation between the predicate concept and the head concept (as in slot filling). That is, the predicate noun is not participating *as a whole* in the resulting combination. It is not playing the role of a slot filler. What role then does it play in such combinations? We suggest that the predicate noun plays at least two other roles besides being a slot filler. First, in a process called *property mapping*, the *filler of a slot in the predicate concept* is used as a filler in the corresponding slot of head concept. Second, in a process called *structure mapping*,[6] the complex structure of the predicate noun guides the creation of new structure or the transformation of existing structure in the head noun. We illustrate these processes using some examples taken from our data. Besides property mapping and structure mapping, we will also discuss complex effects that occur when mass nouns in the predicate position combine with count nouns in the head noun position.

Property Mapping

To illustrate this process, consider the description, "a red snake," that was given by a subject for *robin snake*.This description does not fit the slot filling view since a slot in *snake* is not being filled with the predicate concept *robin*. (A description for *robin snake* that would involve slot filling is "snake that eats robins"). Instead, it appears that the filler *red* of the slot *color* in *robin* becomes the filler for the *color* slot in *snake*, as illustrated in Figure 6. (In this figure and those that follow, we have added unspecified connections between slots to emphasize the importance of relations.) Here, the *color* slot of *robin* is *aligned* (or put into correspondence) with the *color* slot of *snake*. The filler of *color* (*red*) is then mapped across and becomes the filler of the *color* slot in *snake*. As in the standard view of slot filling, a slot in the head concept is affected and the resulting combination *XY* is basically a *Y*, with an additional restriction on the slot. However, property mapping restricts the slot to the *filler of a slot* in the predicate concept rather than the predicate concept. Notice that this process leaves the meaning of the head noun basically intact. However, the predicate concept contributes only a small part of its meaning to the combination.

As described, one could incorporate property mapping into the concept specialization model without altering the the model's basic assumptions. The slot filling mechanism would consider both slot fillers in the predicate concept and the predicate concept itself as potential fillers for head noun slots. Moreover, the mechanism could still operate successfully on a list of slots and fillers

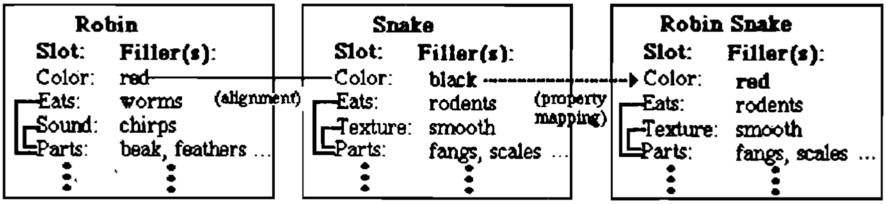

Figure 6. Using property mapping to combine *robin* and *snake* to form *robin snake.*

(the representation for nouns that is emphasized in the model) as long as an alignment of slots could be made to guide the property mapping. However, the next process that we consider is quite different from this augmented view of slot filling and requires more complex representations.

Structure Mapping

We will illustrate this process using three examples taken from our data. The first example is *pony chair*, which was defined as "a small chair." This description does not fit the slot filling view since a slot in *chair* is not being filled with the predicate concept *pony*. One might be tempted to classify this description as an an example of property mapping. Here, the filler *small* of the slot *size* in *pony* fills the size *slot* in *chair*, yielding the interpretation of *pony chair* as a small-sized chair. However, note that the typical pony is actually larger than the typical chair. If one literally interprets *pony chair* as a chair similar in size to a pony, then paradoxically, a pony chair will be larger than most chairs!

The resolution rests on noting that "small" is a relative adjective and that ponies are small relative to other horses. We suggest that *pony chair* literally means a chair that is small relative to other chairs. How then were pony and chair combined to yield this meaning? To interpret *pony chair* in this manner, we suggest that the representations of *pony* and *chair* must be more complex than lists of slots and fillers. In particular, the representation of *pony* includes a relation (i.e., less-than) between its *size* slot and the *size* slot of *horse* which represents the fact that ponies are small relative to other horses. To combine *pony* and *chair*, a similar relation is created between the *size* slot of *pony chair* and the *size* slot of *chair*. The relation represents the fact that pony chairs are

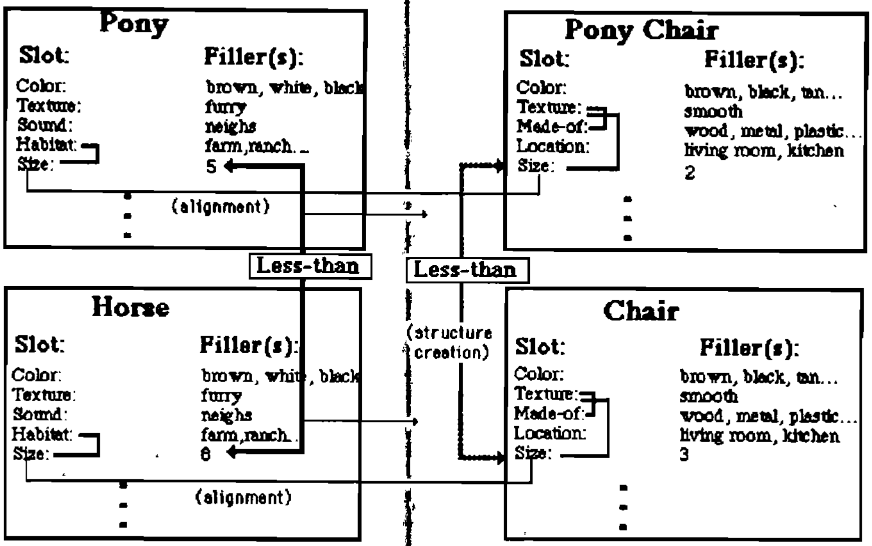

Figure 7. Using structure mapping to combine *pony* and *chair* to form *pony chair*.

small relative to other chairs. This example illustrates a structure mapping which involves aligning (or putting into correspondence) the *size* slot of *pony* with the *size* slot of *pony chair* and the *size* slot of *horse* with the *size* slot of chair (see Figure 7). Then, a *less-than* relation is mapped across between the *size* slots of *pony chair* and *chair*, leading to the notion of a small chair.[7]

There are several important differences between this process and those of property mapping and slot filling. First, neither the predicate concept or a filler of one of its slots functions as a slot filler. Rather, a *structural relation* between one of its slots and the slot of another (closely associated) item guides the creation of a new, similar structural relation in the head concept. Second, the process operates on and creates representations that are more complex than a list of slots and fillers. In this example, the representations include relations between slots of different concepts (i.e., external relations).

A second example is *snake glass*, which was described as a "tall, very thin drinking glass." Once again, this description does not fit the slot filling view. How then were *snake* and *glass* combined? First, note that a *snake glass* resembles the shape of a snake in some way. We suggest that in general, the shape slot specifies a complex structure that relates various aspects of shape to each other (cf. Palmer, 1975; Marr & Nishihara, 1978; Biederman, 1985). In the concept *snake*, this structure might (among other things) indicate that the typical snake is much longer than it is wider. In the concept *glass*, this structure might (among other things) indicate that the typical glass is somewhat taller than it is wider. In this example, we believe that the shape of glass is modified in a way that is *analogous* to the shape of snake. Just as a snake is much *longer* than it is wider, a snake glass is much *taller* than it is wider. Figure 8 outlines how structure mapping might operate to produce *snake glass*. Notice the various parts of structure that have been aligned (e.g. the *length* slot of *snake* has been aligned with the *height* slot of *glass*.) The process results in the *height* of *snake glass* being increased relative to its *width*. Unlike the example of *pony chair*, existing structure in *glass* is being transformed to create *snake glass*

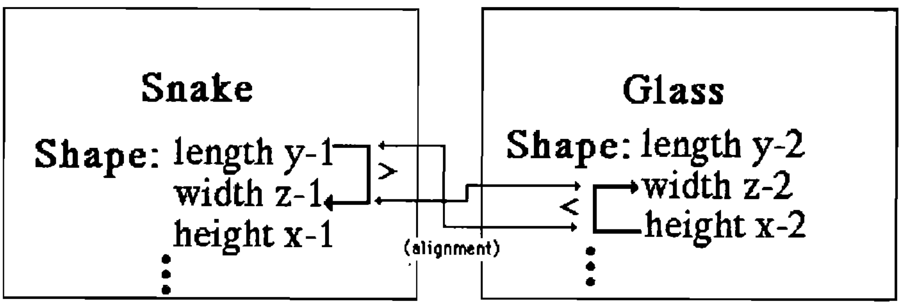

Figure 8. Using structure mapping to combine *snake* and *glass* to form *snake glass*.

rather than new structure being added. On the other hand, interpreting *snake glass* in the manner described might result in the new knowledge that snake glasses are longer than typical glasses. Therefore, one would need to augment *snake glass* with a *longer-than* relation between the *length* slot of *snake glass* and the *length* slot of *glass*. (i.e., an external relation).

A final example is *ladder rake* which was defined as "utensil which is elongated so as to use to reach high places." As in the examples before, this description is not a case of slot filling. We present one possible interpretation of how *ladder* and *rake* are combined. First, assume that the function of *ladder rake* actually shares aspects of both the function of *ladder* and of *rake*. Like ladders, ladder rakes are used to reach high places. Although not specified, one might also surmise that they are used to collect or gather things from high places, thus preserving aspects of the function of rakes. Figure 9 sketches how *ladder* and *rake* might be combined. Notice that the fillers of the *function* slots point to complex structures which we have represented using notions derived from case grammar. To combine *ladder* and *rake*, the function of *ladder* is aligned with that of *rake*. The function of *ladder rake* is created by modifying the function of *rake*. Specifically, the fillers of the *source* and *destination* slots of the *function* of *ladder* are mapped across and become the fillers of the *source* and *destination* slots of the *function* of *rake* (see Figure 9). As a result, a ladder rake has a function that is similar to that of a rake except that one uses a ladder rake to rake in a vertical direction rather than a horizontal direction. Notice also

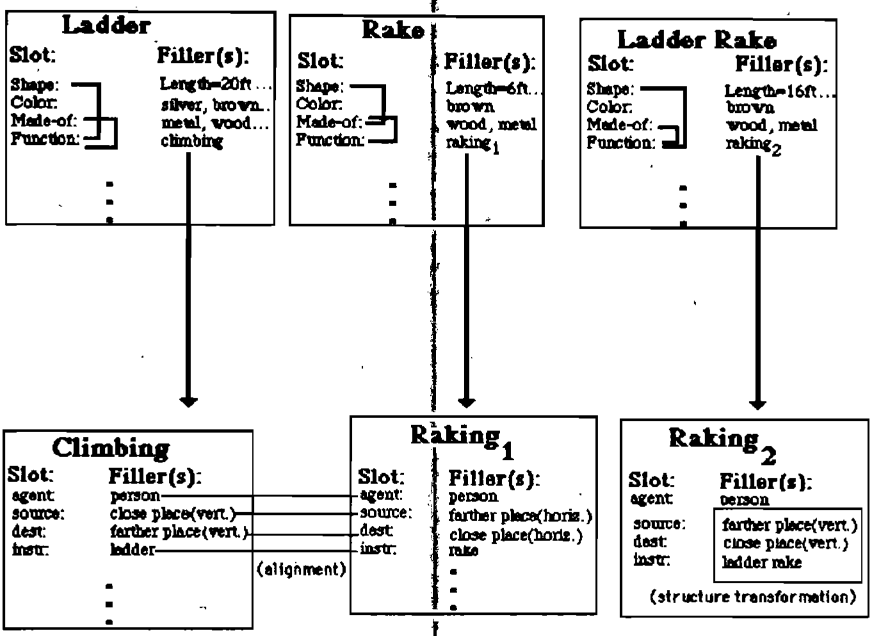

Figure 9. Using structure mapping to combine *ladder* and *rake* to form *ladder rake*.

that the filler of the *shape* slot of *ladder rake* is different from that of *rake* (ladder rakes are longer than rakes because they are used to reach high places). The change in the *shape* slot is an example of interacting properties (see the first section). In particular, the *function* and *shape* slots must include relations between them that capture these interactions.

A variation on the structure-mapping process is that a slot that is filled may actually be one that is inherited by the head concept from the predicate concept. Two examples are *clock tiger* which was defined as "a tiger that can tell time" and *pony frog* which was defined as "a frog that is trained to ride ponies." In both of these (somewhat strange) examples, the relation being asserted is strongly associated with the predicate concept rather than the head concept. These examples illustrate structure creation as a novel slot is being added to the head concept.

Complex Effects of Mass Terms

Mass terms often refer to substances that things are made of. For example, the mass noun "glass" refers to a substance that many objects are made of—windows, vases, bottles, plates, and so on. Therefore, people may be biased to interpret a combination of the form mass term-count term as naming an object that is made of the mass term. In fact, according to our judgment, 59% (59 out of 100) of the descriptions of mass-count terms described an object that was made of that mass term. Some of these descriptions explicitly stated this relation, as in "a clock made out of copper" for *copper clock*. Other descriptions were less explicit, as in "statue" for *stone snake*.

However, it also appears that interpreting a combination in this manner has important effects beyond just indicating the composition of an object. These effects are related to semantic differences in the head nouns that we used in this study. Recall that a head noun was either an artifact or an animal natural kind (see Figure 5). In general, an artifact can be made of a variety of substances (often named by mass nouns) whereas a given animal is generally believed to be composed of one kind of substance. One can often assert that an artifact is made of a variety of different substances and still preserve the identity of the artifact—what appears important though is that the function of the artifact be preserved (Gelman, 1988; Keil, 1986; 1987). A plate for example, can be made of wood, metal, plastic, glass, and so on, and still be a plate (being made of such substances does not affect its function). On the other hand, a dog can't be made of wood or plastic and still be a real dog.

This difference between artifacts and animals suggest different conditions under which head nouns will lose their referential priveleges. When a mass-animal term describes an object made of the mass term, the referent will be some object other than the animal. For example, the description of *chocolate snake* ("chocolate in the shape of a snake") names an object that is made of the mass term. The referent is also the mass term ("chocolate") rather than the animal ("squirrel"). In this example, the head noun has given up its referential priveleges to the predicate noun. Of course, *squirrel* still confers its *shape* on the referent. (In fact, in many contexts, *shape* may be an important property for determining reference. Even though the head noun loses its referential priveleges, it may in a sense, "effect a compromise," by contributing an important property for determining reference.)

In contrast, one can usually intepret a mass-artifact term as an "artifact made of mass term" if doing so would preserve the artifact's function. In these cases, the artifact retains its referential privileges as the head noun. However, if such a description would fail to preserve the artifact's function, then the referent will be some object other than the artifact. Two examples from our data illustrate these different cases. The referent of *clay ladder* ("a ladder made of clay") was judged to be *ladder*. This description also appears to preserve the function of *ladder*—i.e., one could use a clay ladder for climbing. On the other hand, the referent of *candy ladder* ("a long strip of candy") was judged to be *candy*. It appears that function of *ladder* would not have been preserved if the term was literally interpreted as "a ladder made of candy."

As evidence for these hypotheses, recall that there were 59 descriptions of mass-count terms that referred to objects made of those mass terms: 27 of these descriptions involved mass-animal terms and 32 descriptions involved mass-artifact terms. A group of undergraduate judges had determined the referents of these descriptions. In 96% (26 of 27) of the mass-animal terms, the animal term failed to retain its referential privilege. In 66% (21 of 32) of the mass-artifact terms, the artifact retained its referential privilege. For the mass-artifact terms, we have not systematically evaluated whether retaining versus giving up referential privilege corresponds to preserving versus violating an object's function. However, in those combinations in which the artifact term gave up its referential privilege, it did appear that interpreting them as "artifact made of mass term" would have violated their functions.

Summary

These preliminary results suggest that there may be a variety of mechanisms that operate in conceptual combination. Although we did find evidence for slot filling, it was by no means the only strategy that people used. Most notably, another very common strategy was aligning the structures of the two nouns and mapping part of the predicate noun's structure onto the structure in the head noun. Either a filler from the predicate noun was mapped to fill a slot in the head noun (property mapping) or a relation between slots in the predicate noun (or between a predicate slot and a slot in a related concept) was mapped to the head noun (structure mapping). In either case, the predicate noun is (in a sense) dismantled: instead of filling a slot in the head noun it yields part of its meaning in forming a combination. Finally, we also showed that people combine mass and count terms in complex ways. They often interpret a mass-count phrase as naming an object whose composition is indicated by the mass term but whose referent is not always a type of the head concept. In particular, if the head noun names a natural kind, it will lose its referential priveleges although it may contribute an important referential property (i.e., its shape) to the combination. If the head noun names an artifact, it generally retains its referential priveleges unless the composition of the combination violates the function of the artifact.

One possible objection to the present study is that we collected just a single definition for each combination and that our results reflect idiosyncratic responding in our subjects. For example, if we asked a large number of people what a pony robin was, would the majority actually respond "a robin with a tail" (as the subject in our experiment did)? Probably not. However, while a description for a given combination might be idiosyncratic, the different strategies discussed above were not rare occurrences. For example, property mapping and structure mapping were common strategies across subjects.

A second possible objection is that our results are based on unusual or even bizarre noun-noun compounds that one rarely encounters in typical natural language contexts. For example, how often does one encounter a phrase like "chair pony," consisting of a natural kind and an artifact? We have two answers to this objection. First, *individual constituents* of the compounds are not unusual. Therefore, studying how meanings of common words interact (even if their occurring in the same context is unlikely) could shed light on the nature of the meanings themselves. We believe that conceptual combination, like analogy and metaphor, forces words to reveal aspects of their meanings that may not become apparent in more usual contexts. In fact, the interactions between word meanings that we found suggest that the representations of individual constituents need to be structured and complex. Second, a number of the combination types that we used in this experiment (see Figure 5) do appear in our language, as indicated by lexicalized entries in the dictionary. Some examples of natural-kind pairs include: tiger salamander, sparrow hawk, moose bird, dog salmon, and gopher snake. Natural-kind artifact pairs include: whale boat, monkey jacket, book scorpion, carpet beetle, and oyster rake. Artifact natural-kind pairs include: trumpet flower, guitar fish, chimney swallow, razor clam, and pill bug. Mass-count terms include: paper knife, clay pigeon, stone fly, coal fish, and plastic bomb.

The findings raise a number of interesting issues. For example, we have suggested that slot filling is the default strategy for combining noun meanings. An obvious question is when do people adopt other strategies like property mapping and structure mapping? At this point, we can only speculate on the answer to this question. We can think of at least two conditions which might promote the use of property mapping and structure mapping. First, the more similar two objects are, the easier it should be to map properties or structures from one object to the other. In this context, similarity specifies the degree to which one frame can be *aligned* with another. In a combination such as *zebra horse*, high similarity may bias people to map properties from *zebra* to *horse*. Informally, when asked to describe a *zebra horse*, people typically responded, "a horse with stripes" (which is an example of property mapping). Although high similarity may facilitate the mapping of properties and structure, people may still interpret some combinations by slot filling. For example, two very plausible descriptions of *dolphin shark* are "shark with a dolphin-like nose" (property mapping) and "shark that eats dolphins" (slot filling). It may be that a highly salient, plausible relation between two objects can override mapping processes.

A second condition which might encourage mapping processes is the difficulty of finding a plausible relation between objects. As a result, people are not able to meaningfully combine nouns by slot filling (their default strategy) and must consider other strategies. For various reasons, a plausible relation between two objects may not exist. Factors like high dissimilarity of two objects and their low co-occurrence in the environment may rule out such a relation. This condition may apply to some of our examples of structure mapping (such as *snake glass* and *pony chair*).

There are many other interesting questions that future research might address. For example, at some level, can we reliably predict the meanings that subjects will construct for novel combinations? Can we reliably predict which combinations are more difficult to understand than others? Are some meanings of novel combinations "better than others" (as judged by subjects) and why? At this time, we have little to say about these issues. The major goal of the current work has been to determine the strategies that people use and the nature of noun representations required for those strategies. Once we have a better appreciation of these issues, we can begin to address the more difficult questions.

Acknowledgments

This research was supported in part by the National Institute of Child Health and Human Development under Fellowship Award 1 F32 HD07279-01 to the first author. It was also supported by the National Science Foundation under Grant BNS 87-20301 and by the Office of Educational Research and Improvement under Cooperative Agreement No. G0087-C1001-90 with the Reading and Education Center. We have benefited from discussions of this work with Arthur Markman, Douglas Medin, Gregory Murphy and Edward Smith.

Notes

[1] We use italics to indicate concepts or parts of concepts, and reserve quotes for the words that denote those concepts.

[2] We use the terms "head concept" and "predicate" concept to capture an important point. First, as discussed below, a noun-noun combination of the form *XY* typically refers to a *Y* and not an *X* — e.g., a dog apartment is an apartment and not a dog. Therefore, *Y* determines more of the meaning of *XY*. In this sense, *Y* is the head or main concept. In fact, one can view conceptual combination as an abstract function *X(Y)* in which *X* acts as an operator on *Y*. In this sense, it is a predicate.

[3] There are a number of subtleties about frames and frame instances that we will ignore in this paper. For example, the semantics of fillers are different in frames and frame instances. In a frame instance, a slot and its filler specify a fact that is actually true of a particular object. So, (robin-17 color red-13) roughly means that the color of a particular robin is a particular red. (Here, numbers are appended to robin and red to distinguish them from other instances

of robins and their instances of color.) In contrast, a slot and filler of a frame specify a default fact about an object. So, (robin color red) roughly means that the color of robins is typically red.

[4] Smith and Gray (1990) provide an alternative interpretation for some of these feature interactions. Specifically, they suggest that people already may be familiar with some combinations and that feature interactions do not result from a combination process but rather from experience with examples of the familiar combination. So, people may already be familiar with the combination *wooden spoon*, and may have acquired their belief that wooden spoons are large from experience with examples of wooden spoons. In this view, knowledge that wooden spoons are large is not derived from combining *wooden* and *spoon* but rather from examining examples of wooden spoons after the combination process.

[5] While acknowledging that their representations are much like frames, Smith et al. actually use the terms attribute and value for slot and filler and the term prototype for frame.

[6] Structure mapping is the mapping of relational structure from the predicate noun's meaning to the head noun's meaning (as in Gentner's (1983, 1989) discussion of analogy)

[7] We have assumed that the meaning of *pony* was accessed in order to interpret *pony chair* as a small chair. It is plausible that one may have accessed the meaning of *pony keg* instead. Nevertheless, we would claim that similar processing and representational assumptions would still hold. The basic difference would be that structure mapping would operate on a relation between slots of *pony keg* and *keg* instead of between slots of *pony* and *horse*.

References

Biederman, I. (1985). Human image understanding: Recent research and a theory. *Computer Vision, Graphics, and Image Processing, 32*, 29-73.

Brachman, R. J. (1977). What's in a concept? Structural foundations for semantic networks. *International Journal of Man-Machine Studies, 9*, 127-152.

Brachman, R. J. (1978). *A structural paradigm for representing knowledge.* BBN Report No. 3605, Cambridge, MA

Brachman, R. J. (1979). On the epistemological status of semantic networks. In N. V. Findler (Ed.), *Associative networks: Representation and use of knowledge by computers.* New York: Academic Press.

Brautigan, R. (1967). *Trout fishing in America.* New York: Dell Publishers.

Clark, E. V., Gelman, S. A., & Lane, N. M. (1985). Compound nouns and category structure in young children. *Child Development, 56*, 89-94.

Clark, H. H. (1983). Making sense of nonce sense. In G. B. Flores d'Arcais & R. Jarvella (Eds.), *The process of understanding language.* New York: Wiley.

Collins, A. (1978). *Human plausible reasoning*. BBN Report # 3810, Cambridge, MA.

Collins, A. & Michalski, R. S. (1989). The logic of plausible reasoning: A core theory, *Cognitive Science, 13*, 1-50.

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

Cottrell, G. W. (1988). A model of lexical access of ambiguous words. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology & artificial intelligence*. San Mateo, CA: Morgan Kaufmann Publishers.

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

Finin, T. (1980). The semantic interpretation of nominal compounds. *Proceedings of the First Annual National Conference on Artificial Intelligence*, Stanford, CA.

Gelman, S. (1988). Development of induction within natural kind and artifact categories. *Cognitive Psychology, 20*, 65-95.

Gentner, D. (1975). Evidence for the psychological reality of components: The verbs of possession. In D. A. and D. E. Rumelhart (Eds.), *Explorations in cognition*. San Francisco: Freeman.

Gentner, D. (1981). Some interesting differences between verbs and nouns. *Cognition and Brain Theory 4*, 161-178.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. Kuczaj (Ed.), *Language development: Language, cognition, and culture*. Cambridge: Cambridge University Press.

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 and A. Ortony (Eds.), *Similarity and analogical reasoning*. New York: Cambridge University Press.

Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology & artificial intelligence*. San Mateo, CA: Morgan Kaufmann Publishers.

Gentner, D., & Clement, C. A. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation*. New York: Academic Press.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1990). Relations, attributes, and the non-independence of features in similarity judgments, *Cognitive Psychology*, in press.

Halff, H. M., Ortony, A., & Anderson, R. C. (1976). A context-sensitive representation of word meanings. *Memory & Cognition, 4*, 378-383.

Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition, 15*, 55-71.

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

Hirst, G. (1983). *Semantic interpretation against ambiguity*. Technical Report CS-83-25, Computer Science Department, Brown University.

Kay, P., & Zimmer, K. (1976). On the semantics of compounds and genitives in English. *Sixth California Linguistics Association Proceedings*. San Diego: Campile Press.

Keil, F. (1986). The acquition of natural kind and artifact terms. In W. Demopoulos & A. Marras (Eds.), *Language learning and concept acquisition*. Norwood, NJ: Ablex.

Keil F. (1987). Conceptual development and category structure. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press.

Lees, R. B. (1960). *The grammar of English nominalizations*. Bloomington, IN: Indiana University.

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

Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23*, 250-269.

Marcus, M. P. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: MIT press.

Marr, D, & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, B 200*, 269-294.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 37-50.

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

Medin, D. L., & Wattenmaker, W. D. (1987). Category cohesiveness, theories, and cognitive archeology. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press.

Minsky, M. (1975). A framework for representing knowledge. In P. H. Winston, (Ed.), *The psychology of computer vision*. 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 and Language*, in press.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289-316.

Murphy, G. L., & Wisniewski, E. J. (1989a). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 572-586.

Murphy, G. L., & Wisniewski, E. J. (1989b). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Theory and applications*, (Vol. 2), 23-45, Chichester: Ellis Horwood.

Norman, D. A., & Rumelhart, D. E. (1975). Memory and knowledge. In D. A. Norman and D. E. Rumelhart, (Eds.), *Explorations in cognition*. San Francisco: Freeman.

Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition, 9*, 35-38.

Osherson, D. N., & Smith, E. E. (1982). Gradedness and conceptual combination. *Cognition, 12*, 299-318.

Palmer, S. E. (1975). Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction. In D. A. Norman and D. E. Rumelhart, (Eds.), *Explorations in cognition*. San Francisco: Freeman.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd, (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Rips, L. J., & Turnbull, W. (1980). How big is big? Relative and absolute properties in memory, *Cognition, 8*, 145-174.

Rosch, E., & Mervis, C. G. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573-605.

Smith, E. E., & Gray, K.C. (1990). *Mechanisms of conceptual combination*, Unpublished manuscript.

Smith, E. E. & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science, 8*, 337-361.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A modification model. *Cognitive Science, 12*, 485-528.

Smith, E. E., Shoben, E. J. & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81*, 214-241.

Talmy, L. (1976). Semantic casative types. In M. Shibatani (Ed.), *Syntax and semantics*, (Vol. 6). New York: Academic Press.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327-352.

Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General, 113*, 169-193.

Wisniewski, E. J. (1990). *Functional biases in understanding complex noun phrases*, Unpublished manuscript.

Zadeh, L. A. (1965). Fuzzy sets. *Infomation Control, 8*, 338-353.