

Structural Alignment in Relational Interpretations of Conceptual Combinations

Cynthia M. Sifonis (cms1210@acs.tamu.edu)

Department of Psychology; Texas A&M University
College Station, TX 77843-4235 USA

Thomas B. Ward (tbw@psyc.tamu.edu)

Department of Psychology; Texas A&M University
College Station, TX 77843-4235 USA

Abstract

Current theories concerning the comprehension of noun-noun combinations propose that relational interpretations are the result of the modifier filling some role within the head noun (Murphy, 1988; Wisniewski, 1997), whereas property interpretations involve the structural alignment of the head noun and modifier concepts (Wisniewski, 1997). In this paper we argue that structural alignment underlies the formation of both relational and property interpretations of noun-noun combinations. Property interpretations result from the alignment of the modifier with the head noun, whereas relational interpretations result from the alignment of the modifier with a *filler* in the head noun. Modifiers for noun-noun combinations were chosen based on their similarity to relational fillers in the head noun concept. Results indicated that frequency of instantiation of relational interpretations was positively correlated with the similarity of the modifier to fillers in the head noun. Similarity of modifier to fillers in the head noun predicted frequency of instantiation of relational interpretations to a greater degree than the rated ability of modifiers to fill roles in the head noun concept.

Introduction

Research on the comprehension of conceptual combinations has proven valuable to several fields of cognitive science. It has provided information concerning the manner in which combinations are utilized to specify referents of discourse contexts and to extend the vocabulary of a language. It has also provided information concerning the structure of concepts and categories, and the constraints on the utilization concepts which need to be implemented in comprehensive models of conceptual functioning (Markman & Wisniewski, 1997; Medin & Shoben, 1988).

Conceptual combination models have been developed to explain various types of combinations. In this paper we focus on the processes involved in the interpretation of noun-noun combinations and propose a mechanism by which interpretations are formed for such combinations.

Current models designed to explain the interpretation of noun-noun combinations take one of two forms - a thematic relations approach (Gagne & Shoben, 1997; Shoben & Gagne, 1997) or a schema approach (Murphy, 1988; Wisniewski, 1997). Both approaches assume that interpretations of noun-noun combinations are formed through the identification of a relationship that holds between the modifier and head noun in the combination. For example, the combination CATERPILLAR APPLE might be understood as referring to an apple that has a caterpillar living in

it. Interpretations of this sort are known as relation-linking (Wisniewski, 1997) or relational interpretations.

Thematic models propose that relational interpretations are formed by determining which one of a fairly small subset of general relations holds between the two concepts. This approach assumes the relation is chosen based on the combinatorial history of the constituents of the combination (Gagne & Shoben, 1997; Shoben & Gagne, 1997).

Schema models represent the concepts in a combination as schemas or frames. Knowledge associated with each concept is stored in terms of slots and fillers where the slots refer to dimensions of the concept and the fillers associated with those slots consist of the various instantiations of that dimension in the concept. Slots also contain default fillers which are the most frequently instantiated or typical value for a particular slot. Schema approaches assume that interpretations of combinations result from the modifier noun filling a relational slot in the head noun's conceptual space. To explain how a slot is selected to be filled by the modifier, these models propose that associated with each slot is a set of preconditions that must be met by the modifier in order for it to fill that slot (Wisniewski, 1997). For example, in order for the combination BLUEBERRY SPOON to be interpreted as referring to "a spoon for eating blueberries", the modifier of the combination (Blueberry) must meet the constraint of being edible in order to fill the Used For Eating slot of the head noun concept (Spoon).

Although both thematic and schema models predict that combinations will result in relational interpretations, neither predict the existence of another common type of interpretation, namely property interpretations. Property interpretations involve the assertion of a property of the modifier to be true of the head noun (e.g., CATERPILLAR APPLE interpreted as referring to "a fuzzy apple") and constitute about 33% of the interpretations generated for a combination (Costello & Keane, 1997). To extend the theoretical scope of conceptual combination models to account for this phenomenon, Wisniewski (1997) has proposed that property interpretations are formed through the structural alignment of the head noun and modifier, a different mechanism from the one that produces relational interpretations (i.e., role filling). While this account represents an important advance, a more parsimonious explanation would be that one process (structural alignment) underlies the formation of both types of interpretations. Thus, a goal of this paper is to demonstrate that this same mechanism of structural alignment can be used to account for relational interpretations of conceptual

This is not to say that role filling does not occur, ample evidence exists suggesting it does. However, a role-filling explanation of the generation of relational interpretations leaves several questions unanswered. How is it that a slot is selected in the head noun concept for the modifier to fill? Why is it that the combinatorial history of a word influences the relational interpretation chosen?

We propose that structural alignment of the constituents of a combination underlies both processes. Structural alignment of the head noun and modifier is necessary for determining the role the modifier can fill in the head noun. The relational structure of a concept determines how it is going to interact with most other concepts thereby influencing its combinatorial history.

Whether a property interpretation or a relational interpretation results from the alignment process depends on the manner in which the constituents of the combination are aligned. Property interpretations result from the global alignment of the modifier and head noun concepts. In comparison, relational interpretations result from the alignment of the modifier with a filler of the head noun.

Because the concepts involved in noun-noun combinations are complex concepts, those concepts consist of a variety of slots containing fillers which are concepts in their own right (Murphy, 1988). When the modifier is aligned with one of these filler concepts the result is a relational interpretation of the combination.

For example, for the combination CATERPILLAR APPLE, if the concepts of Apple and Caterpillar are aligned such that the slots for texture in both concepts are brought into alignment then "fuzzy" (the texture value for Caterpillar) replaces "smooth" (the texture value for Apple) as the texture for Apple resulting in the interpretation of CATERPILLAR APPLE as "a fuzzy apple". However, if Caterpillar is aligned with the default filler of Worm in the Apple concept (Worm being the default filler of the slot Has Living In It) this would result in the relational interpretation of the combination "an apple with caterpillars living in it".

To reiterate, alignments found at the most global level of the constituents often result in property mapping interpretations of the combination (Wisniewski, 1997). When the concepts are searched recursively and structural alignments are sought within the slots of the constituents, this results in relational interpretations. This is possible because the concepts being combined (especially noun-noun combinations) are frequently complex concepts whose slots and fillers are often concepts in their own right (Murphy, 1988).

Empirical evidence has been provided by Wisniewski (1997) suggesting that property mapping interpretations of combinations are more likely to be generated when the constituents of the combination have very similar representational structures. This similarity facilitates the alignment process (Wisniewski, 1997). If structural alignment operates in the formation of relational interpretations, it should be possible to manipulate the similarity of the modifier with respect to default fillers of the head noun to affect the generation of relational interpretations of the combination. The more similar a modifier is to a particular filler in the head noun, the more completely it will align with the filler

concept and the more likely it will be that a relational interpretation will be generated for the combination in which the modifier fills the role of the filler with which it was aligned.

In this experiment we tested the hypothesis that the similarity between a modifier and a default filler of a head noun is positively correlated with the frequency with which the modifier fills the same role as that default filler in the interpretation of the combination. We also examined whether relational interpretations of the combination would be better predicted by the rated ability of the modifier to fill a role in the head noun rather than by its similarity to the filler of that role.

If relational interpretations of conceptual combinations result from the alignment of the modifier with fillers of the head noun then similarity of the modifier to a default filler in the head noun should be positively correlated with the frequency of instantiation of relational interpretations for the combination. The more similar the modifier is to a default filler the more frequently it is instantiated as participating in the same relation as the default filler in the interpretation of the combination.

Because the ability to determine what role the modifier can fill in the head noun depends on the alignment of the modifier and head noun concepts, it is predicted that similarity and role filling will be highly correlated with each other. However, if structural alignment is the underlying mechanism of role-finding, then similarity ratings will be a more direct measure of this process. Consequently, we predict that the similarity of the modifier to a filler in the head noun will uniquely account for more of the variability in the frequency of instantiation of relational interpretations than will the capacity of the modifier to fill a particular role in the head nouns concept.

The experiment required a series of steps to determine the relationship between similarity, role-finding, and the generation of relational interpretations. The first step involved generating the stimuli set to be used in the experiment. This step consisted of a) choosing a set of nouns to be used as the head nouns in the combinations, b) obtaining feature listings of those nouns to identify their relational features/slots, c) determining the default fillers of those slots, and d) selecting modifiers similar to the default fillers of the head noun that were not already listed as fillers of those nouns.

The next step in the process was to use this stimulus set to determine the relationship between the frequency of instantiation of any given relational interpretation, and the similarity and role-filling capacity ratings of the modifiers and default fillers of the combination. This step involved a) identifying the relational interpretations generated for the noun-noun combinations and determining their frequency of instantiation, b) determining any new default fillers from those interpretations that involved relations that had not been observed during the generation of the stimulus set, c) measuring the similarity of the modifiers to the fillers they replaced and d) obtaining ratings of the capacity of modifiers to fill the roles ordinarily served by the default fillers they replaced.

Step 1: Generating the Stimulus Set

Method

Participants

One hundred twenty-two introductory psychology students participated in generating the stimulus set in exchange for class experimental credit. All were native speakers of English. Fifty-nine students generated the feature listings for the head nouns. Forty students generated the fillers for relational slots in the head nouns and twenty-three students participated in the similarity rating task.

Materials and Procedure

Feature Listing Task

To assure the generalizability of the results we chose nouns that spanned a broad range of categories.

We used a system similar to the one employed by Wisniewski & Gentner (1991). Nouns were chosen that differed along two dimensions: artifact versus natural kind and count noun versus mass noun. Place Nouns (artificial and natural) were also included in order to be able to generalize the findings to more complex structures such as schemas.

Six nouns from each group were selected. To further broaden the range of concepts explored, nouns were chosen from each group that occupied different taxonomic categories (e.g., animal, flower, fruit, tree, insect, vegetable for natural count nouns). This resulted in the following set of words used as the head nouns of the study:

Count nouns/natural kinds: frog, poppy, apple, oak, wasp, lettuce

Count nouns/ artificial kinds: chair, gun, helicopter, rake, pencil, dress

Mass nouns/ natural kinds: clay, copper, sand, milk, mud, fog

Mass nouns/ artificial kinds: pudding, glass, steel, plastic, glue, trash

Place nouns/ natural kinds: Park, Mountain, Beach, Desert, Lake, Forest

Place nouns/ artificial kinds: Home, School, Store, Restaurant, Office, Circus

To choose the words used as the modifiers in the combinations, feature listings for each of the head nouns were obtained. The 59 participants involved in the feature listing task were given a booklet containing 12 of the head nouns in the corpus (2 nouns from each of six groups) and were asked to list the features of those nouns. Word order was counterbalanced across participants to control for order effects.

Filler Generation Task

Because the experiment focuses on the generation of relational interpretations, it was necessary to identify the relational features of each noun listed by subjects in the feature listing task. Relational features are features denoting

roles the noun fills, activities within which it participates, or connections with other concepts. These relational features (e.g., "used for writing on paper" for Pencil, "tossed in garbage can" for Trash) were then used to determine the objects of those relations. Relational objects are the objects of relational features. For example, "paper" would be the relational object of the relational feature "used for writing on paper". The relational features might be considered slots of the concept, and the relational objects of those features are the fillers of those slots.

Because many of the students participating in the feature listing task did not explicitly specify the relational object of the relational feature they listed for a given word (e.g., listing "flies" as a feature of bird as opposed to "flies with wings" or "flies through the air"), the relational features listed for each noun were compiled and a group of forty students were asked to generate the relational objects for each of these relational features. For example, they were asked to fill in the blanks in statements such as "Birds fly with _____." This resulted in a list of relational features and the objects of those features for each head noun in the corpus.

For each head noun, the three most frequently listed relational features (high dominant slots) and the most frequently listed filler (the default filler) for each feature were noted. The three least frequently listed relational features (low dominant slots) and their default fillers were also noted. The addition of low dominant features was included to broaden the range of relational interpretations to which the results could generalize.

The hypothesis of the study states that similarity of the modifier to fillers of the head noun will facilitate the alignment of the modifier with the filler, resulting in a relational interpretation. To test this hypothesis it was necessary to select modifiers for the combinations that were similar to the default fillers of the head nouns but that had not been listed by any participant as a potential filler for a relational slot in the head noun. For example, when "used for writing on paper" was listed as a relational feature of Pencil, "parchment" was chosen as a potential modifier because of its conceptual similarity to the relational object "paper" and because it did not already participate as a filler in that relation. This initial choice of potential modifiers was based on the experimenters' intuitions regarding similarity to the default filler, but those intuitions were subsequently verified via a similarity judgment task. This process produced a list of 216 potential modifiers (Three modifiers chosen from high dominant relations and three modifiers chosen from low dominant relations for each of the 36 head nouns).

Similarity Rating Task

Once potential modifiers were selected, it was necessary to choose those modifiers with the highest similarity to the default fillers from which they were generated.

Twenty-three students judged the similarity of the potential modifiers to the relational objects from which they were derived (e.g., judge the similarity of "parchment" to "paper").

For each head noun, the modifier-relational object pair with the highest similarity rating out of the three high dominant pairs and the modifier-relational object pair with the highest similarity rating out of the three low dominant pairs were chosen to serve as the modifier to be paired with the head noun from which they were generated. This resulted in each of the 36 head nouns being paired once with the modifier generated from the high dominant relational feature and once with the modifier generated from the low dominant relational feature (72 noun-noun pairs in all).

Step 2: Hypothesis Testing

Hypothesis testing involved using the noun-noun combinations generated in Step 1 to determine the relationship between the frequency of instantiation of any given relational interpretation, and the similarity and role-filling capacity of the modifiers to the default fillers of the relation. In this step we a) had subjects generate interpretations to the noun-noun combinations, b) identified the relational interpretations generated for the noun-noun combinations and determined their frequency of instantiation, c) obtained a listing of default fillers for those interpretations that involved relations not produced during the generation of the stimuli set in Step 1, c) measured the similarity of the modifiers to the fillers they replaced and d) obtained ratings of the capacity of modifiers to fill the roles ordinarily served by the default fillers they replaced.

Method

Participants

One hundred and nineteen students from the same subject pool as participants in Step 1 participated in exchange for experimental credit. Fifty-nine students generated interpretations for the conceptual combinations, 24 participated in generating relational objects for relational features that emerged from the combination, but that had not been observed in the previous feature listing task, 19 participated in the similarity rating portion of the experiment, and 17 participated in the capacity rating portion of the experiment. Participants were all native speakers of English.

Materials and Procedure

Interpretation Generation

In order to determine the frequency of instantiation of relational interpretations it was necessary to have subjects generate interpretations to the combinations generated in Step 1.

Each participant received a booklet containing 36 of the 72 conceptual combinations generated in Step 1 for which they had to generate interpretations. Though word order did not vary within the noun-noun pairs, presentation of pairs was counterbalanced to control for possible order effects.

Between 1 and 7 different relational interpretations were generated for each combination. The frequency of instantiation of each interpretation was then determined.

Filler Generation

For each relational interpretation generated for a combination, our intent was to assess the similarity of the modifier in the combination to the default filler (relational object) of the relational feature in question. In most cases, the default filler had been identified from responses to the filler listing task in Step 1. Because some subjects generated relational interpretations for combinations which involved slots for which the default fillers were unknown, it was necessary to compile a list of those newly emergent relational features and have another group of subjects supply the relational objects of those features. For example, subjects given the combination CATERPILLAR APPLE sometimes provided the relational interpretation "an apple made out of caterpillars" Because Made Of was not listed as a relational feature for Apple by subjects engaged in the feature listing task in Step 1, its relational object had not been determined. Consequently, a different group of subjects were given the open ended question "Apples are made of ____?" and were asked to fill in the blank.

Twenty-four subjects participated in the filler generation task. The most frequent response for each relational feature was determined. These were then counted as the default fillers for that relational feature.

Similarity Rating

Once the full set of default fillers were determined for each relation used in the relational interpretations generated for the set of 72 combinations, we determined the similarity of the modifier noun to the filler it was replacing in the interpretation.

The default filler for each relational interpretation was paired with the modifier noun of the combination from which that interpretation was generated. Nineteen subjects were asked to rate the similarity of the complete set of modifier-default filler pairs associated with the combinations. The mean similarity rating was then determined for each relational interpretation.

Capacity Rating

The same modifier-filler pairs were given to another group of 17 subjects who were asked to judge how well the modifier could serve in the same capacity of the filler it was replacing for the head noun (e.g., "How well could parchment fill the same role as paper for a pencil?") The mean capacity rating for each interpretation was then determined.

For the correlational analyses, the frequency of occurrence of each relational interpretation was paired with the mean similarity rating and the capacity rating of the modifier-default filler pair associated with that interpretation.

Because modifiers in this study were chosen based on their high similarity to the default fillers of relational features, interpretations in which the modifiers serve in the same capacity as those default fillers ought to predominate. Should this not occur, it would provide a serious challenge

to the hypothesis under investigation. Thus, to provide a conservative estimate of the correlations, the 72 target relational matches identified in Step 1 were included regardless of whether participants actually gave relational interpretations that included them. If no participants provide relational definitions that made use of a given target relation, that relation received a frequency score of 0 which was matched with the mean similarity and capacity ratings of the modifier and the default filler of the target relation. This was necessary for five of the 72 combinations.

Results

Analysis of the data involved determining the correlation between the frequency of instantiation of an interpretation and the mean similarity and capacity ratings of the modifier to the filler with which it was aligned in the head noun. This involved an item-by-item analysis in which the unit of analysis was the relational interpretations given for the combinations.

A Pearson correlation examined the relationship between frequency of instantiation ($M = 6.48$) and similarity ratings ($M = 5.44$). There was a significant positive correlation $r(225) = .442$, $p < .001$ between frequency of instantiation and similarity suggesting the greater the degree of similarity between the modifier and a default filler, the more frequently the relational interpretation was instantiated in which the default filler was replaced by the modifier.

A Pearson correlation examined the relationship between frequency of instantiation ($M = 6.48$) and capacity ratings ($M = 5.17$). There was a significant positive relationship, $r(225) = .425$, $p < .001$, between frequency of instantiation and capacity suggesting that the greater the degree to which a modifier was judged to be able to fill the same role as the default filler, the more frequently that relational interpretation was instantiated.

Similarity ratings and capacity ratings were also highly and significantly correlated with each other $r(225) = .878$, $p < .001$ suggesting that similar processes or constructs underlie subject's performance on either task.

Because of the high correlation between similarity and capacity ratings, a multiple regression analysis was performed in order to determine which variable is contributing more to the frequency of instantiation of a relational interpretation. Analysis revealed that similarity entered the equation first accounting for 19% of the variance in frequency of instantiation, $F(1, 226) = 54.97$, $p < .001$, and that rated capacity did not account for a significant percentage of the remaining variance.

General Discussion

The hypothesis that there would be a significant positive correlation between frequency of instantiation of a relational interpretation and the similarity of a modifier to the default filler for that relation was supported. There was also a significant positive correlation between the frequency of instantiation of a relational interpretation and the ability

for the modifier to serve in the same capacity as the default filler for that relation in the head noun. However, when the variance due to similarity was accounted for, the correlation between a modifier's ability to fill a role in a slot and the use of that slot in interpreting the combination was no longer significant.

Similarity has frequently been used as a measure of alignability. The more similar two concepts are the better they align structurally. Conversely, the better two concepts align structurally the more similar they are judged as being. The results of this study suggest that similarity and therefore alignability is a factor in the generation of relational interpretations of conceptual combinations.

To extend this argument, the results parallel earlier findings that have been used to support the operation of structural alignment in property mapping interpretations. Specifically, Wisniewski (1997) used the fact that high levels of modifier-head noun similarity were associated with greater use of property mapping to support the idea that property mapping definitions result from a structural alignment process. By the same reasoning, the fact that modifier-relational object similarity predicts the extent to which particular relational interpretations are used supports the use of structural alignment in relational interpretations.

Importantly, we are not suggesting that role filling does not occur. On the contrary, it is quite common. What we are suggesting is that a structural alignment, similarity assessment process may well underlie determinations of exactly which role the modifier will fill in the head noun structure. The fact that similarity accounts for more unique variance than the rated role-filling capacity of the modifier in the generation of relational interpretations suggests that similarity judgment is the more direct measure of the processes at work.

The results suggest that structural alignment operates not only in the generation of property interpretations of combinations but also in relational interpretations. The former involves the alignment of the modifier and the head noun in a global sense, whereas the latter involves the alignment of the modifier to fillers of the head noun.

The fact that one process, structural alignment, can account for both property mapping and relational interpretations, provides a parsimonious account of the processing involved in comprehending conceptual combinations.

References

- Costello, F. & Keane, M.T. (1997). Polysemy in conceptual combination: Testing the constraint theory of combination. Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Erlbaum.
- Gagne, C.L., & Shoben, E.J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23, 71-87.
- Markman, A.B. & Wisniewski, E.J. (1997). Similar and differentiation of basic-level categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 54-70.

- Murphy, G.L. (1988). Comprehending complex concepts. Cognitive Science, 12, 529-562.
- Medin, D.L. & Shoben, E.J. (1988) Context and structure in conceptual combination. Cognitive Psychology, 20, 158-190.
- Shoben, E.J. & Gagne, C.L. (1997). Thematic relations and the creation of combined concepts. In T.B. Ward, S.M. Smith, & J. Vaid (Eds.), Creative thought: An investigation of conceptual structures and processes. Washington, DC: American Psychological Association.
- Wisniewski, E.J. (1997). Conceptual combination: Possibilities and esthetics. In T.B. Ward, S.M. Smith, & J. Vaid (Eds.), Creative thought: An investigation of conceptual structures and processes. Washington, DC: American Psychological Association.
- Wisniewski, E.J. & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In Simpson (Eds.), Understanding Word and Science. North-Holland: Elsevier Science Publishers