Interpretive Semantics and Combinatorial Strategies in Adjective-Noun Conceptual Combination

Zachary C. Estes
University Undergraduate Research Fellow, 1995-96
Texas A&M University
Department of Psychology

APPROVED

Undergraduate Advisor ____________

Exec. Dir., Honors Program ____________

Abstract

Interpretive Semantics and Combinatorial Strategies in Adjective-Noun Conceptual Combination. Zachary C. Estes (Thomas B. Ward), Psychology, Texas A&M University

Conceptual combination is the merging of two separate entities to create an entirely novel concept. Examples range from everyday instances, such as computer desk, to challenging combinations, like upside-down sphere. We study concept combination to achieve a better understanding of language comprehension and extension, and also to investigate theories of concepts and concept interaction, which in turn have major implications for matters of knowledge representation. The present study determines what strategies we use to make sense of concept combinations under certain conditions, varied along the dimensions of relevance and typicality. Also included are oxymoronic (e.g., friendly enemy) and anomalous (e.g., cloudy enemy) combinations. As was expected, the technique of property mapping (attributing a property of one concept to the other concept) was dominant for adjective-noun combinations, but relation linking (exerting a relation between the two concepts) was also used to interpret the phrases. In general, relation linking was shown to occur most frequently under conditions of adjectival irrelevance and atypicality, and also in the anomalous group. In addition, explicit property negation (e.g., pagan marriage = a marriage that *doesn't take place in a church*) is introduced into the literature, and its suggestion of an interactive feedback process of concept interaction is postulated.

Interpretive Semantics and Combinatorial Strategies in Adjective-Noun Conceptual Combination

Introduction

People can make sense of language. Much of what we encounter in the form of language involves organizing and combining concepts into coherent units that properly convey intended meaning. The question of how it is that we understand conceptual combination has recently become one of interest to cognitive psychologists. In essence, the process of conceptual combination involves merging two previously existing concepts into one new entity. For instance, if prompted with the novel combination speaker jacket, one might interpret this to mean “a jacket with a radio speaker built-in,” “a jacket worn by a public speaker,” “a protective covering for a stereo speaker,” or even “a jacket that resembles a speaker.” As you can see, many interpretations are possible for any particular novel combination such as this one.

Studying how people combine concepts serves two main purposes. First, conceptual combination is one way in which we regularly extend the vocabulary of a language and modify the referents of our communication (Downing, 1977; Gerrig & Murphy, 1992). Technology and innovations are constantly providing us with new ideas and products, and we must extend or modify language in order to accommodate these new concepts. For instance, before the advent of the personal computer, there was no need for a table having the specific purpose of housing this nonexistent personal computer. Upon its arrival at the doorstep of every American, though, we had to first create the novel phrase personal computer, and then create a computer desk on which to set this new product. Here, it is clear to see how understanding the mechanisms used to interpret novel conceptual combinations is not only useful for language comprehension, but also language extension. A favorite example of mine is what I refer to as an Internerd (someone who spends way too much time on the Internet).

While this aspect of conceptual combination alone is worthy of investigation, another area of inquiry that is illuminated by the study of combined concepts further justifies such work. Concept combination probes the very nature of concepts by looking at their interaction. These studies can lead us to insights on the mental representation of concepts (Hampton, 1987; Markman & Wisniewski, 1996; Medin & Shoben, 1988). Several competing models of concept representation and combination have been proposed (i.e., Coolen, van Jaarsveld, & Schreuder, 1991; Hampton, 1987; Murphy, 1988; 1990; Shoben, 1993; Shoben & Gagne, in press; Smith, Osherson, Rips, & Keane, 1988) , but none have satisfied completely. More must be done to understand how concepts combine before we can claim to know how they are represented.

The task of understanding how conceptual combination works will not be an easy one. Many factors influence the process of combination, including context (Gerrig & Murphy, 1992; Murphy, 1990) and intraconcept relations (Medin & Shoben, 1988), among others. For instance, Medin & Shoben (1988) showed that altering the material of the concept spoon also changed the size of the spoon. They found that a wooden spoon is large, while a metal spoon was judged to be small.

Though much experimentation has been conducted on concept combination in the last decade and a half, most has focused on examining noun-noun pairs. Among the reasons for this is the assumption that, because these noun-noun combinations follow few set rules for how a concept interacts with another (Downing, 1977; Murphy, 1988; 1990; Wisniewski & Gentner, 1991), understanding noun-noun combinations will facilitate understanding of adjective-noun combinations. Many models further assume that adjective-noun pairs are interpreted by utilizing the same strategies that are used to comprehend and create noun-noun concept combinations (i.e., Gagne & Shoben, in press; Smith et.al, 1988). Consequently, these two assumptions--that understanding noun-noun combinations will facilitate development of theories of adjective-noun pairs, and that these adjective-noun pairs are interpreted by way of the same mechanisms used to comprehend noun-noun combinations--have resulted in neglect of examination of adjective-noun combinatorial strategies.

The Schema Approach

Some models of concepts and conceptual combination (Brachman, 1978; Cohen & Murphy, 1984; Smith et. al, 1988; Murphy, 1988, 1990) take a schema approach. In this view, concepts are represented as schemata or frames, which are lists of various dimensions in the concept--the basic knowledge we have of an entity. These schemata are composed of slots representing the dimensions of an object, and these slots are occupied by fillers, or values for the dimension (Minsky, 1975; Rumelhart, 1980). The schema for apple might have the slot SHAPE, which will usually be occupied by the filler round. Similarly, in this slot filling approach, the concept jeans would have the slot COLOR filled with the value blue in most instances. In fact, Cohen & Murphy (1984) showed that slots can contain default fillers, or values that are typical for a slot. When presented with the concept jeans, it is highly likely that the COLOR slot is filled with blue. As a result, presenting jeans alone (without modification by blue) will most probably lead to the assumption that they are blue, thus illustrating the idea of default fillers.

Some Variations on Adjective-Noun Combinations

Now that we are armed with an understanding of how concepts may be represented in the mind, let us progress to a description of the types of adjective-noun combinations that are used in the set of experimental stimuli for the current investigation (see Appendix). With respect to the head noun, an adjective may vary in its relation to the noun in several ways. The two dimensions of most interest to matters of this study are relevance and typicality. First, an adjective must be either relevant, irrelevant, or anomalous to the noun. A relevant adjective-noun combination is one in which the adjective is a filler for a slot that is present and salient in the schema of the noun to be modified. Blue apple is a relevant combination because the concept apple already contains a slot for COLOR, with blue being a filler for that slot. Cold garbage, on the other hand, is an irrelevant combination because the slot for TEMPERATURE is not typically thought of when presented with the concept garbage. This TEMPERATURE slot, though, can be inferred from general world knowledge. We know that garbage must have a temperature, we simply do not normally think about it or pay attention to it. Nevertheless, garbage must have a temperature. So an irrelevant combination is one in which the adjective is a filler of a slot that is not salient in the noun schema, but that can be inferred from world knowledge. Finally, an anomalous combination is the pairing of two words that basically do not fit together, even by inference from world knowledge. These combinations can be thought of as the most extreme case of irrelevance. A good example here is cloudy squirrel. Not only does the concept squirrel not have a salient slot for CLOUDINESS, it cannot even be inferred from knowledge of the world that a squirrel must either be cloudy or not cloudy. Squirrels seem to be neither cloudy nor non-cloudy. The slot just does not exist in the concept squirrel.

The other dimension of adjective-noun combination relational variability here of importance is typicality. This deals with how typical a particular adjective is of the noun it is modifying. An adjective is typical if the head noun is regularly or at least frequently thought of as possessing the property indicated by the adjective. To take an example from the set of experimental stimuli, a hostile enemy is a typical combination because enemies are hostile more often than not. An indifferent enemy, though, is atypical as a consequence of the fact that enemies are not normally apathetic. That is to say, atypicality is defined by a frequent absence of the property indicated by the adjective in the schema of the head noun. Third, just as anomalous combinations are an extreme version of irrelevance in adjective-noun pairs, oxymora can be viewed as an extreme type of atypicality. In these combinations, the adjective is more than just atypical--it seemingly contradicts the very essence of the head noun. For instance, friendly enemy involves two concepts which don't normally make sense together. In fact, friend and enemy are antonyms.

It is important to note that relevance and typicality are completely independent of one another. Any given adjective must be either relevant or irrelevant and also either typical or atypical for the specified noun (with the exception of anomalous adjectives, because they are irrelevant to the point of not having a typicality). So a relevant, typical pair is hostile enemy. The slot for LEVEL OF HOSTILITY is salient in the concept enemy, and furthermore enemies are frequently hostile as opposed to either indifferent or friendly. Indifferent enemy, on the contrary, is a relevant, atypical combination because, as mentioned above, the slot pertaining to LEVEL OF HOSTILITY is a salient one for enemy, and it is rare that such an enemy is apathetic toward its opponent.

Another potentially helpful way to think of this is to use the idea of default fillers for the slot being modified. Once again, what is meant by the term 'default filler' or its synonym 'default value' is a value that is most common for that slot or that is expected and assumed to fill that slot. For instance, the concept apple has as a default filler for the SHAPE slot the value round. Apples are normally round, and so we assume that any apple mentioned is round unless otherwise stated.

Returning to the issue of relevance and typicality, further explanation might be useful. First of all, if the slot in question is present and salient in the noun concept, then the adjective is relevant to that noun. Second, if a default value for the slot in question is the same as or similar to the property asserted by the adjective, then that adjective is typical for the head noun. If the property to be attributed to the noun is not a default value, then it is an atypical adjective for that noun. Some slots, however, seem to be constrained as to what values may be accepted (Brachman, 1978). For example, the slot for MEANS OF LOCOMOTION in the concept plant is constrained in the sense that it not only has a default value of immobile, but actually also cannot accept the value feet (because the concept plant has as a defining feature immobility, and to give it mobility is to deny it plant-hood). In cases such as this one, if the property to be attributed by the adjective is a value that is not only not a default value for that slot, but even further is a value that cannot be accepted by the noun, then the adjective is oxymoronic to the noun. An example of this is, again, friendly enemy. As we will see later in this paper, there obviously exist ways to get around apparent contradictions like this (as demonstrated by the fact that people can in fact understand oxymora (Gibbs & Kearney, 1994; Shen, 1987)).

If the slot to be modified is not salient in the noun concept, but the value indicated by the adjective is typical for that noun, then the adjective is irrelevant, typical. Though enemies are normally clothed, we do not usually think of whether or not the enemy is wearing clothes. The slot for CLOTHEDNESS is not salient in the concept enemy, yet we know that if we do bring this slot into focus, it will most assuredly be filled with the value dressed. As a result, the combination dressed enemy is irrelevant, typical. On the contrary, undressed enemy is an irrelevant, atypical pair. The CLOTHEDNESS slot is not salient, and when brought into salience, is not normally filled with the value undressed. Undressed is atypical of enemies because the default value for CLOTHEDNESS is dressed. Finally, an anomalous combination is one in which the slot to be modified is so extremely irrelevant that it does not seem to belong with the head noun. The value indicated by the adjective cannot locate a slot in the noun to fill. And because the appropriate slot cannot be located, there is no default value to be found either. There is no way to judge typicality in pairs like this since the slot does not exist and there is no default value to which we could compare the property to be attributed. One example of an anomalous phrase is cloudy enemy. People are challenged to find the CLOUDINESS slot in the concept enemy, and the task proves exceedingly difficult.

Combinatorial Strategies

Now that we have been introduced to all six types of adjective-noun concept combinations used in the present study, we will review the literature for techniques that have been identified for generating sense when confronted with novel word pairs. Psychologists have turned to linguistics for assistance in deciphering this problem of sense generation. Historically, linguists have viewed the task of interpreting novel combinations as involving a taxonomy of abstract relations (Kay & Zimmer, 1976; Gleitman & Gleitman, 1970; Levi, 1978). Levi (1978), for instance, identified sixteen relations used to make sense of these nominal compounds. Among these are the CAUSE relation (e.g., stress headache), the MAKE relation (e.g., auto worker), and the FOR relation (e.g., flower pot). Also, some relations outlined by Levi (1978) have more than one dimension. The IN relation is one such abstraction. To illustrate, a house guest exemplifies the concrete location type of the IN relation, while the summer months involve the temporal aspect of IN.

Recently, Shoben (1993) and his colleague (Shoben & Gagne, in press) have extended this relational taxonomic approach to the psychological research of non-predicating combinations, which include instances of both adjective- and noun-noun pairs, though not all adjective-noun combinations are non-predicating (for instance, white does predicate paper, but does not predicate mystery). This model directly borrows the abstract relations defined by Levi (1978) and augments her taxonomy with the addition of a number of other thematic roles, such as a BY relation (student vote).

Shoben (1993; Shoben & Gagne, in press) has claimed that all non-predicating combinations are interpreted by applying these thematic relations. Other studies, however, have shown the existence of at least two other combinatorial strategies (Wisniewski, in press). Property mapping occurs when one or more attributes of one concept are extracted and then mapped onto the other concept. An example here would be the meaning "a square clock" for the combination box clock. A more extreme and less frequent mechanism utilized for sense generation in concept combination is hybridization. That is, when faced with the task of interpreting a novel combination, we simply fuse the two concepts together. For instance, a skunk squirrel is an animal that has properties of both a skunk and a squirrel.

Wisniewski (1996), though, has since rejected the notion of property mapping in favor of a very similar process he calls property construction. The difference, though subtle, is an important one. He now asserts that properties are not directly mapped, but rather a new property is constructed in the head noun. In the case of a zebra clam, the stripes of a zebra cannot be mapped directly onto a clam because they would be too wide, too long, and too smooth, etc. to fit well on a clam. Instead, the property of striped guides the construction of a new property in the head noun, which has constraints that are different from those of the modifier. In effect, the stripes of a zebra must be "mentally shrunk" to fit the specifications of a clam, and so are not direct mappings. These additional techniques for understanding combinations, though, are intended only to explain noun-noun pairs. The abstract thematic relations mentioned previously stand as the single mechanism posited for the comprehension of adjective-noun combinations.

Purpose

Accordingly, it is my intention to show that, in addition to relation linking, other strategies are employed during the process of interpreting adjective-noun combinations. I will test the application of noun-noun combinatorial techniques, such as property mapping, to adjective-noun pairs, and I will probe for entirely novel strategies used to make sense of these combinations. Also, I will determine which techniques are most likely to be employed under certain conditions of relevance and typicality. Perhaps most importantly, though, I will analyze the findings of this study as well as others to uncover the implications that combinatorial strategies bear on theories of concept interaction.

TYPICALITY RATING TASK

The purpose of the first study was to develop and test a set of experimental stimuli ultimately acceptable for the conceptual combination task. Functionally, this task is actually more like a pretest than an experiment in and of itself, though it does require a unique experimental design and procedure. As is explained in greater detail below, nouns were selected for use in the task, and adjectives that were thought to be typical or atypical and relevant or irrelevant in relation to the head noun were then arbitrarily chosen. For instance, if the noun being modified is enemy, one may examine this concept and select two slots for modification. The first slot must be a salient one in the concept, such as

LEVEL OF AGGRESSION. To fill this slot, one must further anticipate an adjective that will be ranked typical of that noun. Here, hostile was chosen. Another adjective chosen for inclusion in the typicality rating task was indifferent because it was expected to be rated atypical if hostile is typical. Then, an irrelevant slot, or one that is not salient in the noun concept, was identified. The selected slot was CLOTHEDNESS. I predicted that enemies would be rated dressed as typical, and undressed as atypical. In this sense, the typicality rating task does make specific hypotheses about the ratings of each combination. The typicality rating scale went from -2 (very atypical) to 2 (very typical). Accordingly, any mean typicality rating for a combination that is less than zero will be deemed 'atypical,' and any mean rating greater than zero will be 'typical.'

Admittedly, these combinations are arbitrary. This, however, is inconsequential because, again, the goal of this task is simply to verify my intuitions of which pairs are typical and which are atypical. After such verification, the set of stimuli will be proven to represent what is claimed, and its arbitrary nature will become impertinent.

Method

Participants

Eight students in an undergraduate class at Texas A&M University volunteered to participate. These participants were rewarded only with a greeting of appreciation from the experimenter.

Materials

Each volunteer was given a form with written directions, a rating scale, and forty-eight concept combinations. The rating scale ranged from -2 (very atypical) to 2 (very typical), with 0 meaning a neutral typicality rating. Concerning the combinations, twelve head nouns were used in the task, with each one being combined with four different adjectives (relevant, typical; relevant, atypical; irrelevant, typical; and irrelevant, atypical). None of the adjectives were paired with more than one noun. In other words, each of the twelve nouns appeared four times, while each adjective was used only once. Thus, there were forty-eight adjectives in the typicality rating task. The adjectives were matched roughly for frequency of use in the English language using the analysis of Francis & Kucera (1982) to control for familiarity effects.

Of the twelve head nouns, ten were borrowed from the list of direct oxymora contained in Gibbs & Kearney (1994). These oxymora were created in the following manner, as stated by the authors: "We simply selected..pairs of antonyms from *Webster's Dictionary of Synonyms*, such as..living-dead. The second adjective terms were changed into nouns so that, for example, the antonyms living-dead became living-death" (p. 78). These ten oxymora taken from Gibbs & Kearney (1994) were then broken down into their adjective and noun components (i.e., living and death), and the nouns were used in the typicality rating task. The adjectives were not included because the purpose of this task is to determine typical and atypical adjectives for each noun, and we already know that those oxymoronic adjectives are on the extreme end of atypicality. The oxymoronic and anomalous adjectives need not and do not (respectively) fit into this typicality rating scheme. The other two nouns (sky and illness) were extracted from two oxymora constructed by the experimenter using the same technique as described above.

The adjectives in this task were chosen somewhat arbitrarily, but this is not problematic because the task is a pretest. Its function is to determine whether or not the expected (a)typicality is found for each combination so that the most appropriate combinations may be chosen for the actual experiment. To choose the adjectives, I investigated the twelve noun concepts, selected both relevant and irrelevant slots in those concepts, and then finally chose both typical and atypical adjectives. Hence, four adjectives for each of the twelve nouns were selected, yielding a total of forty-eight combinations. The order of the combinations was random.

Procedure

The volunteers were given the Typicality Rating Task sheet, which included the following directions:

> Below you will see many word pairs consisting of an adjective and a noun. I am interested in finding out how typical the particular adjective is of the specified noun. For instance, if the word pair 'green leaf' was presented, you should answer that the adjective 'green' is very typical of the noun 'leaf.' I am not at all concerned with how often you hear the two words together in daily life; I only want to know if the adjective is usually true of the noun or not. Another example here is 'blue carrot.' You should respond that carrots are not normally blue, or in other words, that 'blue' is atypical of the noun 'carrot'...

The participants then rated each of the forty-eight combinations. The task took an average of approximately ten minutes.

Results

Of the forty-eight combinations, the mean typicality ratings of nine of them violated the expected typicality values. Only two head nouns were involved in prediction violation in more than one instance. That is, the nouns shelter and war both were paired with two adjectives that were not rated in the direction of typicality that I expected them to be. All combinations having either shelter or war as the head noun were consequently thrown out and excluded from the mean typicality ratings listed hereafter, as well as being excluded from the set of actual experimental stimuli used in the conceptual combination task later. This exclusion reduced the total number of nouns to ten. With the exception of combinations with these two nouns, only five violated the predictions. Furthermore, only two of these unexpected ratings are substantial discrepancies. Depressing marriage, which I had predicted would be atypical (have a mean typicality rating of less than zero), was judged to have a mean typicality rating of 0.4. This mean typicality rating is still nearer neutrality than typicality. The other substantial deviation was a mean typicality rating of -0.6 (atypical) for the combination dressed enemy, which, based on general world knowledge, warranted a prediction of 'typical' rather than 'atypical.'

The mean typicality ratings (see Table 1) supported my general predictions of which adjectives were typical or atypical for the given noun. The strongest rating was 1.70 for the relevant, typical group. What is important here is that, regardless of relevance, the mean typicality ratings for the typical and atypical conditions were in the correct direction.

Table 1

Mean typicality ratings for combination groups in Typicality Rating Task.

Group	Typical	Atypical
Relevant	1.70	-0.78
Irrelevant	0.78	-1.17

Note. Means do not include ratings of combinations with either shelter or war as the head noun. The scale ranged from -2 (very atypical) to 2 (very typical). Any positive mean is rated 'typical' and any negative mean indicates 'atypical' ratings.

Discussion

To begin, it comes as no surprise that there were not only five deviations from predicted typicality, but also that two of them are substantial deviations at that. Murphy (1990) also found that participants had difficulty judging typicality and relevance when rating combinations of this sort. Consequently, his results were not ideal either. People simply do not understand the task. As can be seen in the directions listed above, participants were indirectly asked to pay attention to such difficulties and to be careful and thoughtful in rating the typicality of a pair. Nevertheless, participants had trouble. For instance, dressed enemy, the most striking violation of my predictions, was given a mean typicality rating of -0.6 (atypical). Accessing world knowledge will surely lead one to realize that when we encounter an enemy, he/she is almost always clothed in some manner or another. Perhaps participants rate the pair atypical because the slot for clothedness is not salient in the noun schema (the adjective is irrelevant) and thus the combination is extremely novel. This, of course, is speculation, and the fact remains that my predictions were not all correct. Furthermore, five combinations that were rated slightly atypical will be included in the conceptual combination task as typical pairs, or vice versa. A cursory glance at the literature, however, shows that such experimental problems are not unusual, and even possibly that they are not real problems with the experiment after all, but rather just difficulties in understanding the task. Perhaps it is best stated in the words of Glucksberg, McGlone, & Manfredi (in press) who, in reference to the fact that people agree that planets are never framed yet have trouble understanding the phrase unframed planet, said “existential possibility is often less important than plausibility.”

ATTRIBUTE LISTING TASK

The purpose of this task was to compile an extensive list of features for each of the concepts to be used in the conceptual combination task. Once more, this is in effect merely a pretest for later use in analyzing the combinatorial strategies used to interpret the experimental combinations. The task makes no hypotheses.

Method

Participants

Fifty-four students from the Introductory Psychology subject pool at Texas A&M University volunteered to participate in the study in exchange for course credit.

Materials

Two separate packet types were constructed for the attribute listing task. Each of the two packet types consisted of a total of thirty-five individual concepts (five nouns and thirty adjectives) distributed across eight sheets of paper. Thus, a total of ten nouns and sixty adjectives were included. Eight of the ten nouns were taken from the direct oxymora stimuli used by Gibbs & Kearney (1994), as described in the Typicality Rating Task section. The other two nouns (sky and illness) were selected in accordance with the technique espoused by those authors for the creation of direct oxymora. Ten of the sixty adjectives were those that were oxymoronic in relation to the nouns used (again, see Typicality Rating Task section for greater detail). Forty of the remaining fifty adjectives were those that were arbitrarily selected for (a)typicality and (ir)relevance purposes and that were chosen as a result of the Typicality Rating Task. The final ten adjectives were selected for purposes of creating the anomalous combinations. The ten nouns were analyzed for slots that simply did not exist in the concept, and on the basis of this analysis, ten adjectives were chosen that are filler values for the non-existent slots in the nouns. None of the concepts appeared in both packet types. The concepts were grouped together according to part of speech--the five nouns were first, and the thirty adjectives followed. The order of the concepts was random within the part of speech groups. The determination of which concepts would be included in which packet was random, with the exception that the oxymoronic adjective for any particular noun was not placed in the same packet type as that noun. Ample space was provided between each concept to list as many attributes as desired.

Procedure

Twenty-seven participants each received one packet type, and twenty-seven different volunteers each were given the other packet. Participants were instructed to list attributes for each concept, and an example of such attribute listing was given. The volunteers were requested to attempt to list six attributes for each noun concept and three attributes for each adjective, though they were also informed that these guidelines were not rigid. After all questions were satisfactorily answered, they began the task. The task lasted approximately forty-five minutes, and participants were not allowed to leave until all students were finished.

Results and Discussion

This phase of the present experiment yields no statistical results. Rather, its significance lies in its ability to resolve questions arising during the later process of analyzing the combinatorial strategies employed by participants to generate sense of the combinations. If the coder is uncertain whether a particular property of the combination was an attribute of one of the component concepts, for instance, this index of features will answer the question. If the property of the combination was in fact included in the attribute list of one of the individual concepts, then the property mapping technique has most likely been utilized.

Participants were asked to list more attributes for the nouns than for the adjectives on the basis of findings by Gentner (1981) and Gentner and France (1988). These findings indicated that noun concepts are more structurally and functionally rigid as compared to verbs and other predicates, such as adjectives. If nouns are more structured, it was assumed for this task that they would contain more set, identifiable concrete features, while the adjectives would have more abstract relations and less concrete features to discern (as explained by Murphy (1990)).

CONCEPTUAL COMBINATION TASK

Finally, this task was the basis or goal of the entire experiment. All else up to this point has been peripheral. The present task, in contrast, is the foundation of the empirical research. The purpose of this task was to determine (1) whether the combinatorial strategies postulated in the noun-noun literature apply to adjective-noun combinations, (2) exactly what techniques are used if not those employed to interpret noun-noun combinations, (3) which strategies are used most frequently in any given situation (varied along the dimensions of relevance and typicality), and (4) what mental processes are involved in comprehending and creating combined concepts.

Despite the fact that this investigation is largely exploratory in nature, several predictions are necessary at this point. We established earlier that three combinatorial strategies have been posited in the literature, though two of these are meant exclusively for noun-noun combinations. First, property mapping or property construction is the extraction or construction of a property from one concept, with that property then being applied to the other concept. A zebra clam is a clam with stripes. This notion of property construction as opposed to mapping is a new one (i.e., Wisniewski, 1996), and as such has not yet been thoroughly tested. Also, I find that the term "property mapping" is simpler to understand. As a result, "property mapping" will be used in place of "property construction" in the remainder of the current manuscript. It should be noted, however, that any instance of property mapping could actually be the related technique of property construction. For purposes of this experiment, though, this distinction will be ignored.

Next, hybridization is an extreme version of property mapping in which a new entity inherits properties of both concepts. A skunk squirrel, again, is "an animal having properties of both a skunk and a squirrel." A real-world example of this strategy is the hatchet--a hatchet is a blend of a hammer and an axe, resulting in a tool that contains both of the individual concepts at the same time. Concerning adjective-noun combinations, hybridization does not appear to be plausible due to the incongruence of the structure of noun and adjective concepts (see Gentner, 1981; Gentner & France, 1988; Murphy, 1990 for arguments of this incongruence).

The third proven technique for generating sense of concept combinations is relation linking. This has been studied mostly by linguists in the past (e.g., Downing, 1977; Levi, 1978), but recently psychologists have emerged on the scene in this domain too (e.g., Coolen et.al, 1991; Shoben, 1993; Shoben & Gagne, in press). These researchers have identified several types of thematic relations that people use to interpret combined concepts, such as the CAUSE relation and the ABOUT relation (e.g., tax law).

First, an immense proportion of property mapping is expected to occur in the relevant, typical combination group (e.g., hostile enemy). Precisely because the adjective concept is relevant and typical, it is a predicating combination. A large X is "an X that is large," to give an example of predication. It is interesting to note, too, that the property of largeness is understood in both large elephant and large mouse, though the largeness is not the same in these two combinations (Halff, Ortony, & Anderson, 1976). With the example of hostile enemy, the adjective concept consists of the slot LEVEL OF AGGRESSION, which is filled with the value hostile. Because this slot is highly salient in the concept enemy, and because the filler hostile is highly typical of the concept, we can expect the property of the adjective to be mapped onto the noun concept. In other words, I predict that these combinations will be interpreted by transferring a filler value from the adjective concept to a slot in the head noun--property mapping.

Next, Murphy (1990) showed that atypical adjective-noun combinations took longer to understand than typical ones. The additional time required to comprehend the atypical pairs seems to imply that simple property mapping is not obvious. On the basis of this, I hypothesize that more relation linking will be used to interpret relevant, atypical combinations (e.g., indifferent enemy) than was the case with relevant, typical pairs. I still expect, however, to find much more property mapping than relation linking within the relevant, atypical group.

Analogously, I predict that more relation linking will be identified in the irrelevant, atypical group (e.g., undressed enemy) than in the irrelevant, typical one (e.g., dressed enemy). In other words, atypical combinations should result in more frequent use of relation linking than typical pairs, regardless of adjectival relevance. Once again, though, I predict that property mapping will be more common than relation linking within the irrelevant, atypical condition. Also, the irrelevant, typical group should exhibit primarily only property mapping, just as is predicted for the relevant, typical pairs. In summary, for these four groups (relevant, typical; relevant, atypical; irrelevant, typical; and irrelevant, atypical), hypotheses are that (1) property mapping will be significantly more frequently used than relation linking in each of these groups, (2) atypicality will increase the number of interpretations via relation linking, regardless of the relevance of the adjective, and (3) irrelevance will likewise increase the frequency of relation linking, regardless of adjectival typicality.

Turning now to the group of oxymora (e.g., friendly enemy), relatively few studies can be found on such contrasting concept combinations. The approach used by most of these researchers involves the domain of social concepts. For instance, Hampson (1990) investigated behaviors attributed to the conjunction of traits that were either congruent (e.g., unsociable and unfriendly) or incongruent (e.g., thorough and haphazard). Other studies (e.g., Hastie, Schroeder, & Weber, 1990; Kunda, Miller, & Clare, 1990) have combined congruent (e.g., Republican bank teller) and incongruent (e.g., Republican social worker) social stereotypes. All of these yielded similar findings--that people tend to try to explain or justify the contradiction. Gibbs & Kearney (1994) summarized that

> "understanding contradictory statements, such as clean pollution, appears to require readers to draw some causal relationship between the first and second terms in just the same way that we often make sense of social concepts that have conflicting implications..by creating some causal theory as to why two different concepts may coexist" (p. 86).

These studies are helpful in predicting the outcome of oxymoronic combinations. If people create causal reasoning to comprehend the conjunction of social categories, we can expect to find similar comprehension processes in simple oxymora. This justification for apparent contradiction seen in social conjunctions leads me to expect to find much relation linking in the contradictory group of the present study. The CAUSE relation described earlier might be prominent among combinatorial techniques for oxymora. Also, because the oxymora are designed to self-contradict, it seems as if one would be hard-pressed to be able to coherently map the adjective onto its noun antonym. Accordingly, I hypothesize that relation linking will be used in the oxymora group more than in the other groups, and that within the group itself, relation linking will be more frequent than property mapping.

Finally, no in-depth research has been done on anomalous combinations (e.g., cloudy enemy). One study (Lukatela, Carello, Kostic, & Turvey, 1988), though, did use a priming paradigm on a lexical decision task involving either congruent or incongruent adjective primes for the noun target. This experiment found that lexical decision for congruent situations (e.g., good aunt) was facilitated, while the incongruent pairs (e.g., slow coat) did not differ from the baseline. Good aunt is equivalent to my irrelevant group most likely, and slow coat is clearly an anomalous phrase. This, like Murphy's (1990) finding that atypical adjectives require more time to comprehend, can be interpreted to mean that a simple property mapping is not found. In the case of slow coat and other anomalous combinations like it, I predict that little property mapping occurs. Instead, either relation linking or some sort of metaphor will be implemented to generate sense of the combination. I would like to reiterate that this experiment is more exploratory than predictive.

Method

Participants

One-hundred fifty students from an introductory psychology course at Texas A&M University volunteered to participate. Each received course credit for participation.

Materials

A total of six different experimental packets were constructed. Each packet consisted of ten conceptual combinations from one of the six combination groups (relevant, typical; relevant, atypical; oxymora; irrelevant, typical; irrelevant, atypical; or anomalous). None of the packets included concept combinations from more than one combination group. In other words, one packet contained ten word pairs, all of which were from the same combination group. Each packet was five pages in length, with two combinations on each page. No packet contained the same concept more than one time. That is, no adjective or noun appeared in the same packet more than once. Also, no adjective was used in more than one packet, though each packet contained the same ten nouns. So only the adjective varied from packet to packet.

As described above in the typicality rating task and attribute listing task, eight of the oxymora were taken from Gibbs & Kearney (1994), and the other two were formed by the experimenter in the same manner. Forty of the remaining combinations were selected from the typicality rating task. The final ten combinations, which were the anomalous ones, were combined on the basis of an examination of the noun concept for slots that do not seem to exist. A value was then selected to fill this slot. The seventy concepts used (sixty adjectives and ten nouns) were the same as those in the attribute listing task.

Procedure

Participants were each given one packet of stimuli. Each of the six types of packets was completed by twenty-five students. They were informed that the combinations may be novel to them, and they were instructed to simply write down a definition that most likely describes the combination. Examples were given and questions were answered to satisfaction. The participants then defined each of the ten combinations at an individual pace. After all definitions had been completed, I read them one more set of directions asking them to now list attributes of the conceptual combinations. They were asked to imagine that they were trying to describe what the combined concept was like to someone who did not know what it was. Again, questions were answered and examples were given. The purpose of this second phase was simply to clarify the meanings of the definitions. Many ambiguous definitions were anticipated, and the inclusion of additional attributes would prove to be of great benefit for analyzing these definitions later. When the last participant finished, the volunteers were debriefed and released. The entire session lasted an average of approximately forty-five minutes.

Scoring

Due to the open-ended nature of this experiment, the coding of these results was a difficult process. A total of one thousand, five hundred definitions were collected, with several attributes for each one. A single rater initially analyzed each of the definitions. For any definition that was not clear as to the combinatorial strategy used, a second rater was asked to interpret the definition. For such a definition, the two raters either reached an agreement through discussion, or if an agreement could not be reached, the definition was placed into an "other" category. Few of the definitions, though, required such measures.

Most of the definitions, which were easily interpretable, were analyzed according to the following process. To begin, each of the definitions was broken down into its component parts. From here, the relation between the components was examined. If either of these two steps was ambiguous, the attribute list for the combinations compiled in the second phase of the conceptual combination task was consulted to achieve a higher sense of understanding for the definition. Also, if it was unclear whether a component of the definition was a representation of one of the concepts, the attribute list for the individual concept from the attribute listing task was reviewed. For instance, if the combination painful joy was defined as "something that feels good," the individual attribute list for painful was scanned to verify that no such attribute was listed for the concept (combinatorial strategies such as this were labeled "concept exclusion," which is explained in greater detail shortly). After all definitions had been quickly scanned by the first rater, it was determined that many of the responses did not correspond well to either relation linking or property mapping. In order to accommodate this realization, some categories of combinatorial strategies were created ex post facto. All categories of strategies will now be described.

Though the entire experiment was based on and geared toward definitions, many participants took it upon themselves to avoid the task in one of three ways. The first and most common way that people evaded defining the combination was categorized as "instantiation." The technique of instantiation was to simply list a probable examplar of the combination, rather than defining the combination. Hostile enemy drew one such response of "my ex-boyfriend." Instantiations may be useful and helpful in everyday life, but to attempt to analyze the relation between hostile and enemy here would be speculative to say the least. As a result, instantiations, as well as the other two strategies of definition avoidance, were not analyzed further into relations between concepts (i.e., property mapping, relation linking, hybridization, etc.).

Another way to escape the requested task of defining combinations was called "concept exclusion," as described above. This occurred when one of the component concepts was completely excluded from the description. In effect, one concept was just repeated or reworded, while the other was ignored. The example of "something that feels good" for painful joy perfectly illustrates this strategy. Painful is left out of the description of the combination.

The last strategy for definition avoidance was "attribute listing." Rather than defining the combination, some participants decided to compile a list of properties to describe it. These students listed attributes not only in the second phase of the experiment, but apparently also in the first. No definition could be found in these cases. One participant, for instance, stated the following for satisfying marriage: "compatibility; love; friendship; partnership; sexual relationship healthy; fun; able to communicate well; fidelity; comfort; acceptance of mutual flaws/shortcomings."

Now, recall that the second step was to examine the relation between the components of a definition. This is the crucial, distinguishing stage of the analysis. A definition was categorized as "property mapping" if a representation of one concept simply possessed at least one property of the representation of the other concept. What is meant by representation of a concept in this situation is what the participant described or reworded the concept as. For example, enemy might be represented as "someone who opposes you." To extend this example, the definition "someone who opposes you and is aggressive toward you" for the combination hostile enemy would be labeled "property mapping." The head noun concept possesses a property of the adjective concept in this case.

In other cases, the head noun concept possesses a property of the adjective concept only under certain conditions. In this respect, the strategy in question is distinct from regular property mapping. We cannot just say the concept X has property y (from concept Y), because it is not always true that X has y. Accordingly, definitions using this combinatorial technique were categorized as "conditional property mapping." This strategy is considered to be a subset of the property mapping strategy rather than an entirely separate technique.

Next, definitions in which a property of one of the component concepts is nullified are labeled "explicit property negation." For instance, acceptable lie was defined as "a nontruth which has no negative stigma associated with it." From examining the attribute list for the individual concept lie, it is clear that a lie is considered negative. And in this combination, the term acceptable is easily construed to mean "no negative stigma." So in the example, the property negative of the concept lie is directly canceled out, or negated. A similar phenomenon, "implicit property negation," was also identified, though it warrants no category of combinatorial strategy in and of itself (for reasons discussed later) separate from regular property mapping. Explicit property negation, like conditional property mapping, is considered only a subset of property mapping (also for reasons discussed later).

"Relation linking," on the contrary, is a completely distinct and separate strategy from any sort of property mapping. In relation linking, a bridge is formed between the two concepts involved. A definition of "a death caused by something horrible" for tragic death utilizes the CAUSE relation. For purposes of this experiment, however, the precise relation used was not determined--only that some relation, be it CAUSE, IN, or something else, was used.

Because relation linking and property mapping are the two most prominent combinatorial strategies in the literature, it may prove beneficial to expand on the difference between the two for the sake of clarity. A simple way to discern the two is to look at how many entities exist in the definition. If it has two concept representations, it is using relation linking because in order to link something, there must be a minimum of two entities to be linked. One cannot link something to nothing. So relation linking requires a minimum of two entities. Property mapping, on the other hand, results in only one entity in the definition. One or more properties is mapped onto the other concept. Interestingly, hybridization (e.g., a skunk squirrel is "an animal with properties of both a skunk and a squirrel) results in one entity as well, but differs from property mapping in that the one entity is not one of the original concepts, but an entirely new concept instead. A skunk squirrel is neither a skunk nor a squirrel, but is actually a new animal with properties of both.

Finally, each definitional attempt in this task was categorized as one of these eight combinatorial strategies (instantiation, concept exclusion, attribute listing, property mapping, conditional property mapping, explicit property negation, relation linking, or other).

Results

Property mapping

One participant apparently did not understand English well enough to do the task, and his responses were excluded from analyses. This left twenty-five participants in each group, with the exception of the group to which this participant belonged, which ended up with only twenty-four. For purposes of statistical analysis, ordinary property mapping, conditional property mapping, and property negation were all combined to form one category of property mapping. To determine the results of the within-category analyses, t-tests for paired samples were used. All but the anomalous group use significantly more property mapping than relation linking. As the means in Table 2 indicate, the relevant, typical and the relevant, atypical pairs were interpreted almost exclusively by property mapping, $t(24) = 15.22$, $p < .001$ and $t(24) = 12.73$, $p < .001$, respectively. The contradictory combinations were also interpreted by property mapping significantly more frequently than by relation linking, $t(24) = 6.35$, $p < .001$. For the irrelevant groups, both typical [$t(24) = 13.97$, $p < .001$] and atypical [$t(23) = 7.17$, $p < .001$] groups resulted in significantly more property mapping than relation linking. The only result that did not reach significance, that of the anomalous group, is actually a trend in the predicted direction--more frequent relation linking in this group.

Next, a series of 2 X 2 (relevance X typicality) ANOVAs were run to determine effects of these variable dimensions on the different combinatorial strategies. For property mapping, there was a significant main effect of relevance, $F(1, 95) = 4.386$, $MSe = 3.778$, $p < .05$. More mapping was found in relevant than irrelevant combinations (see Table 2 for means). Similarly, typicality showed a significant main effect on this combinatorial technique, $F(1, 95) = 4.386$, $MSe = 3.778$, $p < .05$, with typical combinations using more mapping than atypical pairs. The (relevance X typicality) interaction was not significant.

A one-way ANOVA revealed reliable differences in the prevalence of property mapping in relevant combinations along the dimension of typicality (including oxymora), $F(2, 72) = 6.67$, $MSe = 3.53$, $p < .01$. Consistent with the predictions, less property mapping was used in the oxymora group ($m = 5.84$) than in either the typical group ($m = 7.72$) or the atypical group ($m = 7.20$), both $p < .05$ by Tukey's Honestly Significant Difference. The difference between relevant, typical and relevant, atypical was not significant. A one-way ANOVA of the effect of relevance (including the anomalous condition) on the use of property mapping showed analogous results, $F(2, 121) = 12.67$, $MSe = 3.98$, $p < .001$. For this measure, relevance was collapsed across typicality because the dimension of typicality cannot be determined for combinations, such as the anomalous ones, that use an adjective that does not coherently modify the noun. Here too, Tukey's HSD revealed significantly less property mapping by the anomalous group ($m = 5.00$) than in either the relevant condition ($m = 7.46$) or the irrelevant condition ($m = 6.65$), both $p < .05$. There was no significant difference between relevant and irrelevant groups. (Note:

for analyses that have been collapsed across a dimension, the means indicated above will differ from the means of individual values along that dimension in Table 2.)

Relation linking

As stated above, with the exception of the anomalous group, property mapping was significantly more frequent than relation linking in each of the groups. In the anomalous condition, there was no significant difference between the use of the two techniques. To analyze the between-group trends, 2 X 2 (relevance X typicality) ANOVAs were used. There were significant main effects of both relevance [$F(1, 95) = 16.33$, $MSe = .97$, $p < .001$] and typicality [$F(1,95) = 6.89$, $MSe = .97$, $p = .01$] on relation linking in the directions predicted. Again, the (relevance X typicality) interaction was not significant.

In examining the effect of typicality on the frequency of relation linking in relevant combinations (including oxymora), a one-way ANOVA showed a significant difference between the oxymora group and the typical group, $F(2, 72) = 9.60$, $MSe = 1.21$, $p < .001$, as hypothesized. However, there was no significant difference between either the oxymora ($m = 2.00$) and the atypical groups ($m = 1.28$), or the typical condition ($m = .64$) as compared to the atypical condition, again utilizing Tukey's HSD. Another one-way ANOVA probing the effect of relevance (including the anomalous group) on relation linking exhibited reliable differences in the frequency of this strategy between the anomalous condition and the relevant condition, $F(2, 121) = 31.22$, $MSe = 1.70$, $p < .001$, as predicted. Once more, for this analysis relevance was collapsed across typicality. Tukey's HSD showed that relation linking was used significantly more in the anomalous group ($m = 3.48$) than in either the relevant case ($m = .96$) or the irrelevant condition ($m = 1.76$), both $p < .05$. Also, the irrelevant group interpreted the combinations by relation linking significantly more frequently than the relevant group did, $p < .05$.

Table 2

Mean Frequencies of Property Mapping (and Relation Linking) by Groups for Conceptual Combination Task

Group	Relevant	Irrelevant	Anomalous
Typical	7.72 (0.64)	7.20 (1.56)	5.00 (3.48)
Atypical	7.20 (1.28)	6.08 (1.96)	
Oxymora	5.84 (2.00)		

Note. Mean frequencies of relation linking are in parenthesis.

Discussion

People predominantly use property mapping to make sense of adjective-noun concept combinations. To examine this phenomenon in greater detail, and to compare the results to my hypotheses, I will proceed through the analysis condition by condition. To begin, participants in the relevant, typical group used property mapping almost to the complete exclusion of relation linking, as predicted.

In fact, the significant main effects of both relevance and typicality for occurrence of property mapping summarizes the usage of combinatorial strategies in the relevant, typical; relevant, atypical; irrelevant, typical; and irrelevant, atypical groups. These main effects indicate that relevant combinations lead to more property mapping than irrelevant pairs do (regardless of typicality), and that typical adjective-noun combinations also result in more property mapping than atypical ones (regardless of relevance). These findings were predicted. In addition, my hypothesis that property mapping would dominate in all four of these groups was verified by the results.

Consequently, also consistent with my predictions are the main effects of relevance and typicality for relation linking. Atypicality leads to more relation linking than in typical pairs (regardless of relevance), and likewise irrelevance results in this technique more often than in relevant combinations (regardless of typicality). It makes sense that where the use of property mapping is less frequent, use of relation linking should be greater (since these are the only two strategies that occur in meaningful frequency), as these results show.

One condition in which people did not interpret the combinations as I expected is the oxymora group. Property mapping was used significantly more than relation linking here, whereas I predicted the contrary. Relation linking was, however, used significantly more often in this condition than in the typical condition as expected, and was more common than in the atypical group as I predicted it would be (though the difference here did not reach significance). The finding that people still mapped properties from the adjective to the noun even though the adjective contradicted the noun is counterintuitive. It is the basis of property negation. Because the result is not intuitive or predicted by the literature, this property negation was not anticipated and had to be created ad hoc during the scoring phase. And furthermore, because this property negation was not anticipated in advance, my predictions for the group of oxymora were misguided and mistaken.

At last, despite the fact that the anomalous group produced the most relation linking of any group by far, my prediction that it would be more frequent than property mapping was proven erroneous. I seem to have underestimated people's abilities to map properties in adverse situations.

After having discussed these statistical results, justification for the scoring may now be necessary. In particular, the reasoning behind identification and selection of combinatorial techniques should be explained. To begin, it should be evident that relation linking is a legitimate strategy and its existence is uncontested in the literature. The same may be said of property mapping. One question that might arise, then, is why it is that I would choose to study things that are already proven (unless the goal is replication). One answer is that I wanted to dissect property mapping into different, more precise types of mapping. This, though, can again be challenged for a reason. The answer to this, quite simply, is to pursue an interest and more importantly, to expand the boundary of knowledge. When it first became known that relation linking existed, scores of researchers broke this strategy down into its smaller parts such as the CAUSE, HAVE, FOR, MAKE, etc. relations (e.g., Clark, 1983; Coolen et. al, 1991; Downing, 1977; Levi, 1978; Shoben, 1993; Shoben & Gagne, in press). Compared to the frenzy of work done on relation linking, property mapping has been neglected thus far.

With regard to the distinction between (regular) property mapping and conditional property mapping, I think that the distinction is real and that it is an important one. No one could seriously argue that the phrase "it is sunny outside" has the same meaning as "it is sometimes sunny outside." Similarly, "X is my enemy" is different from "X is sometimes my enemy." Conditional property mapping, like that used in defining friendly enemy as "someone who is friendly when I am around, but mean when I am not around," is distinct from regular property mapping cases like "someone who opposes you but is nice" for the same combination friendly enemy. However, even in conditional property mapping, as the term implies, one or more properties is still being mapped. Thus, conditional property mapping is a subset of (regular) property mapping.

Along the same lines, explicit property negation, while distinct in its own right, is assumed to utilize the general process of property mapping and is therefor another subset of property mapping. You may recall that the definition "a nontruth that has no negative stigma associated with it" for the pair acceptable lie is an instance of property negation. The property has negative stigma is explicitly negated from, or crossed out of, the concept lie. In these cases, it is perhaps more obvious to see how property negation is distinct from rather than similar to property mapping. At this point, it will suffice to say that what is assumed to occur with explicit property negation is property mapping, followed by the realization that the property to be transferred contradicts an essential element of the host concept, which then leads to a sort of accommodation, adjustment, or reconciliation of one of the concepts. The adjustment is to one of the two contradictory properties, which in the case of property negation, is to cancel one of them out. More will be said of this later.

Implicit property negation, though, does not proceed in this manner. An example of implicit property negation is the definition "an area protected from the sun that is still warm" for the combination warm shade. The property warm is simply mapped onto shade. However, any cursory review of the attribute list for the individual concept shade will show that the concept has as a salient feature the property cool. This is the reason for referring to it as property negation, and it is implicit because no property is directly negated. Instead, a contradictory one is just mapped onto the other concept. As a result, I have not given this implicit form of property negation a separate category name. It will be included in the regular property mapping category. I suggest that this form of negation differs from the explicit one only in the importance of the attribute that is contradicted. If the feature is not psychologically essential, as cool must not be for shade, then implicit negation will occur. If the attribute is an essential, central, or defining one, then accommodation will be required and explicit negation may result.

Finally, I must justify the "other" category, or the definitions that we could not analyze into one of the specified groups of combinatorial techniques. But I would like first to acknowledge that this category is the result of a failure to develop a perfect scoring design. On the other hand, though, some definitions really were exceedingly difficult to interpret. For example, something as straightforward as blue sky was defined by one participant as "air flying ocean." Whether this is some sort of free association exercise, an indicator of schizophrenia, or an instance of implicit property negation is anyone's guess. But rather than guess, the author chose to include an "other" group. Also, the inclusion of such a cast-out category is common practice in conceptual combination experimentation (See Wisniewski, in press, for an example). A fair 6.8% of all definitions were placed into this group, a percentage in line with that of other similar studies.

General Discussion

Now that we know what mechanisms people utilize to make sense of adjective-noun concept combinations, how these mechanisms operate, and under which conditions any given strategy is likely to be used, let us consider what implications the results of this study might have for theories of concept combination.

First, though, I would like to briefly point out a theoretical inconsistency within theories of concept combination. Up to this point, the focus in the literature of slot filling has been on its tie to relation linking, and in fact some researchers refer to the process of filling a slot with a value from the other concept as relation linking (e.g., Gentner, Markman, & Medin, 1995; Markman & Wisniewski, 1996; Wisniewski & Markman, 1993). In other words, slot filling and relation linking are used as synonyms. Slot filling may be relation linking in cases such as snake robin, which was defined as "a snake that eats robins." Here, the slot EATS in the concept snake is filled with the value robins. But it is equally clear that when red X is defined as "an X that is red," slot filling is property mapping. The concept X contains the slot COLOR, which is filled by red. Red is a property, and it was transferred onto X. Wisniewski (in press) clearly acknowledges that property mapping may involve slot filling, but it is not clear how a mechanism (property mapping) that is contrasted with slot filling (or relation linking) can itself use slot filling.

Recently, it has been postulated that conceptual combination involves a comparison process in which the structures of the two concepts are put into alignment (Wisniewski, 1996, in press; Wisniewski & Gentner, 1991; Wisniewski & Markman, 1993). Markman & Gentner (1993) argued that, by putting the structures of the two concepts into correspondence with one another, differences can be discerned and an interpretation can be formed on the basis of those differences. I too would like to provide support for a comparison process in conceptual combination, and the support that I will provide is qualitatively different from the existing base.

To illustrate why I believe that a comparison process or some sort of structural alignment must occur in combining concepts, take the example of the word pair pagan marriage. This was defined as "a marriage that doesn't take place in a church" by one of my volunteers. In the experiment, this was scored as explicit property negation, because a property of one concept (takes place in a church) was negated by a property of the other concept (see Franks, 1995 for a detailed account of negating privatives in concept combination). Certainly the concept pagan does not contain the property does not take place in church, and a review of the attribute lists for the individual concept pagan showed this intuition to be true. If this is not a property of the concept, then the interpretation could not have taken place by the direct mapping of a property from the adjective concept onto the noun concept. Instead, I suggest that people compare the two concepts, then realize that a marriage takes place in a Christian church and that pagans are not Christians. Here, a difference has been found as the result of the comparison process. Normally, it is a difference found in this way that leads to understanding how the concepts can be combined. For instance, a zebra horse is a horse with stripes. We realize via the comparison process that zebras differ from horses in that zebras have stripes, and this property is then easily mapped onto the horse. But with pagan marriage, finding this difference does not help us see what to map where, as it does with other combinations. Rather, the difference makes it more difficult interpret, because the two properties are incompatible. At this point, something must give-in if the combination is to be understood, so one of the properties is simply negated by the other.

It is hard to see how a property of one concept could be negated by the other concept without some sort of comparison process. It could be possible that one of the concepts in a case like this is simply mapping a negative property instead of negating a property in the other concept, but Nisbett & Ross (1980) found evidence that concepts do not represent the absence of properties. On the contrary, there must be some sort of interactive property attribution in conceptual combination, which parallels that proposed by Glucksberg et. al (1996) for the comprehension of metaphor. It seems that the concepts are not just being compared, but are also interacting in order to reach an adjustment of meaning of one of the concepts. There has to be more than a realization of difference as proposed by the comparison process theory.

There has to be interactive feedback between concepts. Without interactive feedback, pagan marriage would be interpreted in the following manner: comparison of concepts, resulting in realization of difference, followed by mapping a property of one concept onto the other (despite the fact that this property is incompatible with a property in the other concept). With interactive feedback, though, the realization of difference is followed by evaluation of compatibility. If the difference represents compatible values, a property will be mapped. However, if the difference contains values that are incompatible, the concepts receive feedback of this problem, and interaction between concepts is required to adjust one concept so that the property of the other may be accommodated.

Again, without interactive feedback, we would be ignorant to incompatibilities. This is clearly not the case with "a marriage that doesn't take place in a church."

Conclusion

In closing, though much more work needs to be done to investigate these notions of interactive feedback and the comparison process, I have demonstrated how studying combinatorial strategies can lead to insights on theories of concepts. I also have determined which of the strategies for sense generation is most likely to be used under specified conditions of relevance and typicality.

We are all sometimes faced with the difficult task of processing unfamiliar concepts, such as "hostile plant." We interact with many of these combinations in everyday life: notebook, bedroom, baby chair, dust pan, etc. Obviously, they are not all mysterious, ambiguous concepts. It is precisely this sort of practical application of combinatorial strategies that makes investigations of combined concepts important. After all, conceptual combination is the basis for comprehending and extending language. Each word carries its own meaning, or its own set of standard attributes. When two concepts are combined, the words modify and perhaps alter the meaning of those concepts. Together, they create an entirely new concept. By forging ahead with research in concept combination, we head toward a better understanding of language and knowledge representation.

References

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

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.

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

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

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

Francis, W.N., & Kucera, H. (1982). Frequency analysis of English usage: Lexicon and grammar. Boston: Houghton Mifflin.

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

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

Gentner, D., & France (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S.L. Small, G.W. Cotrell, & M.K. Tanenhaus (Eds.), Lexical ambiguity resolution in the comprehension of human language. Los Altos, CA: Morgan Kaufman.

Gentner, D., Markman, A.B., & Medin, D.L. (1995). The structure of similarity. Manuscript submitted for publication.

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

Gibbs, R.W., & Kearney, L.R. (1994). When parting is such sweet sorrow: The comprehension and appreciation of oxymora. Journal of Psycholinguistic Research, 23 (1), 75- 89.

Gleitman, L.R., & Gleitman, H. (1970). Phrase and paraphrase: Some innovative uses of language. New York: W.W. Norton.

Glucksberg, S., McGlone, M.S., & Manfredi, D. (in press). Property attribution in metaphor comprehension. Journal of Memory and Language.

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

Hampson, S.E. (1990). Reconciling inconsistent information: Impressions of personality from combinations of traits. European Journal of Personality, 4, 157-172.

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

Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. Bulletin of the Psychonomic Society, 28, 242-247.

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

Kunda, Z., Miller, D., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. Cognitive Science, 14, 551-577.

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

Lukatela, G., Carello, C., Kostic, A., & Turvey, M.T. (1988). Low-constraint facilitation in lexical decision with single-word contexts. American Journal of Psychology, 101 (1), 15-29.

Markman, A.B., & Gentner, D. (1993). Splitting the difference: A structural alignment view of similarity. Journal of Memory and Language, 32 (4), 517-535.

Markman, A.B., & Wisniewski, E.J. (1996). Similar and different: The differentiation of basic level categories. Manuscript submitted for publication.

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

Minsky, M. (1975). A framework for representing knowledge. In P.H. Winston (Ed.), The psychology of computer vision. 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, 29, 259-288.

Nisbett, R.E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, N.J.: Prentice-Hall.

Rumelhart, D.E. (1980). Schemata: The building blocks of cognition. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), Theoretical issues in reading comprehension. Hillsdale, N.J.: Erlbaum.

Shen, Y. (1987). On the structure and understanding of poetic oxymoron. Poetics, 18, 511-536.

Shoben, E.J. (1993). Comprehending nonpredicating conceptual combinations. In G. Nakamura, R. Taraban, & D. Medin (Eds.), The psychology of learning and motivation, 20, Academic Press.

Shoben, E.J. & Gagne, C.L. (in press). Thematic relations and the creation of combined concepts. In Ward, T.B., Smith, S.M., & Vaid, J. (Eds.), Conceptual structures and processes: Emergence, discovery, and change. Washington, D.C.: American Psychological Association.

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

Wisniewski, E.J. (1996). When concepts combine. Manuscript submitted for publication.

Wisniewski, E.J. (in press). Construal and similarity in conceptual combination. Journal of Memory and Language.

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

Wisniewski, E.J., & Markman, A.B. (1993). The role of structural alignment in conceptual combination, Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, 1083-1086, Boulder, Co.

Appendix: Set of Stimuli for Conceptual Combination Task

Relevant, typical	Relevant, atypical	Oxymora
hostile enemy	indifferent enemy	friendly enemy
harmful illness	harmless illness	healthy illness
cheerful joy	unemotional joy	painful joy
unacceptable lie	acceptable lie	truthful lie
cool shade	warm shade	sunny shade
tragic death	celebrated death	living death
blue sky	purple sky	grounded sky
satisfying marriage	depressing marriage	divorced marriage
long delay	scheduled delay	speedy delay
clear simplicity	unclear simplicity	complex simplicity

Irrelevant, typical	Irrelevant, atypical	Anomalous
dressed enemy	undressed enemy	cloudy enemy
feared illness	desired illness	hinged illness
spontaneous joy	deliberate joy	grassy joy
verbal lie	nonverbal lie	crusty lie
visible shade	invisible shade	liquid shade
urban death	rural death	round death
permanent sky	temporary sky	heroic sky
Christian marriage	pagan marriage	thick marriage
boring delay	entertaining delay	soupy delay
logical simplicity	illogical simplicity	zipped simplicity

Author Identification Notes

The author would like to thank, of course, Tom Ward for giving me the freedom to run in all directions at once, yet guiding me just enough to ensure that I actually arrive somewhere. I also benefited greatly from conversations with Art Markman, Ed Wisniewski, Dedre Gentner, and Rob Goldstone, among others. Special thanks to Missi Wilkenfeld and Debbie Tindell for helping me bring this all together. I couldn't have done it without their help. Nor could I have analyzed my results if not for the generosity of Justin Patterson, who programmed an excellent program, as programmers do. Also, I am very appreciative of the Office of Honors Programs and Academic Scholarships at Texas A&M University for providing me not only with a fine opportunity, but the support needed to get through it all too. Finally, I would like to thank George Killian for his sweet Irish Red.