The role of salience in conceptual combination

JEANNINE S. BOCK and CHARLES CLIFTON, JR.
University of Massachusetts, Amherst, Massachusetts

The present study investigated three factors that affect the interpretation of novel noun–noun (N–N) combinations: simple feature salience, ontological category, and assessed similarity. Participants read and defined a series of novel N–N combinations in which the feature salience of N1 and N2 was manipulated. Participants also rated the combinations for similarity. The combinations were constrained to be within ontological category. All interpretations were scored in terms of the strategies (property mapping vs. relation linking) used to produce the given interpretations. Highly salient features drove property-mapping interpretations based on those features. Natural kinds produced more property-mapping interpretations than did artifacts. There was no correlation between the proportion of property-mapping interpretations and the assessed similarity of the N–N combinations. These results are discussed as an extension of Estes and Glucksberg's (2000) interactive theory of conceptual combination and argue for the importance of feature salience as a factor in conceptual combination.

How people interpret noun–noun (N–N) combinations (e.g., *tiger mouse* or *newspaper spear*) is an important question in language processing. Conceptual combinations allow a speaker to specify a concept or a referent without using grammatically more complex constructions. N–N compounds can be used to expand a vocabulary to include novel terms (Downing, 1977; Gerrig & Murphy, 1992; Wisniewski, 1996; Wisniewski & Gentner, 1991) or to further subcategorize or elucidate representation of existing items (Hampton, 1987; Markman & Wisniewski, 1997; Murphy, 1988, 1990; Smith, Osherson, Rips, & Keane, 1988; Wisniewski, 1996). N–N compounds are common in English and are created, computed, and resolved successfully given sufficient and appropriate contexts. It is, therefore, of interest to define the factors and processes that constrain the construction and interpretation of such combinations.

One aspect of the current literature concerns the strategies that are used in the construction and interpretation of conceptual combinations. Two strategies that have received much attention are property mapping and relation linking (Wisniewski, 1996, 1997a, 1997b, 1998; Wisniewski & Gentner, 1991; Wisniewski & Love, 1998). In property mapping, a particular property or attribute of the modifier noun (N1) maps to the head noun (N2) such that a *robin termite* might be a termite with a red underbelly. In relation linking, a relation is posited between N1 and N2 such that a *robin termite* becomes a termite that eats robins or their nests (Wisniewski, 1996).

In this paper, we assume that a variety of specific factors may favor the use of property mapping over relation linking, whereas relation linking may be a fallback strategy used when none of these specific factors are realized. Although this proposal can be viewed as the opposite of the property-as-last-resort hypothesis advanced by Downing (1977), it permits relation linking to be the preferred strategy for a wide range of conceptual compounds (Gagné, 2000). Property mapping occurs only under certain specified (if realistic and typical; cf. Wisniewski & Love, 1998) conditions. We address three factors that may potentially favor property mapping.

The first factor is the ontological category of the terms being combined. One can easily observe property mapping in such lexicalized natural kind pairs as *zebra finch*, *spider plant*, and, arguably, *swordfish*. In these cases, a stereotypical property of the initial (modifying) noun is used to describe the referent of the compound. Conversely, lexicalized artifact compounds that exhibit relation linking are easily found (e.g., *can opener, card file*, and *park bench*). Wisniewski and Love (1998), in fact, examined a sample of naturally occurring combinations and found that, by and large, combinations of animal names and plant names had property-mapping meanings more frequently than combinations of artifact names did.

Artifacts are, by their very nature, relational entities. They are defined by their function, whereas natural kinds can be seen as defined by their properties. Although artifacts may have highly salient attributes (a knife is sharp; sandpaper is rough), these attributes often are causally related to the function of the item (knives being sharp allows them to cut; sandpaper being rough allows it to smooth; cf. Ahn, 1998). Several researchers have reported data about how both children and adults define terms and how they believe various kinds of changes succeed in transforming one kind of object into another. Artifact terms were more likely to be defined in terms of functions and

relational features, and natural kind terms were more likely to be defined in terms of intrinsic properties and internal structure (Barr & Caplan, 1987; Barton & Komatsu, 1989; Keil, 1989; Rips, 1989). It is reasonable to suppose that characteristics important to the meaning of terms would appear in the meanings of combinations of terms. One goal of our research was to explore the extent to which randomly produced natural kind and artifact pairs differ in preference for property-mapping versus relation-linking interpretations.

A second factor that has been linked to property-mapping interpretations is the similarity of the two nouns in the compound. Experimental evidence (e.g., Wisniewski, 1996, 1997b) has shown that more similar pairs result in more property-mapping interpretations. The structural alignment view (Markman & Gentner, 1993; Markman & Wisniewski, 1997; Wisniewski, 1996, 1997b) argues that the modifier and head concepts are aligned with one another in such a manner as to highlight one or more attributes in which the component nouns contrast with one another. According to this view, similar word pairs generally have easily aligned structures, share many commonalities and alignable differences, and therefore encourage more property-mapping interpretations because the alignable differences allow the appropriate attributes to be selected. Dissimilar pairs have less alignable structures and are less amenable to property mapping.

In our experiments, we investigated the role that similarity might play in predicting the generation of property-mapping interpretations. While Wisniewski (1996, Experiment 2) demonstrated that high-similarity pairs encouraged more property mapping than did low-similarity pairs (72% vs. 48%), he used a rather extreme manipulation of similarity. His high-similarity pairs were generally two basic-level terms within the same immediate superordinate category (e.g., *coat–shirt, radish–onion, whiskey–beer*); his low-similarity terms generally came from two distinct categories within a major ontological category (artifact, natural kind, and, arguably, substance for the category of beverages; e.g., *magazine–shirt, mouse–onion, pistol–beer*). However, the theoretical claims underlying the observed similarity effect seem to suggest that similarity would affect interpretation preferences through the whole of the similarity scale, not just at its extremes. We examined this prediction by eliminating the ends of the scale, ensuring some minimal similarity by never pairing natural kind terms with artifacts and avoiding extremely similar basic-level "sister" terms by randomly pairing terms within ontological category. Our procedures also avoided confounding similarity with within- versus between-ontological category pairs.

A third potential factor is the salience of particular features of the component nouns. It seems natural that if a modifier noun (e.g., *tiger*) has a single highly salient feature (e.g., *stripes*), then that feature should come into play in the interpretation of a compound such as *tiger mouse*. Such a feature might provide the first candidate

for the interpretation of an N–N compound. It seems more parsimonious to take such salience into account than to require an alignment and comparison procedure before selecting appropriate attributes. Only in the absence of such salient features may exhaustive alignment and comparison procedures be required.

Estes and Glucksberg (2000) provide evidence for the role of feature salience but add that N2 must have a highly relevant dimension that is appropriate to the salient feature on N1. Their model emphasizes an interaction between N1 feature salience and N2 dimension relevance. They manipulated feature salience of N1 and dimension relevance of N2 while holding the similarity of the compounds constant and low. Combinations with high-salience features on N1 and highly relevant dimensions on N2 that were appropriate to those features (their H–H condition) resulted in more property-mapping interpretations than did combinations in which either the feature salience or the dimension relevance was low.

We generally accept the premises of Estes and Glucksberg's (2000) interactive model. We remain curious, however, about the simple effect of property salience without the benefit of specifically selected appropriate dimensions. We propose that feature salience alone may have a more important role than was apparent in Estes and Glucksberg's data. While we acknowledge that a highly relevant dimension of a head noun can promote use of a salient feature of a modifier noun, we question whether a head noun must already have a salient dimension to which the salient feature of the modifier is highly relevant in order for this modifier feature to be used in interpreting the compound.

To examine the effect of feature salience, we manipulated the simple feature salience of both N1 and N2. We assessed the salience of particular terms and then randomly paired them in order to investigate the role of feature salience without the constraint of having selected relevant dimensions of N2. If simple feature salience is important, we expect to find property-mapping interpretations driven by salient features of N1 even when N2 lacks a highly relevant dimension. The experiment reported by Estes and Glucksberg (2000) did not include all the conditions needed to test this expectation. They investigated only compounds whose modifier had a highly salient feature that was relevant to a salient dimension of the head (H–H), compounds with a modifier whose highly salient feature was not relevant to a commonly mentioned dimension of the head (H–L), and compounds whose head had a commonly mentioned dimension that was relevant to a relatively nonsalient feature of the modifier (L–H). They found a much higher frequency of property-mapping interpretations for H–H compounds (79%) than for H–L and L–H compounds (23% and 16%, respectively). However, they did not investigate L–L compounds. We propose that a highly salient feature of the modifier will appear frequently in the interpretation of H–L compounds, relative to its appearance in L–L compounds.

SALIENCE NORMS

In order to assess the salience of particular features of nouns, a norming study was run on a list of words taken from Battig and Montague (1969). The norming study contained 320 words, divided equally into natural kinds and artifacts. Participants were asked to read half the words and to list any features that they found to be highly salient. Participants were allowed to list more than one feature if they felt it to be appropriate (e.g., a lemon is both yellow and sour). Immediately after listing the features for a word, participants were asked to rate the salience of each feature on a 7-point Likert scale, with 1 being *low* and 7 being *high*. Each participant responded to 160 words. After an initial pilot study, participants were cautioned against listing word associations rather than salient features (e.g., listing *Mickey* for the word *mouse*, or listing *tamer* for the word *lion*).

Data were collected from 60 undergraduates at the University of Massachusetts, who received academic course credit for their participation. All responses were tabulated in terms of features listed and salience ratings given. We identified items for which our participants generally agreed on a single feature that characterized that item and further agreed that the feature was generally a salient quality of the item. We termed these items *high-salient-feature items* (using the phrase *high-salient items* when the context makes the usage unambiguous). Operationally, we included any term for which a particular feature was listed by at least 40% of our participants (at least 12 of the 30 possible responses), who then gave it a salience rating of 5, 6, or 7. *Low-salience items* had no single feature given by more than 8 of our participants with mean salience ratings of 5 or higher. We eliminated items for which participants primarily defined the item rather than listing any actual features of the item (e.g., for the noun *robin*, such a response would be *bird*) and items for which it seemed that the majority of the participants did not know the definition of the word. Of the 320 words normed, we selected 20 words from each of four categories: high- and low-salience natural kinds; high- and low-salience artifacts. These 80 words were used to construct the experimental combinations. The mean proportions (and their standard deviations) of participants who listed the most salient feature for each word appear in Table 1 (regardless of the salience rating they gave it) for each of the four categories. The mean proportion was significantly higher for high-salience artifacts than for high-salience natural kinds [$t(37) = 2.13, p < .05$], which

Table 1
Mean Proportions and Standard Deviations of Participants Who Listed the Most Salient Feature for a Word

| Feature Salience | Ontological Category | | | |
| | Natural Kind | | Artifact | |
	M	SD	M	SD
High	.62	.13	.71	.14
Low	.31	.10	.27	.08

could (but as it turned out, did not) result in larger effects of salience for artifacts than for natural kinds. The mean proportions for low-salience artifacts did not differ significantly [$t(37) = 1.08$].

EXPERIMENT

We factorially manipulated the degree of salience of the heads and the modifiers in N–N combinations to examine the effect of both on the frequency of property-mapping and relation-linking interpretations of N–N combinations. The word pairs in the present study were constructed with the restriction that both constituents be from the same ontological category (natural kind–natural kind or artifact–artifact). Since the noun pairs were constructed randomly, there were very few combinations in which the component nouns shared the same immediate superordinate. Within both artifacts and natural kinds, N–N combinations were constructed that followed one of four patterns: Both constituents had high-salience features (H–H), both had low-salience features (L–L), the modifier noun had a low-salience feature and the head noun had a high-salience feature (L–H), or the modifier noun had a high-salience feature and the head noun had a low-salience feature (H–L). Examples appear in Table 2, and the full list of combinations (together with the identity of the salient feature for high-salient attribute nouns) appears the Appendix.

Predictions

Natural kinds versus artifacts. Property-mapping interpretations were expected to be favored when the terms being combined denote natural kinds, relative to when they denote artifacts. As discussed earlier, natural kind terms are often seen as defined in terms of the properties of what they designate, whereas artifact terms are defined in terms of the functions of their referents. It was expected that natural kind combinations would often be interpreted in terms of properties of the modifier, whereas artifact combinations would center on function and would therefore result in more relation-linking interpretations.

Salience. In general, if a high-salience feature was present on the modifier noun (H–H, H–L), this salient feature was expected to drive the interpretation of the combination, perhaps especially for natural kind combinations. This was expected because of the assumption that the salient attribute will often be the attribute that distinguishes the particular example from other members of its head category or its class, thereby providing a reason for the construction of the combination. The specific predictions for the four groups of salience pairings follow.

H–L: If there was a high-salience feature on the modifier noun, that feature was expected to drive the interpretation of the combination (provided a plausible interpretation was possible based on this feature). For instance, for the word *robin*, one highly salient feature of the modifier is the fact that a robin has a red breast. Therefore, it was assumed that if *robin* was paired with a low-salience head,

Table 2
Example Combinations

H–H	H–L	L–H	L–L
		Natural Kinds	
tiger mouse	carrot beetle	lark zebra	sand petunia
robin termite	zebra butterfly	trout mouse	mule lark
		Artifacts	
mirror sandpaper	sandpaper plate	newspaper spear	desk roof
ball chair	window shirt	scarf mirror	shovel tie

the feature of red breast might well drive the interpretation (e.g., a robin lark being a lark with a red breast).

H–H: Though the predictions for H–L are relatively straightforward, problems arise when a feature on the head noun has equivalent or higher salience than the salient feature on the modifier noun. For instance, the fact that termites eat wood may be a property of equal or higher salience to the robin's red breast in the combination *robin termite*. The highly salient feature on the head may make it more difficult to arrive at an interpretation that follows the basic modifier–head relationship. Participants may opt for one of a number of options in this case. One possibility is that they may produce a relation-linking interpretation (a termite that eats robins). Plausibility may also play a role, leading to even more elaborated interpretations. For instance, Wisniewski (1996) suggests that this combination may actually be interpreted as a termite that eats robins' nests. Another possibility is that subjects might make a switching error, using the highly salient feature of the head noun to drive the interpretation (a robin that eats wood), thus violating the modifier–head relationship that is assumed to be imposed grammatically.

L–H: For L–H pairs, it was assumed that relation linking would be more prevalent than property mapping, since there are no highly salient features to map from the modifier concept to the head concept. Therefore, in cases such as *lark zebra*, a relation-linking interpretation might be encouraged (a zebra that eats larks or a zebra with larks riding on its back), since *lark* does not lend any high-salience features to drive the interpretation. However, the salience of the feature on the head noun might again drive some of the interpretations, as in the H–H pairs. Thus, switching errors (a *lark zebra* being a striped lark) were also expected in this condition, especially if plausibility is also a factor (larks are not likely to be eaten by zebras).

L–L: Pairs with no salient features on either constituent were expected to show a high proportion of relation-linking interpretations, though not necessarily greater than the L–H pairs since those pairs also did not have salient features to map from the modifier. When property mapping occurs, the properties chosen were expected to be more variable than in interpretations with highly salient features. L–L pairs were not expected to show switching errors, since there are no high-salience features on N2 to drive such errors.

Similarity. The word pairs in this study were assessed for similarity, as described below. It was assumed that Wisniewski's (1996) general finding would be borne out

in that more similar pairs would encourage property mapping more easily than would dissimilar pairs. Therefore, an overall positive correlation of property mapping with similarity was expected. However, by restricting the items in a pair to be members of the same ontological category but not forcing high-similarity pairs to be basic-level "sisters" under the same immediate superordinate term, we eliminated the extremes of the similarity scale examined by Wisniewski (1996). Instead, we selected pairs that represented a middle range of similarity.

Method

Participants. Forty-eight undergraduates at the University of Massachusetts participated for partial course credit. None of the participants in this study had participated in the salience norming study.

Procedure. The participants were presented with lists of 40 N–N combinations and were asked to write down the interpretation that they thought best fit each combination. The complete list of noun pairs appears in the Appendix. After interpreting the 40 combinations, the participants were then asked to rate a second set of 40 combinations (those that had been interpreted by a separate group of participants) for similarity. Since a total of 160 N–N pairs were interpreted, 12 sets of interpretations and similarity ratings were obtained for each combination.

Materials. The final lists from which the combinations were drawn consisted of 20 words from each of eight categories (the factorial combination of high-/low-salience N1, high-/low-salience N2, and artifacts/natural kinds). Measures of the feature salience of the words used appear in Table 1. Word pairs were constructed by taking randomly ordered lists from the separate groups and pairing them to create the combinations. The very infrequent word pairs that were judged to be extremely similar (e.g., *zebra horse*) or to be lexicalized (e.g., *truck door*) were eliminated from the final lists and were replaced by additional randomly generated pairs. The result was 20 combinations in each of the eight groups, for a total of 160 combinations. The combinations were then split into four lists of 40 combinations apiece (5 from each group). Examples appear in Table 2; all pairs are presented in the Appendix.

We evaluated the selected pairs after the fact for similarity (cf. Wisniewski, 1996) and N2 dimension relevance (cf. Estes & Glucksberg, 2000). The mean similarity rating (on a scale of 1 to 7, where 7 is *maximum similarity*) over all categories was 2.64 and did not differ significantly among categories [$F(7,150) = 1.16, p > .30$]. The mean standard deviation within each category was 0.99, with a maximum range of 1.3–6.4.

Relevance of N2 dimension was evaluated in essentially the same way as it was evaluated by Estes and Glucksberg (2000). We expected the random pairing procedure to produce few if any pairs in which any salient feature of N1 would be closely related to the relevant dimension of N2. We evaluated this expectation by having 14 University of Massachusetts students generate up to three questions that they would ask about the nouns used as heads (N2) to distinguish one instance of the concept from another. Estes and Glucksberg (2000) and Glucksberg, McGlone, and Manfredi (1997) used the same task to assess dimension relevance, except that they did not restrict their participants to just three questions. We scored the frequency with which any of the three questions a participant produced to a given N2 was related to the salient feature of the high-salient-feature N1 with which it was paired. We used a liberal scoring criterion, so that (for instance) questions about the N2 *butterfly* that involved the color or texture of its wings were judged to be related to the salient feature of the N1 *zebra* with which it was paired, *stripes*.

On average, 7% of the N2 terms elicited one or more questions to which the salient feature of its paired N1 could be considered relevant (a maximum of 35%, a mode of 0%; 4.7% for items with high-

salience heads, and 8.9% for items with low-salience heads). This figure compares with a mean of 35% such relevant questions for the Estes and Glucksberg (2000) low-relevant heads and a mean of 90% for their high-relevant heads. Our N–N compounds clearly fail to exhibit the relevance of an N1 feature to an N2 dimension characteristic of the compounds Estes and Glucksberg used.
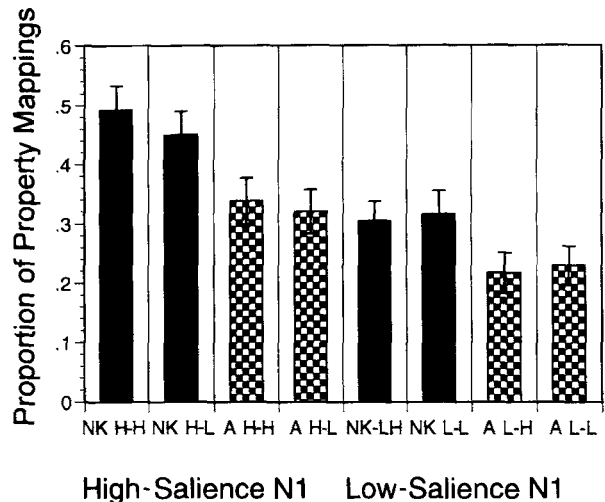
Scoring. Three N–N pairs were eliminated—in one case because one constituent noun proved to be ambiguous (*strainer* was interpreted as something that created difficulty rather than the sieve meaning intended), in a second case because of a spelling error (*tick* instead of *stick*), and in a third case because a constituent noun (*magnolia*) seemed to be unfamiliar to our participants, though this was not the case in the norming study. All responses to the remaining pairs were evaluated in terms of the strategy used to produce the interpretation and any effect salient features had on the interpretation. Responses were scored as property mapping, relation linking, hybrid, switching error, or "unscorable." Interpretations that took a specific attribute from the modifier and applied this attribute to the head noun were scored as property mapping (e.g., a *robin termite* being a red-breasted termite). Interpretations that posited any type of relation between the two nouns were scored as relation linking (e.g., a *robin termite* being a termite that eats robins or robins' nests). Responses that took their features from both constituents and that resulted in an interpretation that was a combination of both constituents were classified as hybrids (e.g., a *robin termite* being a creature that is both a robin and a termite). Interpretations that applied a feature of the head noun to the modifier noun, thus violating the modifier–head relationship, were classified as switching errors (e.g., a *robin termite* being a robin that eats wood). Finally, any interpretations that merely defined one constituent (e.g., *bird; bug*) or that simply listed an attribute of one or the other without describing the role in the combination (e.g., *red; eats*) were coded as "unscorable." The rationale for this was that it was unclear whether these participants were applying the property to the combination rather than just describing one constituent. Also coded as "unscorable" were any anomalous interpretations and blank responses. These responses were left in to avoid biasing the data to appear as though only property-mapping or relation-linking interpretations occurred. Finally, when an interpretation was judged to be derived via property mapping, the effect of any salient features was scored to see whether those salient features drove the interpretation.

A subset of the interpretations (38%) were scored by a second interpreter using the above definitions. Initial interrater agreement was 94%, and all differences were resolved by discussion.

Results

The proportions of each type of response given to each type of N–N compound appear in Table 3. We first compare the responses to natural kinds versus artifacts and then present analyses of the effects of salience and of similarity.

Natural kinds versus artifacts. Figure 1 depicts the proportion of property-mapping interpretations given in all eight groups. Analyses of variance (ANOVAs; one with subjects as the random factor, one with items) were con-

Table 3
Proportions of Response Types Given

Response Type	Natural Kinds				Artifacts			
	H–H	H–L	L–H	L–L	H–H	H–L	L–H	L–L
Property mapping	.50	.46	.35	.35	.30	.31	.23	.24
Relation linking	.14	.20	.17	.23	.30	.32	.40	.38
Hybrid	.01	.00	.01	.04	.00	.01	.00	.03
Switching errors	.12	.09	.14	.12	.16	.07	.16	.08
Unscorable	.23	.25	.33	.33	.24	.29	.21	.27

High-Salience N1   Low-Salience N1

Figure 1. Proportion of property-mapping interpretations in all eight groups. NK, natural kinds; A, artifacts. Error bars indicate 1 SE.

ducted on the proportions of property-mapping responses and the proportions of relation-linking responses. The factors were natural kind versus artifact, feature salience of N1, and feature salience of N2 (within-subjects, but between-items, factors). As expected, natural kinds produced significantly more property-mapping interpretations than did artifacts (42% vs. 27%) [$F_1(1,47) = 31.13$, $p < .0001$; $F_2(1,149) = 25.48, p < .001$]. Conversely, artifacts produced more relation-linking interpretations than did natural kinds (35% vs. 18%) [$F_1(1,47) = 47.19$, $p < .001$; $F_2(1,149) = 35.92; p < .001$]. A further analysis, conducted on just natural kind terms, compared the frequency of property-mapping and relation-linking interpretations. Property-mapping interpretations were more frequent (40% vs. 18%) [$F_1(1,47) = 31.76, p < .001; F_2(1,79) = 39.37, p < .001$]. A similar analysis conducted on artifacts indicated that the difference was significant (in the opposite direction) only by items (26% vs. 32%) [$F_1(1,47) = 2.48, p > .10; F_2(1,76) = 3.99, p < .05$]. The proportion of property mapping for natural kinds never exceeded 50%. This proportion might have been greater but was lowered by the fact that unscorable interpretations, which included blank responses and anomalous responses, were included in computing proportions of each scoreable response. These results lead to the conclusion that, when N–N combinations are successfully interpreted, property mapping is the predominant strategy for natural kinds.

Salience. Overall, without dividing the groups into ontological categories, modifier salience produced the predicted effect. N–N pairs with high-salience modifiers (H–H, H–L) produced more property-mapping interpretations than did those combinations with low-salience modifiers (L–H, L–L) (40% vs. 29%) [$F_1(1,47) = 27.39$, $p < .001; F_2(1,149) = 12.04, p < .001$] (see Figure 1). Figure 1 suggests that the effect of N1 salience may be greater for natural kinds than for artifact terms, but the

interaction of N1 salience and ontological categories was nonsignificant [$F_1(1,47) = 2.07, p > .15; F_2(1,149) = 1.52, p > .20$].

Salience of N2 did not affect the proportion of property-mapping interpretations or the proportion of relation-linking interpretations ($F < 1$ in both cases). We had suggested specifically that H–H terms might elicit fewer property-mapping interpretations and more relation-linking interpretations than H–L terms, but there was no evidence that this was the case (all $ps > .20$). However, in an analysis conducted on the proportion of switching errors, high-salience N2s did lead to more switching errors than did low-salience N2s [$F_1(1,47) = 11.97, p < .001; F_2(1,149) = 11.85, p < .001$]. This effect was stronger for artifacts than for natural kinds [interaction, $F_1(1,47) = 5.92, p < .02; F_2(1,149) = 4.20, p < .05$]. Examples of switching errors produced by the participants include the following: *wasp mouse* being a wasp that lives in the walls; *carrot zebra* being a striped carrot; *scarf mirror* being a reflective or silver scarf; *pencil sandpaper* being a gritty or rough pencil.

Low-salience modifiers produced more relation-linking interpretations than did high—salience modifiers (30% for L–L and L–H pairs vs. 24% for H–L and H–H pairs) [$F_1(1,47) = 11.69, p < .01; F_2(1,149) = 4.230, p < .05$] (see Table 3). This effect was significant as predicted for artifacts (39% vs. 31%) [$F_1(1,47) = 8.43, p < .01; F_2(1,73) = 4.01, p = .05$]. Natural kinds, however, did not show the predicted difference (20% vs. 17%) [$F_1(1,47) = 1.82, p > .15; F_2 < 1$]. Instead, natural kind L–H and L–L pairs produced significantly more "other" responses (hybrids, switching errors, and unscorable responses) than did the H–H and H–L pairs (45% vs. 35%) [$F_1(1,47) = 15.44, p < .001; F_2(1,76) = 7.02, p < .01$], indicating that the participants either were making switching errors or were failing to interpret the combination at all, giving either an anomalous response or no response. Natural kinds do not seem to default to a relation-linking interpretation.

Figure 2 depicts the proportion of property-mapping interpretations that were driven by a salient feature on the modifier. Only combinations with high-salience modifiers are included in these analyses. Overall, in combinations with high-salience modifiers, the property-mapping interpretations given were generally driven by the salient feature on the modifier noun (a mean of 68%). In an ANOVA of the proportion of property-mapping interpretations that involved the salient feature, interpretations were more likely to be driven by the salient feature for natural kinds than for artifacts (75% vs. 61%) [$F_1(1,27) = 11.47, p < .01; F_2(1,74) = 4.03, p < .05$; note that the df was low in the by-subjects analysis because many participants had a zero frequency of property-mapping interpretations in at least one condition, which was true of only a very few items]. However, as Figure 2 reveals, only the artifacts H–L group was below 70%. This observation is supported in the by-items analysis by a nearly significant main effect of head salience [$F_2(1,74) = 3.52, p = .06$] and an interaction of ontological category × head salience [$F(1,74) = 4.02, p < .05$]. However, neither effect approached significance in the by-subjects analysis ($F < 1$), presumably because of the large amount of missing data.

The above effect of head salience was unexpected and prompted post hoc evaluation of the materials in the artifact H–L group. Such evaluation revealed that several items in the artifact H–L group did not behave as expected. This may be the result of constraints imposed by the head noun that are specific to the randomly chosen modifier. One example is the combination *knife carpet*. In isolation, the word *carpet* was rated as low salience because it did not produce one single feature that reached the criterion for high salience. When the combinations were constructed, however, *carpet* was combined randomly with *knife*. This random pairing lead to an implausibility in mapping the salient feature of the modifier (a knife is sharp) to a conflicting dimension of the head (a carpet is something that is generally soft or plush, though neither were high salience in isolation). This implausibility may have blocked the potential property-mapping interpretation based on the salient modifier feature (e.g., a sharp carpet). We argue that, though one may not need a particular highly relevant dimension on N2 in order for property mapping to occur, one cannot have a conflicting dimension. In a more general sense, this also may be because the essential nature of artifacts is one of function, and, therefore, artifact N2s will not accept property attributions that conflict with the artifacts' function.

Finally, though more hybrids were expected in H–H pairs, these responses were extremely rare (see Table 3), and, as a result, no analyses were run on the proportion of hybrid responses. This may be due to the construction of our noun pairs. Wisniewski (1996) found hybrids in his similar pairs (20%) but not in his dissimilar pairs (0%). Our pairs may more closely pattern with his dissimilar pairs. Hybrids may be more likely to occur when the component nouns share the same immediate superordinate. Estes and Glucksberg (2000) also found an extremely low occurrence of hybrids.
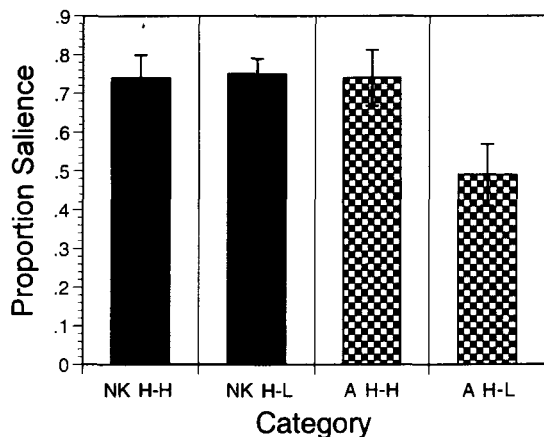

Figure 2. Proportion of property-mapping interpretations that reflected the salient property. NK, natural kinds; A, artifacts. Error bars indicate 1 *SE*.

Similarity. The expected replication of Wisniewski's (1996) similarity correlation was not borne out. There was no correlation between similarity and proportion of property mapping overall ($r = -.06$). Furthermore, there was no correlation within either natural kinds ($r = -.08$) or artifacts ($r = -.22$) alone. Individual correlations within the eight salience groups also failed to reach significance: natural kinds H–H ($r = -.11$); natural kinds H–L ($r = .14$); natural kinds L–H ($r = -.02$); natural kinds L–L ($r = -.52$); artifacts H–H ($r = -.01$); artifacts H–L ($r = -.31$); artifacts L–H ($r = .00$); artifacts L–L ($r = .12$). The most significant correlation ($r = -.52$) was actually in the wrong direction but was largely due to one very deviant point. An argument based on similarity would predict a strong positive correlation between proportion of property mapping and assessed similarity. This lack of replication may have been due in part to the fact that the combinations in this study were constrained to be within ontological categories but were generally not members of closely related categories; the similarity effect observed in previous studies may have been limited to items at extremes of the range of similarity.

Conclusions

The present experiment had three primary goals: (1) to explore how the ontological category of the combination influences the strategy chosen for interpretation, (2) to examine what effects simple feature salience has on the produced interpretations of conceptual combinations, and (3) to investigate further whether or not such effects were correlated with assessed similarity.

Property-mapping interpretations were, as expected, more prevalent for natural kind pairs than for artifact pairs. Conversely, relation-linking interpretations were more frequent for artifacts than for natural kinds. When property mapping was not the chosen strategy, it was assumed that interpretations would revert to relation-linking strategies. This was the case for artifacts but not for natural kinds.

The fact that artifacts did show the expected tendency toward relation-linking interpretations may be due to the relative ease with which relations can be posited between artifacts. Salient properties of artifacts are generally those of their function or properties intricately related to that function. Such functional attributes might well encourage relation-linking interpretations. Relation linking was not consistently the default strategy for natural kinds. When property-mapping interpretations were not chosen, natural kind pairs tended to produce more "other" interpretations.

The manipulation of simple feature salience did have the expected influence on the interpretations produced. Considering our materials as a whole, when there was a high-feature-salience modifier more property mapping was observed; for pairs with low-salience modifiers, more relation linking was observed. Furthermore, when property-mapping interpretations were produced, the salient feature on the modifier was the predominant feature used in the generation of the interpretation. These results

support the view that when salient features are present on the modifier, those features are selected to drive the interpretation of an N–N combination even in the absence of highly relevant dimensions, provided there is not a conflicting dimension on N2. Such salient features may provide the needed justification for the creation of the combination in that they may be the features that distinguish a particular instance from the larger set denoted by its head. Therefore, though we generally accept the model put forth by Estes and Glucksberg (2000), we argue that the model may overstate the need for interaction.

Further evidence for the effect of simple feature salience can be found in our results that show that high-salience N2s also played a role in the interpretations produced. More switching errors were made on N–N combinations with high-salience features on N2. In these cases, the high-salience feature on the head was important or salient enough to violate the standard modifier– head relationship. We did not expect this result, accepting the standard linguistic analysis that N2 must be the head of an N–N compound. Apparently, our participants were willing to overlook this linguistic constraint in their effort to produce an interpretation that treated a highly salient feature as a property of the entity that the compound refers to.

The presence of a high-salience feature on N2 sometimes had a different effect. Consider the example *robin termite*. Here, N1 *robin* has the high-salience attribute of being red-breasted. However, N2 termite also has a high-salience attribute—namely, that it eats wood. Our participants frequently seemed to abstract away from this attribute to an underlying dimension (eating something in this case) and interpret the compound in terms of this dimension. Thus, an occasional interpretation was a termite that eats robins or a termite that eats robins' nests, rather than a red-breasted termite. Even though there was a high-salience feature on N1 that could provide an entirely plausible interpretation, there was a high-salience feature on N2 that was more constraining. The N2 feature exercised its constraint by specifying a highly relevant dimension of N2, much as proposed by Estes and Glucksberg (2000), except that in this case satisfying the dimension resulted in a relation-linking, not a property-mapping, interpretation.

Finally, previous research has revealed strong positive correlations between the proportion of property-mapping interpretations and the assessed similarity of two nouns. This effect was not found in the present study. There were no correlations between property mapping and similarity. The constraint of keeping the word pairs within ontological categories may have increased the similarity of our least similar items in comparison with previous studies, while the lack of a constraint of being closely related basic-level terms certainly lowered the similarity of our most similar terms. As a result, our items spanned a moderate range of similarity. A greater effect of similarity may be observable at the extreme ends of a similarity scale, but such an effect is not present in our data.

This argument is not intended to entirely disavow the possible role of similarity. There is a great deal of research

to support it. Rather, the present experiment suggests that other strategies may take precedence over similarity comparisons. One such strategy relies on the salience of particular features. Information about which features are likely candidates for the creation of an N–N combination may lead to a more economical strategy of interpretation than exhaustive similarity comparisons and alignment procedures. Use of salient features in creating N–N combinations may provide a more economical method than similarity comparisons, since the use of such features can immediately draw attention to the distinguishing characteristics without having to extensively compare two constituents. However, in the absence of high-salience features on either noun, alignment and comparison may remain a viable strategy. Our argument, therefore, would still predict a correlation with similarity in the L–L condition. The fact that such a correlation was not found (and, in fact, was the in the opposite direction at $r = -.52$) suggests that perhaps Wisniewski's (1996) similarity results are limited to his high-similarity items or to the lowest of his low-similarity items.

The present results were obtained using randomly paired nouns, suggesting that simple feature salience is an important factor in its own right even in the absence of head dimension relevance. We argue that the process of interpreting N–N combinations involves the use of highly salient attributes of N1 in combination, if they exist, with highly relevant dimensions of N2. In property mapping, the use of salient features alone can drive the interpretation of a particular combination even if that combination lacks a highly relevant dimension on N2. However, if N2 has a dimension that strongly conflicts with the feature to be mapped from N1, then the process runs into difficulty. While we accept the general premise of an interactive model, we argue that such interaction has its strongest effects at the extremes (in cases where there is either no relevance or no salience) but that, in the intermediate range, simple feature salience alone can drive an interpretation.

REFERENCES

AHN, W.-K. (1998). Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition*, **69**, 135-178.

BARR, R. A., & CAPLAN, L. J. (1987). Category representations and their implications for category structure. *Memory & Cognition*, **15**, 397-418.

BARTON, M. E., & KOMATSU, L. K. (1989). Defining features of natural kinds and artifacts. *Journal of Psycholinguistic Research*, **18**, 433-447.

BATTIG, W. F., & MONTAGUE, W. E. (1969). Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology Monographs*, **80**(3, Pt. 2).

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

ESTES, Z., & GLUCKSBERG, S. (2000). Interactive property attribution in concept combination. *Memory & Cognition*, **28**, 28-34.

GAGNÉ, C. (2000). Relation-based combinations versus property-based combinations: A test of the CARIN theory and dual-process theory of conceptual combination. *Journal of Memory & Language*, **42**, 365-389.

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

GLUCKSBERG, S., MCGLONE, M. S., & MANFREDI, D. (1997). Property attribution in metaphor comprehension. *Journal of Memory & Language*, **36**, 50-67.

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

KEIL, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: Harvard University Press.

MARKMAN, A. B., & GENTNER, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, **23**, 431-467.

MARKMAN, A. B., & WISNIEWSKI, E. (1997). Similar and different: The differentiation of basic-level categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 54-70.

MURPHY, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, **12**, 529-562.

MURPHY, G. L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory & Language*, **29**, 259-288.

RIPS, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou (Ed.), *Similarity and analogical reasoning* (pp. 27-47). New York: Cambridge University Press.

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). Construal and similarity in conceptual combination. *Journal of Memory & Language*, **35**, 434-453.

WISNIEWSKI, E. J. (1997a). Conceptual combination: Possibilities and esthetics. In T. B. Ward, S. M. Smith, & J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 51-81). Washington, DC: American Psychological Association.

WISNIEWSKI, E. J. (1997b). When concepts combine. *Psychonomic Bulletin & Review*, **4**, 167-183.

WISNIEWSKI, E. J. (1998). Property instantiation in conceptual combination. *Memory & Cognition*, **26**, 1330-1347.

WISNIEWSKI, E. J., & GENTNER, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 241-284). Amsterdam: North-Holland.

WISNIEWSKI, E. J., & LOVE, B. C. (1998). Properties versus relations in conceptual combination. *Journal of Memory & Language*, **38**, 177-202.

APPENDIX
Lists of Combinations (With Most Salient Properties of High-Salient Terms)

H–L	L–H	H–H	L–L
		Natural Kinds	
flamingo (pink) potato	sand centipede (100 legs)	mouse (small) flamingo (pink)	potato trout
rhinoceros (horn) berry	pine minnow (small)	centipede (100 legs) rhinoceros (horn)	beetle pine
carrot (orange) beetle	mule rhinoceros (horn)	minnow (small) carrot (orange)	elm horse
leech (sucks blood) sand	sparrow termite (eats wood)	ant (small) wasp (stings)	daisy sparrow
wasp (stings) elm	lily crow (black)	termite (eats wood) rock (hard)	petunia carnation
rock (hard) lily	carnation tiger (stripes)	crow (black) worm (slimy)	coconut butterfly
worm (slimy) petunia	butterfly raisin (wrinkled)	tiger (stripes) mouse	horse daisy
moose (antlers) coconut	berry rock (hard)	banana (yellow) robin (red breast)	acorn magnolia
robin (red breast) mule	magnolia flamingo (pink)	zebra (stripes) raisin (wrinkled)	trout acorn
raisin (wrinkled) lark	crocus rooster (crows)	rooster (crows) leech (sucks blood)	lark crocus
ant (small) carnation	trout mouse	moose (antlers) ant (small)	berry sand
termite (eats wood) sparrow	petunia ant (small)	robin (red breast) termite (eats wood)	sand petunia
crow (black) daisy	coconut moose (antlers)	rhinoceros (horn) crow (black)	carnation coconut
mouse (small) pine	horse robin (red breast)	wasp (stings) mouse	lily mule
centipede (100 legs) horse	potato wasp (stings)	raisin (wrinkled) centipede (100 legs)	sparrow potato
tiger (stripes) crocus	daisy carrot (orange)	carrot (orange) zebra (stripes)	crocus berry
zebra (stripes) butterfly	beetle leech (sucks blood)	rock (hard) tiger (stripes)	butterfly beetle
minnow (small) acorn	acorn worm (slimy)	leech (sucks blood) minnow (small)	mule lark
banana (yellow) magnolia	lark zebra (stripes)	worm (slimy) banana (yellow)	magnolia elm
rooster (crows) trout	elm banana (yellow)	flamingo (pink) rooster (crows)	pine lily
		Artifacts	
chair (sitting) boot	doll triangle (3 sides)	ball (round) chair (sitting)	car scarf
bicycle (2 wheels) shovel	pencil sandpaper (rough)	mirror (reflective) sandpaper (rough)	plate belt
sandpaper (rough) plate	boot strainer (holes)	spear (pointed) strainer (holes)	scarf pencil
stove (hot) magazine	chimney drum (sound)	window (transparent) door (open/close)	roof shirt
strainer (holes) desk	newspaper spear (pointed)	knife (sharp) drum (sound)	belt chimney
door (open/close) truck	bureau hat (on head)	ladder (rungs) lamp (light)	carpet bureau
drum (sound) roof	desk door (open/close)	hat (on head) sweater (warm)	shirt carpet
sweater (warm) newspaper	stick lamp (light)	freezer (cold) triangle (3 sides)	tie television
triangle (3 sides) car	truck sweater (warm)	lamp (light) taxi (yellow)	magazine stick
taxi (yellow) doll	television puzzle (pieces)	puzzle (pieces) stove (hot)	shovel tie
spear (pointed) belt	car ball (round)	bicycle (2 wheels) spear (pointed)	doll shovel
ball (round) bureau	plate bicycle (2 wheels)	triangle (3 sides) ball (round)	pencil boot
window (transparent) shirt	scarf mirror (reflective)	sandpaper (rough) window (transparent)	boot desk
mirror (reflective) chimney	roof stove (hot)	strainer (holes) puzzle (pieces)	chimney plate
knife (sharp) carpet	belt window (transparent)	drum (sound) mirror (reflective)	newspaper magazine
lamp (light) television	carpet ladder (rungs)	stove (hot) knife (sharp)	bureau truck
freezer (cold) stick	shirt knife (sharp)	door (open/close) bicycle (2 wheels)	desk roof
ladder (rungs) pencil	tie freezer (cold)	sweater (warm) ladder (rungs)	stick newspaper
hat (on head) tie	magazine taxi (yellow)	chair (sitting) hat (on head)	truck doll
puzzle (pieces) scarf	shovel chair (sitting)	taxi (yellow) freezer (cold)	television car