

Relational Similarity and the Nonindependence of Features in Similarity Judgments

ROBERT L. GOLDSTONE, DOUGLAS L. MEDIN, AND DEDRE GENTNER

University of Illinois

Four experiments examined the hypothesis that simple attributional features and relational features operate differently in the determination of similarity judgments. Forced choice similarity judgments ("Is *X* or *Y* more similar to *Z*?") and similarity rating tasks demonstrate that making the same featural change in two geometric stimuli unequally affects their judged similarity to a third stimulus (the comparison stimulus). More specifically, a featural change that causes stimuli to be more superficially similar and less relationally similar increases judged similarity if it occurs in stimuli that already share many superficial attributes, and decreases similarity if it occurs in stimuli that do not share as many superficial attributes. These results argue against an assumption of feature independence which asserts that the degree to which a feature shared by two objects affects similarity is independent of the other features shared by the objects. The MAX hypothesis is introduced, in which attributional and relational similarities are separately pooled, and shared features affect similarity more if the pool they are in is already relatively large. The results support claims that relations and attributes are psychologically distinct and that formal measures of similarity should not treat all types of matching features equally. © 1991 Academic Press, Inc.

The question of what makes two objects psychologically similar has considerable significance for much of cognitive psychology. Models of categorization have claimed that an item is categorized as an *A* and not a *B* if it is more *similar* to *A*'s prototype or exemplars than *B*'s (Ashby & Gott, 1988; Medin & Shaffer, 1978; Reed, 1972). Models of memory have claimed that *X* reminds people of *Y* if *X* is *similar* to *Y* (Hintzman, 1986; Raaijmakers & Shiffrin, 1981), and that *similar* items are clustered together in memory. Learning theorists have claimed that learning task *A* is facilitated if it is *similar* to task *B* which is already part of the learner's repertoire. In the problem solving literature, if someone has solved a problem *similar* to problem *Y*, then attempts to solve *Y* will be more

Preparation of this article was supported in part by National Science Foundation Grants BNS-87-20301 and BNS-88-12193, and by the Office of Naval Research under Contract No. N00014-85-0030 awarded to the third author. We thank Donaldson Dulany, Eldar Shafir, Gordon Logan, Brian Ross, Ed Shoben, and three anonymous reviewers for valuable comments on earlier drafts of this article. Correspondence and reprint requests should be addressed to Robert Goldstone, Department of Psychology, 330 Packard Drive, University of Michigan, Ann Arbor, MI 48104.

successful than if no such similar problem has been solved (Gick & Holyoak, 1980; Ross, 1984). The degree to which these theories have explanatory power in their domains partially rests on their ability to specify what similarity is and how to measure it.

This paper investigates the role of relational similarity for overall similarity judgments. Previous researchers have urged psychologists to consider the importance of structural and relational aspects of stimuli (Biederman, 1987; Gentner, 1983; Hock, Tromley, & Polmann, 1988; Murphy & Medin, 1985; Ortony, 1979; Palmer, 1978; Tversky & Hemenway, 1984). Intuitively, the notion of relational similarity can be grasped by the fact that seeing a tiger and her cub in a zoo is similar in some ways to seeing a robin and her nestling in a tree. If there is a sense in which the scenes are similar, it is probably not because the animals themselves are especially similar; tigers and robins may share some features, but no more than are shared by many animals (*two eyes, living, locomotion system, etc.*). An account of similarity is required that includes the notion that the two scenes are similar because the relation "*X is the mother of Y*" is the same in both cases.

One strategy for incorporating relations into similarity is to treat them in exactly the same way that simple features are treated. Under this interpretation, the tiger and the robin scenes are similar because both contain the abstract feature "mother next to child." This paper explores the alternative hypothesis that similarity models cannot adequately predict similarity judgments if relational features and simple features (attributes) are undifferentiated. Given the premise that relations and attributes are both likely to influence similarity judgments, this paper focuses on the factors that affect the influence each has. We will argue for a model of similarity, MAX, which specifically differentiates attributes from relations, and which weights the two types of similarity differently, depending on the other features present in the scene. A theory which simply *lists* the separate features of an object, even if it includes relational features, will miss important generalizations concerning the feature dependencies we observe.

DEFINITION OF ATTRIBUTES AND RELATIONS

The term "Attribute" will be used to refer to any single component or property of a stimulus. Attributes can be thought of as predicates which take single arguments; for example, green is an attribute of grass because we can write "GREEN(GRASS)" where the predicate GREEN takes the single argument GRASS. Attributes do not have to be concrete. Abstract properties such as "evil" and "pretty" could be attributes in a representation system. "Relation" refers to the description of a connection between two or more attributes. Relations are predicates which take two or

more arguments. Unlike attributes, relations can only be ascribed to an object by reference to other attributes/relations. If a picture is composed of a red square next to a red circle, then the picture has the relation "Same shading."

A propositional representation system is often used to indicate how a relation is bound to arguments that are objects or attributes (Palmer, 1975; Gentner, 1983; Winston, 1980). A geometric scene can be propositionally described as "ABOVE (CIRCLE, TRIANGLE)." The order of arguments is critical in propositional representations; the same attributes and the same relation, ordered differently, may yield a different interpretation. While "ABOVE (TRIANGLE, CIRCLE)" describes a triangle above a circle, "ABOVE (CIRCLE, TRIANGLE)" describes a circle above triangle. A second scene may be described as "ABOVE (SQUARE, STAR)." This second scene is *relationally similar* to the first scene, because the relation "ABOVE" is identical. A third scene, "RIGHT-OF (TRIANGLE, CIRCLE)," is *attributionally similar* to the first scene because it contains the same primitive features, although the relation between the features differs. Attributional similarity corresponds to Gentner's (1983) Mere Appearance similarity, and relational similarity corresponds to her Analogy.

It is tempting to think of relational features as global features or properties. Relations are more global than attributes in the sense that a relation may bind two or more arguments. In general, however, relations cannot be equated with global features because a relation may involve only a subset of a scene (e.g., eyes being the same color on a face) and there is no reason why a single large feature could not encompass a larger region than a single relation. We further reinforce the distinction between relational features and global features in the General Discussion of this paper (see also Navon, 1977).

What will count as relations and attributes is intrinsically connected with choices concerning the representation system. One system may treat points of light as primitive attributes; another system may treat lines as primitive attributes. In the former system, lines may be defined relationally as a series of point attributes associated via "Adjacent-to" relations. In the latter system, the lines themselves will be the attributes, and the relations involved will reflect the configurations of lines. "Parallel to" is a possible relation for the system in which lines are attributes, but not in the system with points as attributes. Whether a property is an attribute can only be assessed once a particular representation has been proposed.

FEATURE INDEPENDENCE AND OTHER POSSIBILITIES

Tversky's contrast model (1977) will be taken as a starting point, because of its commanding influence and elegance. The mathematical equation

$$\text{SIM}(A, B) = X * f(A \cap B) - Y * f(A - B) - Z * f(B - A)$$

defines the similarity of A to B to be a function of the features that A and B share, minus the features that A has that B does not have, minus again the features that B has that A does not have. The X , Y , and Z terms are simply weighting terms that depend on the subject's task and the stimuli. Although it is not an inherent requirement of the contrast model, Tversky often makes the further assumption that the function ' f ' satisfies feature additivity such that $f(X \cap Y)$ is expressible as the sum of the measures of all the features that belong to both A and B .

The model relies on three assumptions: matching, monotonicity, and independence. The matching assumption simply states that the similarity of A to B will be a function of their shared features and their distinctive features. This excludes from analysis any features that are absent in both A and B . The monotonicity assumption states that similarity increases when the number of shared features between A and B increases and/or when the number of distinctive features between A and B decreases. The independence assumption states that the joint effect of two feature components in determining similarity is independent of the fixed level of the third component. It is incorrect to equate component independence with feature independence; there are always three components that enter into a similarity judgment: $(A \cap B)$, $(A - B)$, and $(B - A)$. Components are made of sets of features, but it would be possible for feature independence to be violated without having component independence violated. Gati and Tversky (1984) have found violations of feature independence when qualitative and quantitative features occur in the same stimulus and have suggested other possible violations of independence. They discuss examples in which independence is violated because adding one feature causes a global feature to be added in one case but not in another. In their example, \perp and \sqsubset are more similar to each other than are \sqsubset and \square even though adding the shared feature $|$ to both \perp and \sqsubset should increase their similarity as computed by the contrast model. The reason why the additional feature does not increase similarity is that \square has a global feature (closure) that \sqsubset does not. Consequently, Gati and Tversky argue that feature independence appears to be violated when new features combine with old features to produce emergent features.

Figure 1 presents a *prime facie* counterexample to feature independence, analogous to Gati and Tversky's example. When subjects ($N = 28$) are asked to circle the figure set that is more similar to the triangles in A, 89.3% of subjects circled the pair of squares. However, in B 100% of subjects circled the figure with two circles. In going from A to B the same physical feature was added to each picture—a single square. Adding the same feature to each set *seems* to have reversed the original similarity




FIG. 1. Violation of the independence assumption, assuming a feature vocabulary of square, triangle, and circle.

judgments. Statistically speaking, there is an interaction between the bottom element and the top two elements; each is not factored in separately in determining similarity.

In order to explain Fig. 1's judgment reversal, some notion of relations between features is required. However, within the framework of the contrast model one could simply claim that relations should be understood as features. In the past, researchers have considered features that were global (Navon, 1977), abstract, holistic (Kemler, 1983; Nickerson, 1972), and integral (Garner, 1974; Monahan & Lockhead, 1977; Smith & Kemler, 1978). Perhaps relations could be thought of as another type of feature, or reduced to one of these proposed types. As features, relations would be included in Tversky's original equation, and these "relational features" would consequently have an impact on similarity judgments. In accordance with this suggestion, Tversky and Gati (1982) have suggested that relational properties such as "Symmetry" can be treated as features. So, in Fig. 1B, the stimulus with squares would have the features "square on top" and "vertically aligned," and would also have the *relational feature* of all three shapes being identical. If this reformulation of relations as features succeeds, the contrast model would require no algebraic adjustment to account for relational similarities.

We shall argue that the "relations as features" hypothesis fails in two distinct ways. First, if one combines the contrast model with the feature additivity assumption, one can show violations of feature independence.

similarity. The middle element of *A* is closer to *T*'s middle element than is *B*'s middle element. Exactly the same feature change occurs between *C* and *D*.

By examining the differences between similarity ratings for each pair, we can see whether changing a single feature has an uniform effect on similarity ratings. Three possible rating patterns are possible:

$$\begin{aligned} \text{SIM}(D,T) - \text{SIM}(B,T) &= \text{SIM}(C,T) - \text{SIM}(A,T) \\ \text{SIM}(D,T) - \text{SIM}(B,T) &> \text{SIM}(C,T) - \text{SIM}(A,T) \\ \text{SIM}(D,T) - \text{SIM}(B,T) &< \text{SIM}(C,T) - \text{SIM}(A,T). \end{aligned}$$

The first expression asserts that the difference between how similar *D* is to *T* and how similar *B* is to *T* will be equal to the difference between how similar *C* is to *T* and how similar *A* is to *T*. This is the result required by similarity models that assume feature additivity and a direct mapping of featural differences onto similarity ratings. The same feature change occurs between *B* and *D* and between *A* and *C* ("X" → square). If features are independent with regard to their effect on similarity, then the effect of changing an "X" to a square will be the same for *A* and *B*.

The second expression implies that an independence assumption is incorrect; increasing attributional similarity increases rated similarity *more* when there are attributional similarities between the stimuli than when there are relational similarities. We will argue for a simple model, MAX, which predicts these results, stating that (1) attributional similarities are pooled together, and relational similarities are pooled together, and (2) the weight that a similarity has on the final similarity judgment increases with the size of the pool to which it belongs. If there are more shared relations between *A* and *B*, then when making a similarity judgment, each relation will get more weight; if there are more shared attributes, then attributes will be relatively more weighted. MAX consists of a functional differentiation claim and a selective attention claim. MAX claims that relations and attributes are psychologically distinct, that similarities are classified as relational or attributional types, and that people attend more to similarities belonging to whichever similarity type is greatest.

The third expression would also violate an independence assumption, because adding the same attributional similarity increases similarity *more* if it is in the context of strong relational similarities than if it is in the context of strong attributional similarities. This is the result predicted by a MIN model, which states that the similarities belonging to the smaller similarity type are selectively weighted during similarity judgments.

A pithy description of the models is

MAX: "When relationally close, weight relations more; when attributionally close, weight attributes more."

MIN: "When relationally close, weight attributes more; when attributionally close, weight relations more."

Independence: "How you weight something doesn't depend on what's already there."

The two important points about the two models (MAX and MIN) are that they both violate independence of features, and they both require two distinct similarity types: relational and attributional. If there are systematic biases in favor of either MAX or MIN, then a minimal requirement of any similarity model is that each feature be classified as to whether it is a relational or attributional feature. A model that treats relations as it treats attributes would not accommodate the nonindependence finding. Moreover, even a model which selectively weights relations or attributes in only a *context-independent manner* would not capture MAX or MIN effects.

EXPERIMENT 1

Method

Materials. Stimuli consisted of questions of the form "T IS LIKE Y OR Z." T, Y, and Z were pictures of geometrical objects or simple natural objects; each picture contained between two and five elements, with elements varying on size, shading, shape, and position. There were eight groups of pictures, each group containing five pictures. Each group had a single T picture for each comparison. The other four members of the groups were of the four types in Fig. 2. Comparing horizontally, members A and C shared one attributional similarity with T that B and D did not. Vertically, C and D shared one relational similarity with T that A and B did not. In addition, A was attributionally closer to T than was C, and B was attributionally closer to T than was D in the same way. See Appendix A for the complete stimulus set. For each group, six "T IS LIKE Y OR Z" questions were posed, with each member of {A, B, C, D} being compared with each other member (yielding: A or B, A or C, A or D, B or C, B or D, and C or D). Examples of relational similarities include Same-shape-as, Same-shading-as, Taller-than, Larger-than, and Same-orientation-as. Examples of attributional similarities include Square, Circle, Striped, Black, 2 inches tall, and Upward-pointing.

Whether picture A or B came first in the question was randomized for one test set, and the order was reversed in a second test set. Six comparisons were arranged randomly on a page, with the restriction that no two comparisons from the same group be on the same page. Interspersed among distractor questions were two questions for each group of the form "T IS LIKE A OR B" and "T IS LIKE C OR D." The stimuli were designed on the Macintosh computer and printed on a laser printer.

Procedure. Participants were given the handout and were told to circle the choice which was most similar to the comparison figure. The task took approximately 15 min to complete.

Subjects. Twenty-six subjects were given six groups of pictures (the first six groups in Appendix A). Forty-nine subjects (including the 26 subjects above) were given two other groups of pictures (the last two groups in Appendix A). Participants were University of

Illinois undergraduate students taking part in the experiment for partial fulfillment of a course requirement. Groups of 10 to 16 subjects were tested simultaneously.

Results and Discussion

The forced choice similarity judgment task permits only an ordinal measure of similarity. Since MAX and MIN are only supported if changing the same feature in two stimuli causes a decrease in relative similarity in one case and an increase in relative similarity in the other, we have a very conservative measure of the number of subjects who act in accordance with MAX or MIN. Consequently, it is inappropriate to compare the number of times the independence model is supported with the number of times MAX and MIN are supported. The number of responses falling into the "independent" group will contain responses based on independent features, *plus* all responses based on MAX or MIN that are too small to be picked up by an ordinal measure. Instead, systematic departures from independence are measured by comparing the number of MAX responses with the number of MIN responses. A certain number of "noise" responses could be expected that would yield MAX or MIN responses, but these error responses should be equal for MAX and MIN under the independence assumption.

Collapsing over all 49 subjects, and all nine picture groups, MAX effects were observed 26 times, MIN effect were observed 5 times, and 223 responses fell into the independence group. The number of responses supporting MAX is significantly greater than the numbers supporting MIN ($Z = 3.78, p < .001$). The number of MAX responses and MIN responses is broken down by picture group in Appendix A.

Pictures with relational similarities were in general judged to be more similar to *T* than pictures with attributional similarities. Out of the 223 responses that did not show MAX or MIN effects, 179 of these responses showed a preference for relational similarities over attributional similarities, compared to 44 preferring attributional over relational similarities ($Z = 8.96, p < .001$).

The significantly greater support for the MAX model over the MIN model suggests that attributional similarities are more important for similarity judgments when the two objects being compared are superficially similar (and conversely for relational similarities). In Fig. 2, for example, subjects showed a tendency to think *A* and *D* were the most similar to *T*. *D*, the relationally similar choice, was picked when the overall attributional similarity was low. *A*, the attributional choice, was picked when the overall attributional similarity was higher. The result is particularly striking because of the within-subject nature of the choices. MAX was only supported when the *same subject* that chose *D* over *B* as being more

similar to *T* picked *A* over *C*. Therefore, simple subject bias cannot explain the results; in fact, MAX is only supported when a subject responds once relationally and once attributionally for a given group of pictures.

Whether an attributional similarity is preferred over a relational similarity depends on the other attributes of the object. Furthermore, relations seem to be stored and/or processed differently from attributes. A fully predictive model for similarity judgments cannot treat relations and attributes as undifferentiated.

EXPERIMENT 2

Experiment 1's finding that an increase in attributional similarity increases the weight of other attributes relative to other relations is consistent with the hypothesis that attributes and relations are considered separately for similarity judgments. However, since a majority of the responses did not show MAX effects, one could assume featural independence at the level of similarity computation with nonindependence being introduced by the judgment function mapping similarities onto choices. One plausible possibility in this regard is the notion that choices map onto the ratio of similarities rather than the differences. With a ratio rule, the contribution of a particular component to performance will decrease as the total similarity increases, as in a Weber fraction. Referring again to Fig. 2, *B* and *D* have less total similarity to *T* than do *A* and *C*. Given that subjects as a group gave more weight to relational matches than to attributional matches, with a ratio rule this greater weight should be more manifest on *B* vs *D* choices than on *A* vs *C* choices. Consequently, people would choose *D* over *B* more often than they would choose *C* over *A*, because the relational advantage of *C* and *D* is more apparent when total similarity is low.

Experiment 2 tests the ratio rule interpretation by varying both attributional and relational similarities. Figure 3 shows a sample stimulus set. The important change from the first experiment is that as one moves either down or to the right one both takes away an attributional similarity and adds a relational similarity. Therefore, *A* and *C* do not necessarily possess more total similarity to the target than *B* and *D*. Indeed if relations continue to receive more weight than attributes, then *B* and *D* will be seen as more similar to the target than *A* and *C*. In this circumstance, the ratio rule interpretation of nonindependence will predict MIN violations. The MAX hypothesis, of course, continues to predict that subjects will tend to choose *A* and *D*, the choices with the greatest attributional and relational similarities, respectively.

Method

Materials. Each stimuli set consisted of 12 pages of questions of the form "*T* is like *Y* or




FIG. 3. Sample stimuli and experimentally observed similarity ratings from Experiment 2.

Z'. *T*, *Y*, and *Z* were all geometrical pictures containing between two and five elements, with each element varying on size, shading, shape, and position. The pictures came from one of eight groups, each group containing five pictures. Each group had a single *T* picture for each comparison. The other four members of the groups were of the four types in Fig. 3. An attributional similarity was defined as an agreement between two pictures in the size, shape, or shading of any of their elements. The first member, *A*, contained two attributional similarities in common with *T* that *D* did not, and one more attributional similarity than either *B* or *C*. *A* relational similarity is an agreement in one of the following relations: Same-shape as, Same-shading-as, Taller-than, Larger-than, and Numerically-more-than. *D* contained two relational similarities in common with *T* that *A* did not, and one more relational similarity than either *B* or *C*. *B* and *C* each had one relational and one attributional match with *T*. The eight picture groups are given in Appendix B.

In six separate "T IS LIKE Y OR Z" statements, each picture *A*, *B*, *C*, and *D* was compared with every other picture (*A* or *B*, *A* or *C*, *A* or *D*, *B* or *C*, *B* or *D*, *C* or *D*). Thus, each of the eight groups generated six different comparisons. Which picture filled the *Y* slot and which filled the *Z* slot was randomized, and counterbalanced in a second test form. Six comparisons were placed on a page, with the restriction that no two comparisons from the same group be on the same page. The stimuli were designed on the Macintosh computer and printed on a laser printer.

Procedure. Participants were told to circle one of the two figures of each set, whichever picture was more similar to the left figure. The task took approximately 20 min to complete.

Subjects. Subjects were 23 University of Illinois undergraduate students taking part in the experiment for partial fulfillment of a course requirement. Groups of 12 and 11 subjects were tested simultaneously.

Results and Discussion

As in Experiment 1, the ordinal measure was a conservative measure of

the number of independence assumption violations. MAX is supported if $SIM(D,T) - SIM(B,T) > SIM(C,T) - SIM(A,T)$ and MIN is supported if the first term is less than the second term. For a given group of $A - D$ pictures plus target T , 16 ordinal results are possible:

If $SIM(D,T) > SIM(B,T)$ and $SIM(C,T) < SIM(A,T)$ then MAX is supported

IF $D < B$ and $C > A$ then MIN is supported

IF $D > B$ and $C > A$ and $D > C$ and $A > B$ then MAX

IF $D > B$ and $C > A$ and $D < C$ and $A < B$ then MIN

IF $D < B$ and $C < A$ and $D > C$ and $A > B$ then MAX

IF $D < B$ and $C < A$ and $D < C$ and $A < B$ then MIN

Otherwise: Independence.

MAX and MIN are each supported by three different pairwise orderings. For MAX or MIN to be supported, a positive-negative interaction effect must occur. For example, if $D > B$ (shorthand for "If the similarity of D to T is greater than the similarity of B to T ") and $C < A$, then clearly $(D - B) > (C - A)$ because $(D - B)$ is positive and $(C - A)$ is negative. Alternatively, if both $D > B$ and $C > A$, then MAX can still be supported if $D > C$ and $A > B$ because this pattern requires that $(D + A) > (C + B)$, which can be rewritten as $(D - B) > (C - A)$. The only time that MAX or MIN is supported is when adding the same feature to two stimuli causes a reversal in which stimulus is deemed more similar to the target.

Collapsing over all 23 subjects, and over all eight pictures, MAX was supported 45 times, MIN was supported 17 times. The other 122 response patterns were consistent with independence. Because of the conservative bias favoring placement in the independent group, this group cannot be compared with the other two. The number of MAX confirmations is significantly greater than the number of MIN confirmations (two-tailed $Z = 3.55$, $p < .001$). Given the choice of picking B or D as more similar to T , the relational picture (D) is picked over twice as often as the attributional picture (126 to 58). This advantage virtually disappears when subjects choose between A and C —the relational picture is picked only 8 more times than is the attributional picture (96 to 88).

Pictures with relational similarities are generally judged to be more similar to T than pictures with attributional similarities (Relational pictures were picked 608 times, attributional pictures 312 times). Therefore, the ratio rule interpretation of nonindependence must predict that MIN violations will dominate MAX violations, contrary to our data. Another explanation which makes an improper MIN prediction is an overall similarity bias which claims, "When scenes are similar in any way, then attributes are more highly weighted." While this hypothesis is supported by Experiment 1, it wrongly predicts that attributional choices will be

more common among *B* and *D* comparisons than *A* and *C* comparisons, because *B* and *D* are, overall, more similar to *T*.

The tendency for subjects to respond according to MAX is also apparent when the analysis is broken down by subject. Seventeen subjects made more MAX than MIN judgments, three subjects made more MIN judgments than MAX judgments, and three subjects made an equal number of MAX and MIN judgments. The difference between the number of subjects tending to be MAX biased and the number of subjects tending to be MIN biased is significant ($Z = 3.15, p < .005$). Only one subject failed to make at least one MAX or MIN response.

One objection to Experiment 2 is that the same feature that is added from *A* to *B* may not be, in some sense, the same feature that is added from *C* to *D*. For example, in Fig. 3, going from *C* to *D* effectively adds the feature "Shape AND Shading symmetry" whereas in going from *A* to *B* only "Shape symmetry" is added. Although this interpretation does not seem plausible for the stimuli in Fig. 3, four judges isolated four out of the eight picture groups as having a possible interpretation where the *C* - *D* feature change was not the same as the *A* - *B* feature change. In all four of these groups, when both relations were present, higher-order relations such as "The figure with fewer lines is bigger" were the additional features. The four groups that were judged as possibly manifesting higher-order relations are asterisked in Appendix B. When these four asterisked groups were eliminated from analysis, MAX was supported 22 times and MIN was supported 8 times ($Z = 2.55, p < .05$). Consequently, even after the suspicious items were eliminated, the difference between *D*'s and *B*'s similarities to *T* is greater than the difference between *C*'s and *A*'s similarities to *T*.

This asymmetry in similarity judgments is consistent with the MAX model. The weight of a shared relation in determining similarity is higher if it occurs in pictures that are already relationally similar. Likewise, the weight of a shared attribute in determining similarity is higher if it occurs in pictures that are already attributionally similar. In Fig. 3, for the subjects that responded in concordance with MAX, there is *no context-independent* statement concerning the relative weights of the relation "SAME-SHADING" and the attribute "BLACK." If the pictures are attributionally alike (*T* and *A*), then the attribute "BLACK" is more important than the relation "SAME-SHADING", and *A*, not *C*, will be judged as more similar to *T*. If the pictures are relationally alike (*T* and *B*), then the relation "SAME-SHADING" is more important than the attribute "BLACK", and *D*, not *B*, will be judged as more similar to *T*.

EXPERIMENT 3

The second experiment undermines the ratio rule interpretation of non-

independence. So far, however, we have not demonstrated ordinal violations of independence that cannot be attributed to the function mapping similarity onto judgments. In Experiments 3 and 4, we attempt to reveal stronger violations of independence and correspondingly stronger support for the MAX hypothesis. The forced choice paradigm necessitates presenting choices in the context of each other. Because context sensitivity seems to be the rule in similarity judgments (Birnbaum, 1982; Parducci, 1965; Tversky, 1977), there may be a general concern that the MAX effects found in Experiments 1 and 2 are a product, in some way, of context effects. Consequently, we shifted from a forced choice task to a rating task for the next two experiments.

In Experiment 3, these difficulties are overcome by presenting subjects with a similarity rating task of the form: "On a scale from 1 to 9, how similar are X and Y ?" The same abstract design is used as in Fig. 2, and support for MAX can come in two forms. The weak form of support occurs if the similarity ratings follow the $(D - B) > (C - A)$ rule (where D is shorthand for $SIM(D, T)$). For example, if the similarity ratings were

| | |
|-------------|-------------|
| $SIM(A, T)$ | $SIM(B, T)$ |
| 5.6 | 6.0 |
| $SIM(C, T)$ | $SIM(D, T)$ |
| 5.8 | 6.8 |

then $(D - B) = .8$ and $(C - A) = .2$. This pattern supports MAX in that going from attributional similarity to relational similarity ($A \rightarrow C$ and $B \rightarrow D$) has a larger increase on similarity when attributional similarity is low. Alternatively put, the attributional choice is almost as good as the relational choice (compare $SIM(A, T)$ with $SIM(C, T)$ when attributional similarity is high, but when attributional similarity is low, the relational choice is much more similar to T).

The reason that this is only weak support for MAX is that it falls short of the ordinal effect that was required for MAX support in Experiments 1 and 2. The above results could be explained without invoking MAX; for example, it could be that subject's similarity judgments are more sensitive in the 5-6 range than they are in the 6-7 range. That is, the rating scale may not satisfy the assumptions of an interval scale. So, the difference between 5.6 and 5.8 may be psychologically equal to the difference between 6.0 and 6.8.

The strong form of support that a rating task could yield is exemplified by the following ratings:

| | |
|-------------|-------------|
| $SIM(A, T)$ | $SIM(B, T)$ |
| 6.0 | 4.0 |

$$\begin{array}{c} \text{SIM}(C,T) \\ 5.0 \end{array}$$

$$\begin{array}{c} \text{SIM}(D,T) \\ 4.5 \end{array}$$

Here, not only is $(D - B) > (C - A)$ but $(D - B)$ is positive and $(C - A)$ is negative. This pattern represents ordinal support for MAX. Making the same change from attributional similarity to relational similarity has *opposite* effects on similarity ratings, depending on the level of attributional similarity between the compared items. As long as the similarity rating scale is assumed to be monotonically increasing, this pattern of results is strong support for the MAX effect.

Method

Materials. There were six groups of pictures fitting the abstract design of Fig. 2. Figure 4 shows an example set of stimuli (set 2). Each group consisted of a target picture *T*, and four choice pictures, *A*, *B*, *C*, and *D*. Each picture was made up of between two and five geometric forms. The forms varied on shape and shading. The relations used were Same-shape and Same-shading. Attributional similarity was changed from low to high by changing the shape or shading of one or more of the geometric forms in such a way that it did not disrupt the other attributional or relational features of the picture. All pictures were displayed on Macintosh SE computers.

Procedure. Subjects were instructed to rate how similar the two picture sets displayed on the screen were. They were told to press the "1" key if the pictures were not very similar, press the "9" key if the pictures were highly similar, and press the other number keys for




FIG. 4. Sample stimuli from Experiment 3. Feature independence is violated if subjects choose *D* as more similar to *T* than *B*, and choose *A* as more similar to *T* than *C*.

the various intermediary degrees of similarity. On each trial, a target picture and a picture from its group (A, B, C, or D) were displayed on the screen, side by side.

Participants were first shown 25 pictures for 3 s each, to give them some idea of the stimulus variability. During this familiarization period subjects were not prompted for ratings. Next, the subjects began the rating task. The order for displaying the target/test picture pairs was randomized. The left/right position of the target and the test picture was randomized, and counterbalanced on a second display to the subject. Consequently, each target picture was displayed eight times—two times with each of the four test pictures within its group.

Subjects. Subjects were 29 University of Illinois undergraduate students taking part in the experiment for partial fulfillment of a course requirement. Groups of 1-4 subjects were tested simultaneously.

Results and Discussion

The mean similarity ratings for each of the four picture types, broken down by picture group, are shown in Table 1. MAX predicts that $SIM(D,T) - SIM(B,T)$ is greater than $SIM(C,T) - SIM(A,T)$. In other words, MAX predicts that the expression $(D - B) - (C - A)$ in Table 1 will be positive. MIN, on the other hand, predicts that $(D - B) - (C - A)$ will be negative, and an assumption of independence predicts that $(D - B) - (C - A)$ will not be biased either positively or negatively. For five out of six picture sets, $(D - B) - (C - A)$ is greater than zero. The mean similarity ratings, collapsing over the six picture groups, show the same trends, yielding a value of .6 for $(D - B) - (C - A)$; this value is significantly greater than zero ($df = 5$, $t = 3.448$, $p < .02$), confirming the interval prediction made by MAX. Two of the individual picture sets (sets 2 and 3) have values for $(D - B) - (C - A)$ that are significantly greater than zero. (Set 1 and set 6 show effects of comparable magnitude but the greater variability in these conditions prevented the difference from achieving statistical reliability.)

TABLE 1
Similarity Ratings with Respect to Target T in Experiment 3

| | Picture set | | | | | | Mean |
|---------------------|-------------|-------|-------|-------|-------|-------|------|
| | Set 1 | Set 2 | Set 3 | Set 4 | Set 5 | Set 6 | |
| A | 5.8 | 6.4 | 5.5 | 5.9 | 4.3 | 6.7 | 5.8 |
| B | 3.4 | 3.1 | 4.4 | 4.0 | 2.1 | 6.0 | 3.8 |
| C | 5.2 | 5.8 | 4.9 | 6.2 | 4.4 | 6.8 | 5.6 |
| D | 3.6 | 3.6 | 4.5 | 4.8 | 2.1 | 6.8 | 4.2 |
| $(D - B) - (C - A)$ | .8 | 1.1 | .7 | .5 | -.1 | .7 | .6 |
| No. MAX Ss | 10 | 14 | 11 | 11 | 12 | 8 | 11 |
| No. MIN Ss | 5 | 3 | 5 | 6 | 9 | 3 | 5.2 |

The data also show an ordinal crossover between relational and attributional similarities. An ordinal crossover occurs if *D* is more similar to *T* than is *B*, and *A* is more similar to *T* than is *C*. This result was found for three of the picture sets (1, 2, and 3). The observed similarity ratings for picture set 2 are shown in Fig. 4. Moreover, the mean similarity ratings for all picture sets display an ordinal crossover. The *A* picture's mean similarity to *T* (5.8) is greater (nonsignificantly) than *C*'s similarity to *T* (5.6), and *D*'s similarity to *T* (4.2) is significantly ($df = 5, t = 2.7, p < .05$) greater than *B*'s (3.8). The fact that a crossover of similarity ratings was obtained is suggestive of the MAX effect that was found earlier.

Further evidence of ordinal support for MAX comes from examination of individual subjects' data. For each picture set, each subject was labeled as a MAX subject if he/she consistently (i.e., on both trials) rated *A* as more similar to *T* than *C*, and *D* as more similar to *T* than *B*. They were labeled MIN if the opposite pattern applied. If the subject's two ratings of a pair did not depict the same ordinal pattern with respect to its comparison rating, then the response was not counted as either MAX or MIN. Because some subjects were therefore left unlabeled, the number of MAX subjects plus the number of MIN subjects in Table 1 is less than the total number of subjects. Using this analysis, no scaling assumptions are required except that the higher rating given, the more similar the pair of objects are. The number of MAX subjects was greater than the number of MIN subjects for each picture group, and the difference is significant (exact binomial probability = .0156). The total number of subjects who gave more MAX responses than MIN responses is 22 out of the 29 subjects ($Z = 7.76, p < .01$).

Again, we find that whether an attribute or a relation match has more weight for similarity judgments depends on the number of other attribute matches that the objects share. When attribute match *A* is pitted against relation match *R*, *A* increases similarity more than *R* only when there are other attributes in common between the objects. When attributional similarity is low, *R* increases similarity more than *A*. Experiment 3 eliminates the possibility that this effect depends on the forced choice paradigm. In addition, the evidence in favor of MAX satisfied the strong, ordinal criteria, as well as the weaker, interval criteria.

EXPERIMENT 4

Experiment 4 was designed to answer two objections to the last experiments. The first objection is that the MAX effect may be due to the effect of shared *identical* elements between picture sets. Applying the objection to Fig. 4, the claim would be that in going from *T* - *C* to *T* - *A* we add the shared feature match of CIRCLE, but we *also* add the match IDENTICAL OBJECT. In other words, *T* and *A* have the identical form in the

second position—a white circle. In going from $T - D$ to $T - B$ we add the shared CIRCLE match without adding the IDENTICAL OBJECT match. Consequently, if identity along all features has significance above and beyond the sum of the individual feature matches, then going from $T - C$ to $T - A$ does not represent the same change as going from $T - D$ to $T - B$ does. Smith (1989) has recently suggested that adults do give particular weight to identity. This objection cannot explain all of the observed results (it does not explain the MAX effect found in stimuli like Fig. 2), but if identity is having a large effect on similarity judgments we may have overestimated our MAX effects.

The second objection is that all of our stimuli have consisted of detached parts which form a unit only in that they constitute the same "scene." The geometric shapes are detached perceptually and cognitively. They are perceptually detached in that one part of a scene does not touch any other. They are cognitively detached in that the scene does not form a Gestalt (Wertheimer, 1958); that is, a coherent, unified whole.

The first objection is met by designing stimuli that never give an identity advantage to the A picture as compared to the C picture in the $[T, A, B, C, D]$ picture group. An example of one such set is shown in Fig. 5. In this set, going from $T - C$ to $T - A$ does not result in an identical part match because none of parts of A are identical to any of the elements of T —they are all larger. In other sets, A 's parts are all nonidentical to T 's parts due to their shading or shape.




FIG. 5. Sample stimuli and observed similarity ratings (for picture set 3) from Experiment 4. None of the shapes in pictures A , B , C , or D are identical to any of the shapes in T .

The second objection is met by designing sets in which the parts combine to form a face which is perceived as a coherent, whole object. Such a set is shown in Fig. 6. These sets are compared with sets which are identical except that the large face-forming circle is eliminated, creating scenes like the scenes used in the early experiment.

Method

Materials. Fourteen picture sets were used that subscribed to the design of Fig. 5. For all of these sets, the parts of the A scene were never exactly identical to any of the parts of the T scene. In six of the sets the parts were larger or smaller, in five of the sets the parts were differently shaded, and in five of the sets the parts had different shapes (these numbers exceed 14 because in some cases more than one feature was changed to insure nonidentity). Four of the sets formed to create the appearance of a face with two eyes, two ears, a nose and a mouth. Four other sets were identical to the face sets except for the elimination of the profile-defining circle. For all sets, the relations used were Same-color and Same-shape. The picture parts were defined by three attributes: shading, size, and shape. For the

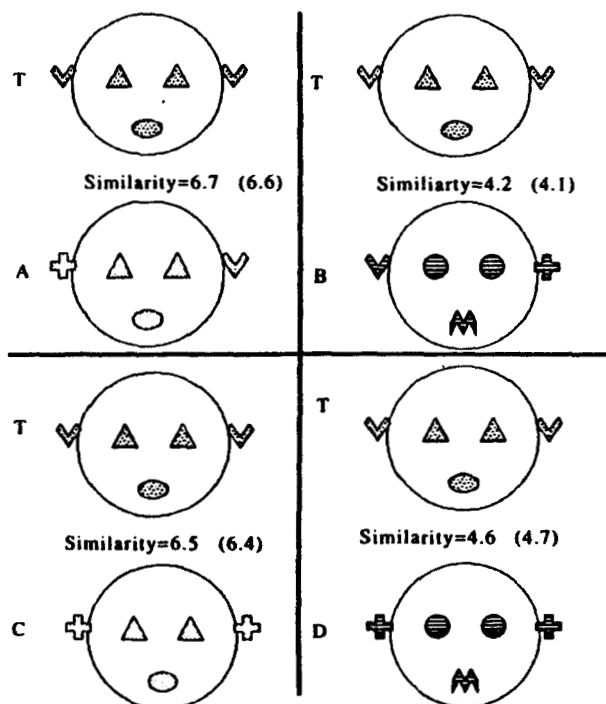


FIG. 6. Sample stimuli and observed similarity ratings from Experiment 4. The numbers outside of the parentheses show the observed similarities for picture set 4. The numbers in parentheses show the ratings for set 4 when the face-defining circle is removed from A, B, C, D, and T.

face sets, the relations involved either two eyes, two ears, an eye and an ear, or an eye and the mouth. All pictures were displayed on Macintosh SE computers.

Procedure. The same procedure to collect ratings data was used as was described in Experiment 3.

Subjects. Subjects were 42 University of Illinois undergraduate students taking part in the experiment for partial fulfillment of a course requirement. Groups of 3-6 participants were tested simultaneously.

Results and Discussion

Table 2 shows the mean similarity ratings for each picture set, and the number of MAX and MIN subjects. For 11 out of 14 pictures set, $(D - B) - (C - A)$ is positive, the interval criteria for MAX effects. In only one picture set is the term negative. The overall mean difference between $(D - B)$ and $(C - A)$ is .51, a value significantly greater than the difference of zero predicted by feature independence ($df = 13$, $t = 4.05$, $p < .01$).

Out of the 11 picture sets, there are five ordinal crossovers between relational and attributional responding in the MAX direction while there are none in the MIN direction. In these tests, *D* gets a higher rating than *B* while *A* gets a higher rating than *C*, with respect to similarity to *T*. Four

TABLE 2
Similarity Ratings with Respect to Target T in Experiment 4

| | Picture set = Geometric scenes | | | | | | Mean |
|---------------------|--------------------------------|-------|-------|-------|-------|-------|------|
| | Set 1 | Set 2 | Set 3 | Set 4 | Set 5 | Set 6 | |
| A | 5.4 | 7.2 | 6.0 | 3.9 | 7.2 | 5.9 | 5.9 |
| B | 3.2 | 3.1 | 3.0 | 2.0 | 3.2 | 3.1 | 2.9 |
| C | 6.0 | 7.3 | 5.6 | 4.6 | 7.4 | 6.3 | 6.2 |
| D | 4.3 | 3.5 | 4.1 | 2.7 | 3.4 | 3.2 | 3.5 |
| $(D - B) - (C - A)$ | .5 | .3 | 1.5 | 0 | 0 | -.3 | .46 |
| No. MAX responses | 42 | 32 | 50 | 39 | 29 | 32 | 37 |
| No. MIN responses | 28 | 28 | 13 | 34 | 24 | 35 | 27 |

| | Picture set = Face sets and controls | | | | | | | | Mean |
|---------------------|--------------------------------------|--------|-------|--------|-------|--------|-------|--------|------|
| | Set 1 | Set 1C | Set 2 | Set 2C | Set 3 | Set 3C | Set 4 | Set 4C | |
| A | 6.6 | 6.4 | 6.2 | 6.0 | 7.0 | 6.6 | 6.7 | 6.6 | 6.5 |
| B | 3.3 | 3.0 | 3.6 | 3.1 | 3.0 | 2.6 | 4.2 | 4.1 | 3.4 |
| C | 6.9 | 6.6 | 5.4 | 4.8 | 6.2 | 6.3 | 6.5 | 6.4 | 6.1 |
| D | 3.9 | 3.7 | 3.2 | 3.0 | 3.2 | 2.8 | 4.6 | 4.7 | 3.6 |
| $(D - B) - (C - A)$ | .3 | .5 | .4 | 1.1 | 1.0 | .5 | .6 | .8 | .7 |
| No. MAX Ss | 43 | 40 | 39 | 44 | 49 | 44 | 46 | 47 | 44 |
| No. MIN Ss | 27 | 29 | 30 | 27 | 24 | 27 | 20 | 21 | 26 |

Note. Control sets are marked with the suffix "C" and do not have circles representing the outer facial profile.

of these crossovers occurred in the face stimuli—two in the full-face condition, and two in the control, faceless sets.

In Experiment 3, each subject's response pattern to a picture set was labeled as a MAX response if *both* of his/her judgments rated *A* as more similar to *T* than *C*, and *D* as more similar to *T* than *B*. In order to achieve slightly greater sensitivity in Experiment 4, the subjects' two judgments were analyzed separately. Thus, "No. MAX" responses in Table 2 is the same as "No. MAX" subjects in Table 1, except that each individual response is paired (the *first T - A* judgment is paired with the first *T - C* judgment, and the *second T - A* judgment is paired with the second *T - C* judgment) and considered separately. MIN labels were assigned to picture sets in which *C* was rated as more similar to *T* than *A*, and *B* was rated as more similar to *T* than *D*. The number of MAX subject responses was greater than the number of MIN subject responses for 13 out of 14 picture sets, and this difference is significant (binomial two-tailed probability = .0136). We can also label each subject as either a MIN or a MAX subject depending on whether more of their fourteen set responses are ordinarily MIN or MAX. If we do so, 34 subjects are labeled MAX, seven subjects are labeled MIN ($Z = 17.78$, $p < .01$), and one subject had exactly identical numbers of MAX and MIN responses. This is a particularly important finding in that it dismisses the possibility that MAX effects are caused by just a few strongly MAX-biased subjects. The bias toward increased relational weighting during times of low attributional similarity seems to be displayed by the large majority of people, on the large majority of stimuli.

In addition to finding strong, ordinal evidence of MAX effects, Experiment 4 answers the two objections previously raised. There is no evidence that the stimuli that were presented as unified faces were less likely to cause MAX effects than either the faceless control sets or sets which were used in prior experiments. If we analyze the data by subject or by picture, there is no significant difference between the MAX effects for the face sets and the control sets. MAX effects do not seem to rely on presenting stimuli which are cognitively or perceptually detached. In addition, even when the influence of identical forms is removed, we still obtain MAX effects. The finding that attributes count more than relations when objects are already attributionally similar seems to be a fairly robust phenomenon.

GENERAL DISCUSSION

Experiments 1, 2, 3, and 4 support the claim that the relative importance of relations and attributes in object similarity judgments depends on the types of similarities already present in the objects. If attributional similarities predominate, then attributional matches will have their

weights increased, relative to relational matches. If attributional similarities are decreased (Experiments 1, 3, and 4), or relational similarities increased (Experiment 2), then relational matches will gain in importance relative to attributional matches.

This pattern of results suggests (1) adding similarities between two objects has different effects on perceived similarity, depending on whether relations or attributes are added, (2) there is context sensitivity in how shared features are integrated to form an overall similarity judgment—how one feature affects similarity depends on the nature of the other features present in the object, (3) selective attention mechanisms can function on the basis of relation/attribute similarity type, and (4) not only *can* mechanisms function to selectively attend to relations or attributes, these mechanisms *naturally* function even when no instructions are given to bias their operation.

A general response bias cannot explain our results. MAX violations only arise when subjects make relational choices on some occasions and attributional choices on other occasions. A subject who consistently responds relationally, for example, is counted as acting in accordance with the independence assumption.

MODELING THE OBSERVED NONINDEPENDENCE

Tversky's contrast model coupled with the assumption of feature additivity does not anticipate the feature nonindependence that was observed. Feature additivity claims that $f(X \cap Y)$ is expressible as the sum of the measures of all the features that belong to both X and Y . The demonstration of the additive contrast model's failure to explain the interaction effect found is given in Appendix C. In the demonstration, it is granted that features may be basic shapes and shadings, and also relations such as "Top and bottom shapes are same." The version of the additive contrast model cannot predict, taking Fig. 2 as an example, that picture D is more similar to T than is B , if C is to be more similar to T than is A . Although feature additivity was assumed by Tversky in his application of the contrast model to categorization, clustering, and prototypicality (Tversky, 1977), the contrast model is in no way constrained to adopt feature additivity.

It is important to bear in mind that nothing in the present data directly undermines the general form of Tversky's contrast model because the model only embraces feature additivity as a special case. Our results do suggest an important processing principle that so far has not been incorporated into similarity judgment models. The observed pattern of feature nonindependence could be accommodated by Tversky's contrast model and the MDS model if a processing assumption were included which selectively weighted components depending on the other components of

the scene. In terms of Tversky's model, the claim would have to be that $f(X \cap Y)$ is not the sum of the features shared by X and Y when different types of features (e.g., attributional versus relational) are involved.

The exact functional definition of $f(X \cap Y)$ is underdetermined by the present experiments. The empirical constraints merely require similarities from one type (relational or attributional) to increase the weight of similarities of their own type relative to the weights of the other type. One promising possibility along these lines is that similarity increases nonlinearly as a function of the number of matching features *within a single pool*, and these within-pool similarities are simply added together to determine overall similarity. Several researchers have already suggested nonlinear functions relating similarity with number of shared features such that adding a shared feature increases similarity more if the number of shared features is already high [Hintzman's (1986) cubic function, Medin & Shaffer's (1978) multiplicative rule, and Shepard's (1974) exponential law]. If we accept this assumption of nonlinearity, and also assume that these functions operate only on the similarities within single pools, then MAX effects are predicted. For example, using Medin and Shaffer's multiplicative rule for within-pool influences and an additive between-pool function, we could get

$$f(X \cap Y) = \left(\prod_{k=1}^n A_k + \prod_{k=1}^m R_k \right),$$

where n is the number of shared attributes, m is the number of shared relations, and all A and R terms are greater than 1 and refer to the amount that type of similarity is increased by when the feature is shared.

There are also less direct ways of producing MAX effects within the framework of the contrast model. For example, one could assume that attributes and relations correspond to two distinct dimensions on which values can be ordered (for a set of three relations, one relation can be said to lie between the other two). MAX effects might be obtained if *mismatches* within dimensions are subadditive, whereas interdimensional mismatches are additive. Tversky and Gati (1982) argue for intradimensional subadditivity by demonstrating violations of the corner inequality; they find that the corner path $(ap, ar) + (ar, cr)$ is often psychologically shorter than the diagonal path $(ap, bq) + (bq, cr)$. Burns, Shepp, McDonough, and Weiner-Ehrlich (1978) and Shepard (1964) also present evidence that two differences lower similarity less if they fall on the same dimension than if they fall on different dimensions. MAX effects might also be obtained if *matches* within a dimension are superadditive. One difficulty with these explanations of MAX effects is their assumption that

relations and attributes are dimensions. While it makes intuitive sense to order objects on size and brightness dimensions, it is not at all obvious that it even makes sense to order object values according to their "relationality." Is LARGER-THAN more or less relational than SAME-SHAPE, and does SAME-SIZE lie between them? Ordering all possible attributes along a single dimension is even more problematic, given the enormous variety of attributes and the lack of a criteria for determining which of two attributes is "more attributional." The MAX effects presented here may be viewed as extending the previous violations of the corner inequality to domains where the features are organized around categories, not around dimensions.

A final method of handling MAX effects circumvents the problems associated with positing relational and attributional dimensions. The method requires the decomposition of attributes and relations into microfeatures, coupled with the assumption that any two relations or any two attributes tend to share more features than an attribute and a relation do. Suppose that in determining the similarity of two stimuli (*A* and *B*) two matching relations (*R1* and *R2*) are detected. The key idea is that in computing the number of matching features there are contributions from both corresponding and noncorresponding microfeatures. That is, most of the common features will derive from *R1* in *A* matching *R1* in *B*, and *R2* in *A* matching *R2* in *B*. However, there will also be microfeatures from *R1* in *A* that match some of the microfeatures from *R2* in *B*, because *R1* and *R2* are somewhat similar. This assumption will produce superadditivity within sets of relations or sets of attributes and, therefore, MAX effects. Because any two *arbitrarily* chosen features from within a dimension tend to augment each other more than two arbitrarily chosen features from between dimensions, the assumption has to be made that any two features within a dimension share microfeatures. Even relations as different as "TALLER-THAN" and "SAME-SHADING" would have to have microfeatures in common—microfeatures referring to the fact that both features are *relational*. As such, this account still assumes that features must be tagged as to whether they are relational or attributional.

In fact, all of these variations have in common, at one level or another, distinctions between attributes and relations. Any similarity model will have to divide features into (at least) relational and attributional types, and the relations between features within a similarity type will not be the same as the relations between features of different similarity types. It is useful to think of attributes and relations as existing in two different pools, with the mutual influence of features within a pool exceeding the influence of features across pools.

Difficulties arise when trying to fit our data with a conventional multidimensional scaling model. In multidimensional scaling (MDS), stimuli

are represented as points in a space, and the similarity between two points decreases monotonically as the distance between them increases. One general MDS model (Carroll & Wish, 1974; Nosofsky, 1987) computes the dissimilarity of stimuli I and J via

$$d_{ij} = \left[\sum_{k=1}^N W_k |x_{ik} - x_{jk}|^r \right]^{1/r},$$

where $0 \leq W_k \leq 1$ and $\sum W_k = 1$. W_k is the weight of dimension k , X_{ik} is the value of stimulus i on dimension k , and N is the number of relevant dimensions. By varying the value of r , Euclidean ($r = 2$), city block ($r = 1$), and other metrics may be modeled.

The demonstration in Appendix D shows that if this MDS model is to predict how A could be more similar than C to T in Fig. 2, then it has to make assumptions about the dimension weights and values that make it impossible for it to explain how D could be judged as more similar than B to T . This incompatibility persists as long as $r > 0$. Furthermore, the MDS analysis cannot be saved by ignoring relational dimensions. In general, MDS models are constrained such that if C is further away from (more different than) T than A , then displacing C and A the same amount along a dimension cannot reverse this distance relation.

A more surprising result is that, even if the MDS model were to include our assumption that relations and attributes are pooled separately for similarity judgments, it still would not be able to predict MAX effects. In fact, an MDS model that assumes a "relations" dimension (where all the shared *relations* between two scenes are pooled) and an "attributes" dimension actually predicts MIN effects. The demonstration of this failure, formally presented in Appendix E, is intuitively borne out by Fig. 7. A pooled MDS analysis of Fig. 3 would predict that B and C are closer to T than are A and D . This is because when similarities of the same type are pooled together in the same dimension, their effects will be additive; their effects will be *less* than additive (unless $r = 1$) if the similarities belong to different dimensions. B and C , each possessing differences that belong to two dimensions, will be placed closer to T , and will accordingly be judged more similar to T . This, of course, is exactly the opposite pattern of that which was empirically obtained—there was a tendency for subjects to consider B and C the *least* similar of the choices.

COMPARISON WITH EXISTING THEORIES

The results support Gentner's (Gentner, 1983; Gentner & Landers, 1985) distinction between responses based on "true analogy" and those based on "mere-appearance" similarity. In true analogy mode, only re-




FIG. 7. The general multidimensional scaling (MDS) model, supplemented by the assumption that relational similarities are pooled separately from attributional similarities, predicts a pattern of results that is the opposite of that predicted by MAX.

lational predicates are mapped from a base to a target. In mere-appearance mode, object attributes alone are mapped. In "literal similarity" mode both relations and attributes are counted when performing the mapping. People making judgments about the goodness of an analogy were found to selectively attend to the objects' relational properties. In the current experiments, subjects preferred to base their similarity judgments on either just relations or just attributes; they did not choose as similar objects that had one relational and one attributional similarity. The "pooling" notion inherent in the MAX model supports Gentner's observation that people in "analogy mode" tend to select only the relational commonalities between systems. Furthermore, our experiments suggest that subjects "put all their eggs into one basket" such that high similarity in one similarity class is preferred over medium similarities in two classes. This corresponds to a tendency to make similarity responses using either "true analogy" or "mere appearance" criteria. Analogies are only possible if the attributional chaff is separated from the relational grain, and the current experiments suggest that this sorting occurs naturally. Once the two similarity types are separately pooled, then similarity judgments can selectively weight the evidence procured from each pool.

Medin, Goldstone, and Gentner (1990) argue for the psychological distinction between relations and attributes on independent grounds.

Their subjects made forced choice similarity judgments in which one choice was attributionally similar to a target and the other choice was relationally similar to the target. They found that the relationally similar choice tended to be chosen as both more similar to the target and as more different from the target. Similarity and difference judgments appear to treat relations and attributes differently; to predict the differential effect a feature will have on both types of judgment, it is necessary to know whether the feature is a relation or an attribute.

Our work is similar to Palmer's in a number of respects. Palmer (1978) argued against the "segment-independence assumption" in light of evidence that line segments enter into larger units—figure goodness and naturalness are functions of properties which depend on more than one segment. Likewise, we argue against feature independence on the grounds that an attribute or relation does not have a fixed salience or weight; it's weight depends on the other features that the objects possess. Palmer (1977) has also stressed the need for features at multiple levels of organization. At least three levels (whole figure, multisegment parts, and individual line segments) must be postulated to explain how figures are divided, rated for goodness, and synthesized. Similarly, our model assumes that attributes and relations between attributes are simultaneously extracted, and that both can have major influences on similarity judgments.

CAN THE ATTRIBUTES/RELATIONS DISTINCTION BE REDUCED TO A DIFFERENT DISTINCTION?

One objection to drawing the distinction between attributes and relations is that it is confounded with other distinctions. The literature in psychology makes a number of distinctions regarding features. It might be thought that the relations vs attributes dichotomy could be better characterized by one of these distinctions. We will consider two such distinctions: holistic vs analytic (Brooks, 1978; Ward & Scott, 1987) and global vs local (Navon, 1977). We argue that while these distinctions are related to the relations/attributes distinction, they neither singly nor jointly exhaust its meaning.

Kemler (1983) distinguished between holistic and analytic modes of processing on the grounds that the former involves the overall, unitary aspects of an object, while the latter apprehends the separate dimensions/properties of an object. She gave evidence that as people mature they go from a relatively heavy reliance on the holistic mode to more reliance on the analytic mode. For example, young children who see a piece of paper cut in half often state that the paper has changed its color also. Kemler argued that the children perceived the difference between the old piece of piece of paper and the transformed piece, but were unable to analyze the

difference to determine the specific dimension which was responsible for the change.

While it might seem tempting to equate such holistic properties with relational features because they both involve the synthesis of elements, there is developmental evidence against this suggestion. Gentner (1988b) presents evidence that children give attributional interpretations to comparisons to which adults give relational interpretations. Given "A cloud is like a sponge," the five year old typically explains, "They both are round and fluffy" whereas the adult typically responds, "They both hold water and give it back later." Thus, sensitivity to the holistic properties of objects, far from constituting a relational bias, is completely compatible with a general attribute bias; children show exactly this pattern. Relative to children, adults are simultaneously more analytic (Kemler, 1983) and more sensitive to relations (Gentner, 1988b).

Further evidence against equating holistic with relational processes comes from adult reaction times. Adults under time pressure base judgments more on holistic properties than do adults without time pressure (Smith & Kemler-Nelson, 1984). On the other hand, our preliminary results show that time pressure increases the proportion of attributional responses. In a pilot experiment, subjects were given either 5 s or unlimited time to make a forced choice between an attributionally and a relationally matching picture. For the subjects required to respond within 5 s, the attributionally similar picture was picked as more similar 42% of the time. This percentage is significantly higher than the 35% selection proportion we find for subjects who have no time limit. In other words, the more time subjects have to respond, the more likely they will be to prefer the relationally similar choice over the attributional choice. Thus, it seems plausible that holistic properties are different from relations; holistic properties are quickly (developmentally and chronologically) apprehended, whereas sensitivity to relations develops along a slower time course.

As to the second distinction, global aspects of an object are said to involve the whole object and can be processed without breaking down the object into parts (Navon, 1977; Palmer, 1980; Pomerantz, 1983). Navon relates the globality of a feature with its place in a construction hierarchy—if feature *X* is built up from (in part) feature *Y*, then *X* is more global than *Y*. Navon gives an example of a crescent moon next to a star; the most global feature of this scene is the 'blob' obtained by enclosing both of the celestial bodies. This global blob can be processed before the scene is resolved into its component objects. In fact, Navon's hypothesis is that global aspects of an object are processed before the more local aspects are. He presented letters that were composed of smaller letters (see Fig. 6 for an example). Even though the local and global features

were equally perceptible, identification of the local letters was influenced by the identity of the global letter. The reciprocal effect of local letters on global letters was not found.

The argument could be made that relations should be thought of as simply one kind of global feature. Global features, like relations, involve parts, and global features and relations are often more than "just the sum of their parts" (Wertheimer, 1958). However, many relational features, including the ones we use here, are *not* global in the sense of involving the whole object. In Fig. 3, the relation "Same-color (second object, fourth object)" only involves two out of five of *C*'s elements. Furthermore, the most global feature of *C*, according to Navon's blob analysis, would be the general horizontal linearity of *C*. The "thin horizontal blob" global feature can be (and is, if Navon's global precedence hypothesis is true) processed before *C* is even analyzed into separate objects. The Same-color (second object, fourth object) relation, however, requires that object 2 and object 4 be distinguished in order to see the relation. While global objects can be perceived before their components are (An "F" composed of many small "R"s can be perceived before the "R"s are perceived), it would be surprising if the "SAME-SHAPE" relation were perceived before the actual shapes were perceived. "Surprising" is not the same thing as "impossible." It *might* be the case that humans see that "XOX" has a SAME-SHAPE relation without seeing that it has a "X" shape, if we had a primitive symmetry or identity detector. However, this would require a mechanism that is quite different from the mechanism by which horizontal linearity is perceived before individual components are. The simple global mechanism that Navon had in mind was that we first apprehended gross features and only later focus in on detail. This mechanism is not adequate to explain our relational/attributional distinction, since relations are neither always gross nor perceptually independent of their parts.

A distinction which is not dissociable from the attributes/relations distinction in our stimuli is that of abstract/concrete. Though the distinction is difficult to define, we have the intuition that the "Darker-than" relation is more abstract than the attribute "gray." Our results, if framed in terms of abstractness, suggest conditions under which abstract properties are more likely to be noticed than concrete/superficial properties. Namely, if two scenes (1) are abstractly similar or (2) are not superficially similar, then further abstract commonalities will be relatively more likely to be the basis for judgment. If one wishes to highlight a particular abstract similarity between two events as opposed to a concrete similarity, then one would be wise to downplay their superficial correspondences and stress their other abstract correspondences. Our reason for preferring the relation/attribute distinction over the abstract/concrete distinction is that its

terms are explicitly defined, it more precisely maps unto the design of our stimuli, and it permits a more precise statement of our findings. There are many types of abstraction ("mammal" is more abstract than "dog"; " $D = R * T$ " is more abstract than " $120 \text{ miles} = 60 \text{ miles/h} * 2 \text{ h}$ ") that our results do not address. The relation/attribute distinction seems to be the most accurate and conservative way to characterize the differences between our stimuli.

If it is true that the relation/attribute distinction cannot be reduced to the two dichotomies discussed above, these dichotomies can still be investigated using the same procedures detailed in Experiments 1-4. For example, to investigate the validity of the global/local dichotomy, the paradigm illustrated in Fig. 8 can be used. The *A* picture would contain two large-scale similarities to *T*, *B* and *C* would contain one large-scale and one small-scale similarity to *T*, and *D* would contain two small-scale similarities to *T*. If *D* is judged to be more similar to *T* than *B*, and *A* is more similar to *T* than is *C*, then local and global similarities are selectively attended in accordance with MAX. Such a result would suggest another nonindependence of features, and would establish the two similarity types as functionally distinct. Preliminary experimentation has yielded 27 MAX responses and only four MIN responses; this systematic bias supports Kimchi and Palmer's (1985) postulation of a psychologically motivated distinction between global and local features. Work in progress by the authors suggests that the technique used in this paper can be enlisted to support such diverse distinctions as shapes vs shades, curvy




FIG. 8. The logic of Experiment 3 can be applied to other psychological distinctions. Here, a "two-pools" account of global and local similarities is suggested if *D* and *A* are judged to be the most similar to *T*.

shapes vs straight shapes, and left vs right. In short, although we have used the MAX hypothesis to support the distinction between attributional and relational features, the MAX principle appears to have a broader potential.

CONCLUSIONS

Our results show that shared relations have greater effect on similarity judgments when they are in the context of other relational similarities (and likewise, for attributes). Our interpretation of this phenomena has been that similarities of the same type (relational type or attributional type) are pooled together, and similarities that are pooled together tend to mutually increase the weight of one another in similarity judgments. An equally plausible alternative to the separate pooling of similarity types is for all similarities to be intermixed, and selectively accessed by *processes* that are differentially specified for relations or attributes. Finding empirical evidence for deciding between separate structures and separate processes will most likely prove to be an extremely difficult task (Palmer, 1978). The commonality between the two interpretations is that they both assert that structures/processes of the same type are mutually reinforcing whereas structures/processes of different types are less reinforcing or even inhibitory.

In claiming that the relations and attributes of scenes are separately pooled, and treated differently for similarity judgments, we are not claiming that these are the only two pools to which features can belong. There may be many overlapping pools that are constructed on the basis of multitudinous distinctions. However, it is interesting that relations and attributes are two such pools; attributes as different as "fifth object is striped" and "first object is large" mutually influence/enhance each other, as do relations as different as "first two objects have same shading" and "last object is darker than first." People seem to be sensitive to more than the whether a feature's content concerns shading, size or shape information—they are also sensitive to the structural role of the feature.

APPENDIX A



T

Max violations: 5
Min violations: 1



A



B



C



D



T

Max violations: 3
Min violations: 1



A



B



C



D



T

Max violations: 1
Min violations: 0



A



B



C



D



T

Max violations: 2
Min violations: 1



A



B



C



D



T

Max violations: 1

Min violations: 0



A



B



C



D



T

Max violations: 5

Min violations: 0



A




B



C




D




T

Max violations: 5


Min violations: 0



A



B



C



D



T

Max violations: 4

Min violations: 1



A



B



C



D

APPENDIX B

Note. Stimuli that were excluded from the second analysis are asterisked.



T

Max violations: 5

Min violations: 1



A



B



C



D



T

Max violations: 8

Min violations: 1



A



B



C



D



T

Max violations: 4

Min violations: 2



A



B



C



D



T

Max violations: 9

Min violations: 0



A



B



C



D



T

Max violations: 4
Min violations: 2



A




B



C




D



T


Max violations: 9
Min violations: 0




A



B



C



D



T

Max violations: 4
Min violations: 5



A



B



C



D



T

Max violations: 2
Min violations: 6



A



B



C



D

APPENDIX C

Modeling with the Additive Contrast Model

We can, without loss of generality, analyze Fig. 2:

Attributes: $A1$ = "X" on top

$A2$ = star in middle

$A3$ = "X" on bottom

$A4$ = circle in middle

$A5$ = square on top

$A6$ = square on bottom

$A7$ = all attributes shared by T , A , B , C , and D .

Relations: $R1$ = top and bottom shapes are same

$R2$ = all relations shared by T , A , B , C , and D

$R3$ = top and bottom shapes are different.

Given these conventions, the figures sets can be defined as

$$T = A1 + A2 + A3 + A7 + R1 + R2$$

$$A = A1 + A2 + A6 + A7 + R2 + R3$$

$$B = A1 + A4 + A6 + A7 + R2 + R3$$

$$C = A2 + A5 + A6 + A7 + R1 + R2$$

$$D = A4 + A5 + A6 + A7 + R1 + R2.$$

The contrast model is expressible as

$$\text{SIM}(Q, R) = X * f(Q \cap R) - Y * f(Q - R) - Z * f(R - Q).$$

The choice-target similarities according to the feature additivity version of this model can be written as

$$\text{SIM}(A, T) = X(A1 + A2 + A7 + R2) - Y(A6 + R3) - Z(A3 + R1)$$

$$\text{SIM}(B, T) = X(A1 + A7 + R2) - Y(A6 + A4) - Z(A2 + A3 + R1)$$

$$\text{SIM}(C, T) = X(A2 + A7 + R1 + R2) - Y(A5 + A6) + R3) - Z(A1 + A3)$$

$$\text{SIM}(D, T) = X(A7 + R1 + R2) - Y(A4 + A5 + A6) - Z(A1 + A2 + A3).$$

The results we obtained suggest that $\text{SIM}(D, T) > \text{SIM}(B, T)$ while $\text{SIM}(C, T) < \text{SIM}(A, T)$. These results require $\text{SIM}(D, T) - \text{SIM}(B, T)$ to be positive while $\text{SIM}(C, T) - \text{SIM}(A, T)$ is negative, whereas the above expressions dictate

$$\begin{aligned} \text{SIM}(D, T) - \text{SIM}(B, T) &= X(R1 - A1) - Y(A5) - Z(A1 - R1) \\ &= \text{SIM}(C, T) - \text{SIM}(A, T); \end{aligned}$$

That is, the model predicts that $\text{SIM}(D, T) - \text{SIM}(B, T)$ must equal $\text{SIM}(C, T) - \text{SIM}(A, T)$.

The demonstration results are unchanged if any of the A or R terms are

zero, if the features are treated as binary (circle in middle vs star in middle), or if additional feature terms are added. The equality also holds where the values of X , Y , and Z are allowed to be different for attributes and relations.

APPENDIX D

Modeling with a Generalized MDS Model

To fit our results into the MDS model, it is first necessary to list the relevant dimensions. Again using Fig. 2 as a general example, four dimensions were extracted:

Dimension 1 = Top element, taking values of "X" or Square

Dimension 2 = Middle element, taking values of Star or Circle

Dimension 3 = Bottom element, taking values of "X" or Square

Dimension 4 = Same-shape relation, taking values of present or absent.

The weights for these dimensions are W_1 , W_2 , W_3 , and W_4 , respectively. If two figures have the same value on a dimension k (both T and A have the value "X" on dimension 1) then the value of $|X_{ik} - X_{jk}|$ is zero. If the two figures have different values on a dimension, then the value for $|X_{ik} - X_{jk}|$ will be referred to as V_k . Using these conventions we can assign dissimilarity (distance) expressions between each of the four figures and T :

$$\begin{aligned} d_{A,T} &= [W_1*0 + W_2*0 + W_3*V_3' + W_4*V_4']^{1/r} \\ d_{B,T} &= [W_1*0 + W_2*V_2' + W_3*V_3' + W_4*V_4']^{1/r} \\ d_{C,T} &= [W_1*V_1' + W_2*0 + W_3*V_3' + W_4*0]^{1/r} \\ d_{D,T} &= [W_1*V_1' + W_2*V_2' + W_3*V_3' + W_4*0]^{1/r}. \end{aligned}$$

The systematic finding we reported was that subjects chose D over B , and A over C , as being most similar to T . Because the MDS model assumes similarity decreases as distance increases, our data constrains the MDS model so that

$$d(D,T) < d(B,T) \text{ AND } d(A,T) < d(C,T).$$

However, it is impossible for the MDS model to satisfy both of these constraints simultaneously. Under the MDS model, if $d(D,T) < d(B,T)$ then it must be the case that

$$\begin{aligned} W_1V_1' + W_2V_2' + W_3V_3' &< W_2V_2' + W_3V_3' + W_4V_4'. \\ \text{So, } W_1V_1' &< W_4V_4'. \end{aligned}$$

But, if $d(A,T) < d(C,T)$ then

$$\begin{aligned} W_3V_3' + W_4V_4' &< W_1V_1' + W_3V_3'. \\ \text{So, } W_4V_4' &< W_1V_1', \end{aligned}$$

which contradicts the previous requirement.

APPENDIX E

Modeling with MDS, Assuming Pooled Features

Considering the picture group in Fig. 3, the four manipulated differences divide into two groups:

Attributional differences—

A1 = The first element: Circle or Square

A2 = The second element: Black or Striped.

Relational differences—

R1 = The first and last elements: Same-shape or Different-shape

R2 = The second and fourth elements: Same-shading or Different-shading.

Each *A* and *R* term is constrained to have a positive weight. If the two types of differences are pooled together, then a simplified MDS formula results:

$$d_{ij} = [|X_{iA} - X_{jA}|^r + |X_{iR} - X_{jR}|^r]^{1/r}.$$

So, if the differences are pooled, we only have to include two dimensions, *A* for the attributes dimensions, and *R* for the relations dimensions. We can neglect the weighting terms (*W*s) without affecting the results, since weighting biases can be achieved by manipulating the values of A1, A2, R1, and R2, as long as these values do not change depending on the picture in which they occur. The specific similarities within a dimension are simply added together. Given this metric, we can again assign differences between *T* and each of the four choices in Fig. 4.

$$d_{A,T} = [0 + (R1 + R2)^r]^{1/r} = R1 + R2$$

$$d_{B,T} = [A1^r + R2^r]^{1/r}$$

$$d_{C,T} = [A2^r + R1^r]^{1/r}$$

$$d_{D,T} = [(A1 + A2)^r + 0]^{1/r} = A1 + A2.$$

It is not strictly correct to say that there are no attributional differences between *T* and *A*. The analysis only requires that there are no attributional differences between *T* and *A*, *relative* to the other three choices. Our MAX results require $d_{D,T} - d_{B,T} < d_{C,T} - d_{A,T}$. Swapping sides yields the inequality $d_{D,T} + d_{A,T} < d_{C,T} + d_{B,T}$. This inequality requires that

$$(A1 + A2) + (R1 + R2) < [(A2)^r + (R1)^r]^{1/r} + [(A1)^r + (R2)^r]^{1/r}$$

which can be rewritten as

$$(A2 + R1) + (A1 + R2) < [(A2)^r + (R1)^r]^{1/r} + [(A1)^r + (R2)^r]^{1/r}.$$

However, this statement is never true while $r \geq 1$. In fact, if $r > 1$, then exactly the opposite inequality is predicted by MDS. If anything other than a city-block metric is used, then MDS (plus the pooling assumption) requires that

$$(A2 + R1) + (A1 + R2) > [(A2)^r + (R1)^r]^{1/r} + [(A1)^r + (R2)^r]^{1/r}$$

since

$$(A2 + R1) > [(A2)^r + (R1)^r]^{1/r} \text{ and } (A1 + R2) > [(A1)^r + (R2)^r]^{1/r}.$$

Consequently, MDS predicts MIN, *not* MAX, for all values of $r > 1$.

REFERENCES

- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 33-53.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Birnbaum, M. H. (1982). Controversies in psychological measurement. In B. Wegener (Ed.), *Social attitudes and psychological measurement* (pp. 401-485). Hillsdale, NJ: Erlbaum.
- Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Burns, B., Shepp, B. E., McDonough, D., & Wiener-Ehrlich, W. K. (1978). The relation between stimulus and analyzability and perceived dimensional structure. *The Psychology of Learning and Motivation*, 12, 77-115.
- Carroll, J. D., & Wish, M. (1974). Models and methods for three-way multidimensional scaling. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 2, pp. 57-105). San Francisco: Freeman.
- Garner, W. R. (1974). *The processing of information and structure*. New York: Wiley.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. *Cognitive Psychology*, 16, 341-370.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1988a). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity, analogy, and thought*. New York: Cambridge Univ. Press.
- Gentner, D. (1988b). Metaphor as structure mapping: The relational shift. *Child Development*, 59, 47-59.
- Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man, and Cybernetics*. Tucson, AZ.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-429.
- Hock, H. S., Tromley, C., & Polmann, L. (1988). Perceptual units in the acquisition of visual categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 75-84.

- Kemler, D. G. (1983). Holistic and analytic modes in perceptual and cognitive development. In T. Tighe & B. Shepp (Eds.), *Perception, cognition, and development*. Hillsdale, NJ: Erlbaum.
- Kimchi, R., & Palmer, S. E. (1985). Separability and integrality of global and local level of hierarchical patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 673-688.
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1990). Similarity involving attributes and relations: Judgments of similarity and difference are not inverses. *Psychological Science*, 1, 64-69.
- Medin, D. L., & Shaffer, M. M. (1978). A context theory of classification learning. *Psychological Review*, 85, 207-238.
- Monahan, J. S., & Lockhead, G. R. (1977). Identification of integral stimuli. *Journal of Experimental Psychology: General*, 106, 94-110.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, 9, 353-383.
- Nickerson, R. S. (1972). Binary Classification reaction time: A review of some studies of human information-processing capabilities. *Psychonomic Monograph Supplements*, 4 (No. 6), 275-317.
- Nosofsky, R. M. (1987). Attention and learning process in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 87-108.
- Ortony, A. (1979). Beyond literal similarity. *Psychological Review*, 86, 161-180.
- Palmer, S. E. (1975). Visual perception and world knowledge. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition*. San Francisco: Freeman.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9, 441-474.
- Palmer, S. E. (1978). Structural aspects of visual similarity. *Memory & Cognition*, 6, 91-97.
- Palmer, S. E. (1980). What makes triangles point: Local and global effects in configurations of ambiguous triangles. *Cognitive Psychology*, 12, 285-305.
- Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review*, 72, 407-418.
- Pomerantz, J. R. (1983). Global and local precedence: Selective attention in form and motion perception. *Journal of Experimental Psychology: General*, 112, 516-540.
- Pomerantz, J. R., & Garner, W. R. (1973). Stimulus configuration in selective attention tasks. *Perception and Psychophysics*, 14, 565-569.
- Raaijmakers, J. G., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93-134.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382-407.
- Ross, B. H. (1984). Reminders and their effects in learning a cognitive skill. *Cognitive Psychology*, 16, 371-416.
- Shephard, R. N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology*, 1, 54-87.
- Shepard, R. N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39, 373-421.
- Smith, J. D., & Kemler-Nelson, D. G. (1984). Overall similarity in adult's classification: The child in all of us. *Journal of Experimental Psychology: General*, 113, 137-159.
- Smith, L. B. (1989). A model of perceptual classification in children and adults. *Psychological Review*, 96, 125-144.

- Smith, L. B., & Kemler, D. G. (1978). Levels of experienced dimensionality in children and adults. *Cognitive Psychology*, 10, 502-532.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.
- Tversky, A., & Gati, I. (1982). Similarity, separability, and the triangle inequality. *Psychological Review*, 89, 123-154.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113, 169-192.
- Ward, T. B., & Scott, J. (1987). Analytic and holistic modes of learning family-resemblance concepts. *Memory & Cognition*, 15, 42-54.
- Wertheimer, M. (1958). Principles of perceptual organization. In D. C. Beardslee & M. Wertheimer (Eds.), *Readings in perception*. Princeton, NJ: Van Nostrand.
- Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of association for computing machinery*, 23.
- (Accepted December 14, 1989)