Medin, D. L. & Goldstone, R. L. (1995). The predicates of similarity. In C. Cacciari (Ed.), Similarity (pp. 83-110). Milan: Bompiani.

The Predicates of Similarity

Douglas L. Medin and Robert L. Goldstone
University of Michigan Indiana University

Abstract. Traditional approaches to similarity in cognitive psychology treat it as a two-place predicate. Nelson Goodman, however, has pointed out that for similarity to be meaningful it must be a three-place predicate of the form "A is similar to B in respects to C," we argue that similarity statements must be expanded to include at least four other predicates which are directly or indirectly linked to the comparison process. We conclude that the joint consideration of structure and processing principles gives similarity explanatory power.

Address correspondence to:
Douglas L. Medin
Psychology Department
University of Michigan
330 Packard Road
Ann Arbor, MI 48104-2994
email: Doug,Medin@um.cc.umich.edu

The Predicates of Similarity

Douglas L. Medin and Robert L. Goldstone
University of Michigan Indiana University

I. Introduction

A. Similarity as a two-place predicate

Although similarity seems like a fairly straight forward notion, it lies at the interface of a number of deep and important issues. It is straight forward in that intuition readily suggests that similarity of two things is based on shared parts or properties. In the words of William James (1890, p. 529), "To abstract the ground of either difference or likeness (where it is not ultimate) demands an analysis of the given objects into their parts". The general idea is that similarity is based on either identity or similarity underlying parts and properties. Identity versus similarity of parts or constituent dimensions corresponds to two of the major psychological theories of similarity. Featural models of similarity describe the similarity of two entities as an increasing function of matching properties (or features) and a decreasing function of mismatching properties (e.g. Tversky, 1977). Multi-dimensional scaling models represent the similarity of things in terms of the proximity of their values on underlying dimensions. Identity of value of some attribute corresponds to the special case of zero distance on that dimension.

The reader may have noticed, however, that the terms part, property, attribute, value, feature, and dimension have been left undefined. Problems begin to arise when we try to make these notions more precise and they seem to become unmanageable when we try to specify just what features or properties some entity "has."

Both the featural and multi-dimensional scaling approaches to similarity treat it as a two-place predicate. That is, the similarity of A to B, S(a,b), is assumed to be a meaningful relation. There are critics, however, who argue that this notion of similarity is incoherent because it does not sufficiently constrain the ways in which things can be similar or different. These critics claim that similarity is inherently a three-place predicate. That is, it is only meaningful to say that A is similar to B with respect to C.

The remainder of this paper is organized as follows. We first lay out the argument in favor of viewing similarity as a three-place predicate and derive the corresponding implication for the use of similarity as an explanatory construct in psychology. (To anticipate, the news is not good.) We then argue that similarity should be viewed as a multi-place predicate involving not only structure but also processing principles. It is these processing principles, we claim, that allow similarity to be a meaningful, coherent construct with explanatory power. We illustrate our claims with a variety of evidence and then end the paper by drawing out some implications of our framework.

B. Similarity as a three-place predicate.

The philosopher Nelson Goodman (1972) calls similarity a "pretender" and an "imposter". He argues that the similarity of A to B is an ill-defined, meaningless notion unless one can say "in what respects" A is similar to B. In Goodman's words, "we must search for the appropriate replacement in each case; and 'is similar' to function as little more than a blank to be filled (p. 445)." In short similarity is conceived as a three-place predicate.

Given that meaning is conveyed by specific respects, "similarity" seems to disappear when it is analyzed closely. That is to say, the explanatory work is being done by whatever it is that determines respects. In this sense similarity may be more like a dependent variable than an independent variable. For example, we may see things as similar because we classify them together rather than vice versa.

Have similarity theories successfully fixed respects? Despite the various successes of similarity research, we believe that the answer must be "no." Multi-dimensional scaling models are often used in an exploratory manner to discover underlying dimensions. Once obtained, scaling solutions are often quite successful at describing other patterns of performance such as recognition memory or categorization (e.g. Nosofsky, 1989). Multidimensional scaling does not address the question of how dimensions are determined. In a closely analogous way featural models of similarity describe how matching and mismatching features are integrated to form judgments but they beg the question of how features are fixed. In Tversky's words, "when faced with a particular task (e.g. identification or similarity assessment) we extract and compile from our data base a limited set of relevant features on the basis of which we perform the required task. Thus the representation of an object as a collection of features is viewed as a product of a prior process of extraction and compilation (Tversky, 1977, pp. 329-330)." Nelson Goodman might argue that unless this prior process is well-understood, similarity will have severely limited utility.

Although we agree that "extraction and compilation" are critical processes, it is not the case that the only issue of interest in similarity research is filling in the respects clause. Even when the constituent properties are well-known, there are important questions concerning how respects are integrated to form judgments, as the work on similarity, induction, and categorization attests (e.g. Krumhansl, 1978; Gati and Tversky, 1984, Osherson, Smith, Wilkie, Lopez, and Shafir, 1991; Nosofsky, 1992). When Goodman writes, "to say that two things are similar in having a specified property in common is to say nothing more than they have the property in common" (p. 439) we believe that he has oversimplified the role that similarity typically plays. In most uses of similarity, more than one property is relevant, broad properties such as "visual appearance" are instrumental, and exact identity on a property is not required. Furthermore, as we shall see, similarity comparisons may be directional and involve both assertions and inference processes. Fixing respects remains a fundamental issue, but there's more to similarity than having a property in common.

Another reaction to the respects issue is to point to the empirical success of similarity. For example, one could argue that the perceptual system determines respects in a fairly rigid manner. That is, "similarity" is just shorthand for "similarity as constrained by the perceptual system". On this view, Goodman's critique is technically correct but of little practical value.

But similarity critics are unlikely to be persuaded by the empirical success of similarity. Experimenters may restrict their attention to just those experimental contexts where their conjectures about the relevant aspects of similarity are likely to be correct. As such, the apparent stability of similarity may be illusory. Moreover, large individual differences might be expected to exist. People from different cultures, with different levels of expertise, or with different knowledge bases will likely disagree on their similarity judgments, despite their basically equivalent perceptual systems.

We also believe that to draw a sharp distinction between hard-wired perceptual processes and more conceptual-driven processes is a mistake. Even when judging the similarity of purely visual displays, people show sensitivity to abstract relations such as "this display has elements of increasing size, just as this display has elements of increasing darkness (Goldstone, Gentner, and Medin, 1989)." Nor do measures of perceptual similarity completely converge on a single construct. Similarity relations may vary with processing time (Goldstone, 1991) and Palmer (1978) reports that the influence of common structural relations depends on factors such as whether the stimuli are presented simultaneously or sequentially. It appears that similarity is tied up with processes in a way not captured by viewing similarity as fixed (see also Smith and Heise, 1991, for arguments for the view that even perceptual similarity is dynamic rather than static).

Our overall impression, then, is that Goodman's argument that similarity must be a three-place predicate has serious implications for psychology. It seems that concerted efforts are needed to rescue the concept of similarity from empty circularity. We believe that similarity is often quite flexible and that Goodman's critique is potentially very damaging. But we will also argue that the reason similarity sometimes seems chaotic is that too little attention has been paid to the processing side of similarity. Our thesis is that "respects" are systematically fixed by similarity comparison processes and fixed sufficiently that similarity is far from being an empty explanatory construct. From a processing perspective similarity should be viewed as a multi-place predicate.

C. Similarity as a multi-place predicate.

We believe that similarity statements are quite complex. Specifically, the statement that A is similar to B in respect C is an incomplete, misleading analysis of similarity. At a minimum, similarity statements need to be expanded to include "according to comparison process D, relative to some standard E, mapped onto judgments by some function F, from perspective G". There may be other predicates; the claim is that similarity judgments are at least this complex. In the remainder of this paper, we outline a framework for comparison and then describe evidence bearing on these predicates of similarity.

II. A Framework for Comparison.

A. Similarity as comparison.

Similarity is a type of comparison that often arises in normal discourse. Glucksberg and Keysar (1990) suggest that similarity comparisons and similes (A is like B) show a number of common properties. For example, both similes and similarity statements appear to be directional. To say that surgeons are like butchers means something different than to say that butchers are like surgeons (the former criticizes surgeons and the latter compliments butchers). Tversky (1977) has shown that similarity judgments may be asymmetric. For example, Tversky found that people rate the similarity of Red China to North Korea to be less than the similarity of North Korea to Red China. According Tversky's contrast model, asymmetries arise because the distinctive features of A in an (A, B) comparison receive more weight than the distinctive features of B. Later on we consider the possibility that the common or shared features change as a function of the direction of the comparison (again, see Glucksberg and Keysar, 1990).

Both similarity statements and similes may also be anomalous. For example, we believe anomaly is involved for the statements "robins are similar to birds" and "robins are similar to questions". The former is anomalous because robins are not similar to birds, they are birds. (We'll ignore the fact that cognitive psychologists often ask participants to make such judgments). The latter seems strange because the term, similar, seems to presuppose some amount or type of similarity, neither of which is apparent.

Observations about anomaly are important because they make the point that similarity judgments involve something more than a calculation. To the extent that similarity judgments include intuitions about anomaly, they must involve more than a weighting function on common and distinctive properties. At the very least, there seems to be a process that monitors the comparison for particular relations between compared items and for extreme values of similarity.

A third shared property between similarity statements and similes is that one expects them to be at least somewhat informative. To say that butchers are like surgeons is to assert that some salient property of surgeons. (e.g. skill in cutting) is also true of butchers (see Ortony, 1979). We believe that similarity comparisons may also involve assertions where properties of one entity become candidate properties of the other. Upon hearing that quaggas are similar to zebras, a person may correctly infer that quaggas have four legs. If so, then similarity is more than identity in certain respects.

In short, similarity seems to be a type of comparison with properties distinct from those associated with the view that similarity is only a determination of shared and distinctive respects. In addition to matching and mismatching properties, similarity may involve other processes such as directionality, anomaly, and implicit understanding about pragmatics and informativeness. These processes are crucial for understanding similarity statements.

B. A Framework for Comparison.

1 Property activation and comparison. Suppose you are asked to rate the similarity of North Korea to Red China. What information will enter into the comparison? It is very unlikely that you would access all of your knowledge about these complex entities. Furthermore, there is no guarantee that corresponding pieces of information will be activated. For example, you might think of Chinese rugs and be uncertain about whether or not rugs are also made in North Korea. Sjoberg (1972) argues that similarity judgments involve an active search process where one looks for ways in which the things under consideration are similar. Along these lines, we suggest that information activated for one concept will tend to be evaluated or act as a candidate inference for the other concept (see Clement and Gentner; 1991, for related ideas in the domain of analogy). For example, instead of treating "manufactures rugs" as true of Red China and not true of North Korea one might draw the plausible inference that rugs are, in fact, manufactured in North Korea. One might also expect that when similarity comparisons are stated directionally, as in "A is similar to B," properties of the B term will receive more attention as candidate properties of the A term than vice versa. This would follow from the principle that similarity statements are, at least in part, informative assertions.

2. Alignment and Structure.

There must be some process that brings properties into correspondence to determine whether they match or mismatch. We refer to processes by which entities are placed into correspondence as alignment. Our claim (see also Gentner, 1989; Goldstone and Medin, in press) is that alignment is nontrivial and needs to be considered explicitly. For example, if some hypothetical animal A has one tail and animal B has three tails how many potential matches and mismatches do we have? If only local feature matches are made, without any concern for global consistency, then each of B's three tails can be matched with A's tail. On the other hand, if a match between A's tail and one of B's tails is noticed, then this match may "lock" or bind A's tail to B's tail, making it ineligible for other matches.

We believe that the alignment process for similarity may be more or less identical to the alignment process for analogy (Falkenhainer, Forbus and Gentner 1990; Markman and Gentner, 1991). That is to say, global consistency and structure rather than only local matching are important for alignment. Again the implication is that similarity theorists need to be explicit about processes - in this particular case, processes that determine what correspondences will count toward similarity.

3 Summary.

We suggest that the entities being compared mutually constrain the "features" which are activated or inferred. Furthermore, the alignment process is dynamic and driven by global constraint satisfaction. Like Shannon (1988), we believe that similarity judgments involve creative, constructive processes. Similarity is not simply determined by features; the features themselves may be fixed by the similarity comparison. In the next section we bring out the implications of this general framework for similarity predicates.

III. Filling in Predicates

A. According to some comparison process D

We have been arguing that similarity assessment is a dynamic, context-specific process that includes the alignment of entities. That is, similarity may be less a computation over some feature space than a search and discovery process. As we shall see, a number of implications follow from this view.

1 Similarity comparisons may be directional. One way in which similarity comparisons might be informative is for properties of the base term to be asserted or emphasized with respect to the target term. When A is compared to B, properties of the base or standard (the B term) should be more likely to be activated than properties of the target (A term). This leads to the prediction that common or shared features associated with a comparison may differ as a function of the direction of the comparison, (see also Glucksberg and Keysar, 1990). In a recent experiment we found evidence consistent with this prediction.

The study was run in the following way. Two groups of participants rated the similarity of the 17 pairs of words shown in Table 1. One group rated the similarity of the X terms to the Y terms (e.g. England to the United States). They were instructed to "Consider Y. How similar is X to Y? What features do Xs have in common with Ys?" where X and Y were filled in with the words from Table 1. The other group rated the similarity of the Y terms to the X terms (United States to England) and listed features that Ys had in common with Xs.

Recall that our hypothesis is that common properties will differ as a function of the direction of comparison in being more closely associated with the base term. To evaluate this prediction, we transcribed all feature listings onto a computer spreadsheet, randomized their order and then gave them to two judges who were blind to the purpose of this experiment to score. Each judge was given the word pairs and asked to categorize each feature into one of the five following categories:

1 equally applicable to both concepts

2 biased toward the meaning of the left concept

3 biased toward the meaning of the right concept

4 true of the left concept but not the right one

5 true of the right concept but not the left one

Categories "2" and "3" were to be used if the feature were true of both words, but the feature seemed more appropriated for or applicable to one of the words. For example, for the comparison of Albert Einstein and Benjamin Franklin, "scientist" is true of both men but our two judges decided that this feature was more closely associated with Einstein. Categories "4" and "5" were reserved for clear-cut cases where a feature was true for only one of the terms. Thus, "famous American" was judged to be true of Benjamin Franklin but not Albert Einstein.

Summing over the 17 pairs the judges rated 2103 features and agreed on 87 percent of their category assignments. Judge 1 rated 432 features as biased toward the base, 377 as biased toward the target and 1294 features as neutral or equal. Judge 2 rated 413 features as biased towards the base, 363 toward the target, and 1327 features as neutral. Integrating across judges these differences are statistically reliable. Categories "4" and "5" were rarely used by the judges but again favor the base category (28 times a common features was judged to be true of the base but not the target compared with 21 times in the reverse direction). A clear example of the bias in feature listings is that participants were more likely to mention the property "found on farms" when asked about the similarity of dogs to cows than when asked about the similarity of cows to dogs.

Although asymmetries in similarity rating were not the focus of this study, they are summarized in Table 1. The average absolute difference in ratings (0.96 points on a 20 point scale) is roughly of the magnitude observed by Tversky and Gati (1978). Some of the comparisons fit very nicely with the view that asymmetrical comparison are like assertions. Consider the largest asymmetry which involved pencils and crayons. A salient property of pencils is that you can write with them and one can write with crayons almost as easily. A salient property of crayons is that you can color with them which seems considerably less applicable to standard lead pencils. Indeed the rated similarity of crayons to pencils (15.13) was considerably higher than the rated similarity of pencils to crayons (12.68).

Insert Table 1 about here

These results support the prediction that the common properties activated for a comparison depend on its direction. The current claim is somewhat different from Tversky's (1977) claim concerning asymmetries. According to Tversky's Contrast model, asymmetries arise because the distinctive features of the target (called "subject" by Tversky) are given greater weight than the distinctive features of the base ("referent"). Our results hinge on the asymmetrical processing of shared, not distinctive, features. Asymmetries due to common features are not predicted by the Contrast model if the similarity of A to B and the similarity of B to A involve exactly the same common features. Unlike Tversky's asymmetry, our results indicate that features associated with the base/referent are more likely to be considered in the comparison than features associated with the target/subject. Both Tversky's and our conjectures may be correct, if shared features are more salient if they belong to the base and distinctive features are more salient if they belong to target.

This observed asymmetry requires two things to be true: that properties activated of one concept be evaluated with respect to the other <u>and</u> that this activation be biased toward properties of the base concept. The magnitude of the second process represents a serious underestimate of the likelihood of the first process. Indeed, we suspect that comparison-dependent properties are fairly common. One striking example of context-dependent features comes from recent follow-up study by Cindy Aguilar, Evan Heit, and Douglas Medin. For a comparison involving <u>stomach</u> and <u>brains</u> one participant listed "stomachs hold food, brains hold food thoughts." Had participants been asked to list properties of the term <u>brains</u> in isolation, we think it extremely unlikely that anyone would have listed the property that brains can think about food. Again these observations fit nicely with the general argument that respects are importantly fixed by the comparison process itself.

2. Comparison-dependent features.

Although it is common to assume that entities "have" features, the framework we have been discussing is compatible with features being at least partially comparison-dependent. On our view, properties activated of one entity in a similarity comparison are evaluated as candidate properties of the other entity. If properties are sometimes ambiguously present then the features attributed to an entity may depend on what it is compared with. In one of our recent studies an object was placed into two different comparison contexts and we attempted to show that the properties inferred to be present in one context would contradict the inferred properties in the second context.

The logic of the experiment was quite simple. Participants were asked to compare a Stimulus B either to Stimulus A alone or to Stimulus C alone. They were asked to list common and distinctive features in each comparison. The stimuli were visual forms like those shown in Figure 1. The stimuli were constructed with the idea that a property inferred as present in B for an AB comparison might be incompatible with a property attributed to B in its corresponding BC comparison. For example, in the first triplet B might be said to have three "prongs" or "fingers" when compared to A but four prongs when compared to C. The A and C picture were chosen to: 1 be fairly similar to B, 2 clearly reflect the two opposite interpretations of B and 3 be approximately equally similar to B. Two different groups of participants were run and a given group saw only the AB or BC comparison (never both).

Insert Figure 1 about here

The results conformed to our predictions. Descriptions of B features that matched the stimulus it was paired with (e.g. saying that the top B stimulus in figure 1 had three prongs when compared with A) were classified as "comparison-consistent properties". Descriptions of B features that mismatched the value of the stimulus it was compared with were called comparison-inconsistent properties. Overall, the average number of comparison-consistent properties, 15.75, was substantially greater than the number of comparison-inconsistent properties, 4.86. Even when differences are listed, comparison-consistent properties appear more frequently than comparison-inconsistent properties. For example, people often listed "right-most prong is warped in B" as a difference between B and C in the top triad of figure 1. That is, they give B an interpretation that is consistent with C's property of having four prongs.

One incidental finding is the predominance of metaphors for similarities relative to differences. Metaphors were five times as likely to appear in similarity listings than in difference listings. For example, several people cite "both look like pincers" as a similarity between A and B in the second triad; very few people list "B looks like a pincer, but C does not" as a difference between B and C. Although we did not anticipate this result, we will offer a speculative interpretation. Differences may be derivative of similarities (see also Markman and Gentner, 1991). For example, one might note that a tennis ball and a baseball are both round, used in sports, bounce, and can be hit and then note as differences that baseballs tend to be hit further and that tennis balls bounce better. If differences tend to be derived from (aligned) similarities, then the properties listed as differences will tend to be more specific and less abstract than those listed for similarities. The relative prominence of metaphors for similarities would be just one instance of a more general tendency for similarities to be more abstract than differences.

The main conclusion from this study, however, is that the properties of an object depend on what it is being compared with. Our study used the extreme case where the properties attributed to the B stimulus when compared to the A stimuli contradicted the properties attributed to it when it was compared to the C stimulus. In short, features may be comparison-dependent.

3 Comparison requires alignment. Markman and Gentner (1990) have investigated the relation between similarity and alignment, using scenes such as the one depicted in figure 2. Note that the woman in one scene is perceptually similar to the woman in the other scene but that they play different roles. In the scene on the left the woman is the recipient whereas on the right she is the donor. The scenes involve crossmapping in the sense that the most natural perceptual correspondences conflict with the relational correspondences (Gentner and Toupin, 1986).

In Markman and Gentner's study the experimenter pointed to the cross-mapped object and asked the participant to select the object in the other scene that "went with" that object. One group of participants first made similarity judgments and then mapping judgments and another group only made mapping judgments. The group that first made similarity judgments was far more likely to map according to roles(e.g. from the woman in the first scene to the squirrel in the second scene) than was the control group. Other control conditions rule out familiarity with the materials as the reason for the differences. The Markman and Gentner study shows both that similarity judgments involve alignment and that alignment is driven by global (relational) structure rather than the best local matches (see also Goldstone and Medin, in press).

Insert Figure 2 about here

---

4 "Matches" and "mismatches" depend on alignment.

Suppose we are evaluating the similarity of Person 1 who has a striped shirt to Person 2 who has both a striped shirt and striped pants. If striped is a feature do we count just one match or two? If only the striped shirt to striped shirt match is allowed, what happens when Person 1 is compared with Person 3 who has a plain shirt and striped pants. Does the fact that both Persons 1 and 3 have stripes contribute to similarity and, if so, should the match count exactly the same as the striped shirt to striped shirt match?

Now consider a slightly more complicated situation. Imagine Person 4 with a black and white striped shirt and red and green checkered pants compared with Person 5 who has on a black and white checkered shirt and red striped pants. Putting aside questions about abstract features like "both are strangely dressed," we ask should we count both the red and green matches and the striped and checkered matches or are they mutually exclusive? For example, if we align checkered with checkered do we get a mismatch for colors?

The above examples raise the issue of simple local matches versus structure and global consistency. We already know from the Markman and Gentner (1990) study that alignment does not depend on local matches alone. Indeed, the alignment process for similarity may be more or less identical to the alignment process for analogy (Falkenhainer, Forbus, and Gentner, 1990; Falkenhainer, 1990; Markman and Gentner, 1991). In a separate paper (Goldstone and Medin, in press) we deal extensively with the alignment process for similarity comparisons and the relationship between alignment and similarity comparisons. In that paper we describe a computational model for alignment and experiments which bear on it. The upshot is that the effect of matching features on similarity judgments is dynamically determined and is driven by multiple constraint satisfaction and global consistency. For example, the extent to which a striped shirt on one person and striped pants on another person contributed to similarity critically depends on other matches present in the scene. The influence that a matching feature has on similarity depends on whether it belongs to aligned parts. In turn, parts are placed in alignment if they have many matching features in common, and if they are consistent with other developing correspondences.

A final example may serve to illustrate the complexity of the relationship between "matching" or "mismatching" and alignment. Suppose we are comparing drawings of people. Drawing A is missing a hand, Drawing B is not missing any body parts, Drawing C is missing a hand and a foot and Drawing D is missing a foot. The fact that A is missing a hand and B is intact decreases their similarity. One would also expect that the similarity of A and C is reduced by the fact that C is missing a foot that is not missing in Drawing A. So far these intuitions fit a frame-based comparison process where corresponding body parts are compared. But intuition (and some informal observations by Wisniewski and Medin, 1991) suggests that Drawings A and D are actually more similar than Drawings A and B (or B and D) despite the fact that A and D differ in the status of two appendages. Informally speaking, the level of comparison shifts and we say that both A and D are missing an appendage, thereby converting two frame-based mismatching values into a higher level match. This example shows that just what gets aligned with what and, therefore, whether something counts as a match or mismatch, is comparison-specific. A frame-based alignment process simply is not dynamic enough to do the work of similarity comparison.

5 Alignments (and therefore similarity comparisons) use structure.

The final point we wish to make about similarity comparisons is that they critically depend on structure--not only lower-level features or attributes but also relations among attributes and even higher order relations of relations (see Gentner, 1983, 1989). Although for some purposes it is useful to describe similarity in terms of independent feature lists, the fact that relations (e.g. A is above B) take either attributes or other relations as arguments makes the constituents of similarity inherently non-independent. A simple example serves to make this point. Let Drawing E be a triangle above a circle, Drawing F be a circle above a triangle and Drawing G be a star above a square. Both E and F have a triangle, a circle, and the relation *above*. E and F are not , however, identical and we may try to capture this difference by treating the arguments (circle, triangle) as bound to the relation above. That is, E is more similar to itself than it is F because E and F have a different ordering of arguments. On the other hand, E and F also resemble G in that all three contain the abstract relation of aboveness. To capture this similarity, we would want to treat the *above* relation as unbound. For *above* to be treated as unbound means that the order of the arguments within the relation is not relevant. One could describe the situation in terms of attributes, bound relations, and unbound relations but to do so would be to abandon the idea that features should be independent and non-decomposable (or at least not functions of features in the same descriptive language). We see no alternative to directly addressing structure, even if it entails giving up the idea of independent features.

6 Summary. The critical point of the predicate, "comparison process D" is the focus on process. One could summarize this section with the manifesto, "no representation without taxation." That is to say, one cannot develop models of similarity structure without worrying about processing. We have seen that the properties that enter into a comparison may depend on the direction of the comparison, that the properties an entity has may be comparison dependent, and that whether some property comparison entails a match or a mismatch depends on a dynamic alignment process driven by global constraint satisfaction. In short, one cannot begin to understand "A is similar to B in respects C" without considering the nature of the comparison process.

B. Relative to some standard E

Much of the work on similarity judgments implicitly assumes that they are made relative to the absolute standard of identity. That is , the highest possible degree of similarity is sameness or identity. We shall argue, however, that similarity judgments may often be relative to some implied contrast set rather than relative to identity. When someone says "Monica Seles and Steffi Graf are very similar" the contrast set or standard might consist of other professional female tennis players or even other successful, professional, female tennis players. The implied contrast set functions (in part) to determine which properties are relevant to the comparison. For example, if the contrast set for the Seles-Graf comparison is other professional, female, tennis players, then the fact that they are both women and both professional tennis players may contribute little to their judged similarity. Rather the key respects should tend to be properties that Graf and Seles share that other female professional tennis players do not possess (so-called diagnostic features--see Tversky, 1977).

Tversky's (1977) research on the extension effect can be viewed as showing the importance of implied contrast sets. He had people rate the similarity of various European countries to each other and the similarity of various North and South American countries to each other. Two groups of participants rated only one set whereas a third rated both sets. The group rating both sets gave significantly higher similarity ratings than the other two groups. One interpretation of this result is that the group that rated both sets included members of the alternative set as potential contrast set members (despite the fact that they were not asked to make any cross-set similarity judgments) which increased the weight given to properties that all members of a set shared.

We believe that changes in implied contrast sets can produce otherwise counter-intuitive effects on similarity judgments. Imagine evaluating the similarity of two scenes, one of which consists of a boy standing on a chair to reach a jar of cookies and the other of which consists of a boy standing on a stool to reach a box of cereal. The scenes would appear to be fairly similar. Now assume that we add a box of cereal to the first scene and a jar of cookies to the second scene. Although these additions increase the number of object correspondences, they serve to highlight the fact that the two boys have different goals (no longer reaching something inaccessible but rather reaching for one thing rather that another) and the sense of similarity decreases. In this case the contrast set is made explicit but more generally it may also be implicit.

Context and contrast sets. As a final example of standards or contrast sets we describe a recent study where our goal was to show that contrast sets influence "respects" and thereby, rated similarity. In the experiment to be described we attempted to vary contrast sets by presenting participants with either two-way or three-way comparisons. For the two-way comparisons, either A and B or A and C are presented in separate contexts. In the three-way comparisons the two pairs appeared in the same context. Abstractly, then, our design roughly corresponded to Tversky's experiment on the extension effect. We attempted to vary standards or contrast sets by varying relationships between A and B and A and C. Specifically, the comparisons involves antonymical, metaphorical, or categorical relations. The general idea was that whether or not one type of relation influenced judgments involving a different type of relation would depend on whether the comparison context was combined or separated.

The actual study involved a number of antonymic sets (e.g. sunrise, sunset; black, white) which were paired with categorical relations (sunrise, sunbeam; black, red) and metaphorical sets (e.g. Rolls Royce; Champagne; Monday, January; sun, light bulb) paired with categorical relations (e.g. Rolls Royce, Volkswagon; Monday, Sunday; sun, moon). For the three-way comparisons, the AB and AC appeared simultaneously on a computer screen in the form of an isosceles triangle with item A at the top. Participants were instructed to look at both comparisons before making their similarity ratings. For two-way comparisons the AB and AC pairs were not presented at the same time.

The results were generally as expected. Consider first the antonymic sets. When a pair like "black-white" is presented in isolation, the implied contrast set might be intermediate greys. In addition, in typical discourse contexts, antonymic comparisons presuppose similarities and assert differences. When a categorical pair (such as black-red) is added, the new comparison set serves to highlight the idea that black and white are both monochromatic. Hence, one might expect that similarity ratings for antonymic pairs to be greater in the combined than in the separated context. The mean similarity rating (on a 9-point scale) increased from 4.5 in the separate context to 5.5 in the combined context. The corresponding increase for categorical comparisons was from 5.0 to 5.3. The interaction of comparison type with context was statistically reliable. As a concrete example of the interaction, for the separate context, two-way comparisons the ratings were sunrise-sunset=4.6, sunrise-sunbeam=5.2; for the combined context, three-way comparison the ratings were sunrise-sunset=6.4, sunrise-sunbeam=5.7.

Contrast sets should work somewhat differently for the metaphorical set. The addition of categorical comparisons to the metaphorical comparison should not necessarily boost similarity ratings for two reasons: the (implied) contrast sets or standards are likely to be other possible metaphorical bases and the "respects" are likely to be fixed. That is, metaphorical comparison presuppose differences and assert similarities. The presence of metaphorical comparisons should, however, increase the rated similarity for categorical comparisons. And that is what obtains. The similarity ratings for metaphorical pairs averaged 4.4 and 4.5 in the respective two-and three-way contexts, compared with corresponding ratings of 5.0 and 5.8 for the categorically related pairs. One striking example involves the triplet skin, hair, and bar. For separate contexts the ratings were skin-bark=6.6, skin-hair=4.7. For the combined context the ratings were skin-bark=5.3; skin-hair=5.5. Again the interaction of comparison type with context was reliable.

This study is consistent with the idea that different standards of comparison may be used as a function of the comparison context. The combined contexts introduced contrast sets that apparently were not naturally evoked in separated contexts. A natural consequence of different standards of comparison is that different "respects" will enter into similarity judgments. As we have seen, the interactions are sufficiently large that in some cases one can observe reversals in the A to B versus A to C similarity as a function of the contrast sets which are triggered by a comparison context. These and other results support the general claim that similarity judgments and the respects that enter into them depend on the standard or contrast set. Most importantly the contrast set is context-dependent rather than fixed. This further underlines the importance of both structure and process for understanding similarity comparisons.

C. Mapped onto judgments by some process F.

1. The relativity of judgment. Let us focus for a moment on the rating process itself. The psychological similarity of two entities associated with comparison often must be mapped onto some quantitative or categorical scale. That is, how does one decide whether the similarity of cats to dogs should be associated with a rating of 4, as opposed to 6, on a 9-point scale? Parducci (1965) has argued that the entire set of comparisons determines the rating that a particular pair receives. A moderately similar pair of things may get a rating of 7 in the context of many highly dissimilar comparisons, but get a rating of 4 in a context of highly similar comparison. Parducci's (1965) range-frequency theory is designed to capture these context effects.

Although context effects on ratings are an important methodological consideration, one might wonder if they have anything to do with (psychological) similarity itself. There is suggestive evidence that the answer is yes. The idea is that psychological similarity may influence ratings and that the ratings in turn may feed back to influence apparent similarity. It is as if the rating is used to make an attribution about the similarity. If one gives a pair of entities a high similarity rating, one might tend to think that the things are, in fact, highly similar. Although we know of no direct evidence for this conjecture there are fairly close parallels in the social judgment literature. For example, Schwarz and Strack (in press) report that context effects induced for people's ratings of how often they watch television influenced their later judgments about how happy they were with the recreational aspects of their lives. Specifically, the scales were manipulated such that a given amount of TV watching appeared to be relatively high or relatively low. People who rated in the high context reported being less happy with their recreation than people who rated in the low context. It appears, then, that rating contexts are not neutral with respect to either rated similarity or psychological similarity.

2 Non-independence of feature weighting.

A recent series of experiments by Goldstone, Gentner and Medin (1991) indicates that feature weighting is not independent of the outcome of the comparison process. Figure 3 shows a sample set of stimuli from one of their experiments. Participants were shown different pairs of alternative stimuli and asked to judge which of the alternatives was more similar to the standard. The alternatives varied with respect to the types of matches they shared with the standard. For example, A contains the most attributional matches (triangle, circle, shading), D contains the most relational matches (two figures with the same shape, all figures with the same shading) and B and C are intermediate in both attributional and relational matches.

Insert Figure 3 about here

Note that as one moves from A and C to B and D, one attributional match (shading) is removed and one relational match (same shading between two elements) is added. Likewise, as one moves from A and B to C and D one attributional match is taken away and one relational match is added. This design carries the implication that if features are evaluated independently, then A should be picked as more similar to T than C to the same degree that B is picked as more similar to T than D. In addition, A should be picked over B to the same extent that C is picked over D. What should not happen is preferring A to B and C and also preferring D to B or C.

The results revealed systematic violations of independence. To be specific, the violations were in the direction of choosing A and D over B and C on the paired tests. Goldstone et al referred to these violations as "max" effects. The idea is that whatever type of similarity predominates will have a disproportionate influence on judgments. If the choices are attributionally similar to the standard then an extra attributional match has more weight than an extra relational match; if the alternatives are relationally similar then the relative weighting of attributes and relations reverses. The idea is that attributes and relations form distinct "pools" or types of similarity and shared features affect similarity more if the pool they are in is already relatively large.

The correct interpretation of the basis for these Max affects is not entirely clear. One idea is that participants find the judgment task to be ambiguous and assume that whichever type of similarity is maximized must be what the experimenter had in mind. A plausible alternative is that Max effects are directly tied to the comparison process. What is clear is that the weight given to a particular match depends importantly on the other matches in the scene (see also Goldstone and Medin, in press). Again we must attend to the processes by which matches and mismatches map onto judgments.

D From perspective G. In order to know the similarity of two things, the particular perspective of the person making the comparison must be known. similarity appears to be influenced both by momentary shifts in a persons perspective and ones permanent point of view. Barsalou (1982, 1983; also see Roth & Shoben, 1983) has demonstrated the subjects rate raccoons and snakes as being more similar when evaluating the animals as possible pets than when given no specific purpose. Even the similarity of items as apparently different as children and jewelry increases under the appropriate perspective (i.e., "things to retrieve from a fire"). When subjects are explicitly instructed to adopt a particular viewpoint (e.g., Chinese, redneck, or housewife), the similarity of items to categories varies widely.

Expertise and culture can produce relatively permanent changes in perspective that influence similarity. Chi, Feltovich, and Glaser (1981) find that novices group together physics problems on the basis of superficial characteristics (e.g., whether there is an inclined plane). Experts group physics problems on the basis of the underlying physical law instantiated (e.g., whether the problem involves the second law of thermodynamics). Experience with a language can increase or decrease the similarity of two sounds, where similarity is measured by peoples ability to discriminate the sounds (Liberman, Harris, Hoffman, & Griffith, 1957). Other examples of cultural influences on similarity are presented by Whorf (1941).

Examples of perspective-dependent similarity are abundant, and provide much support for the thesis that similarity is constructed by the comparison-maker, and is not only a physical property of the compared items. Once again, the dependence of similarity on perspective does not imply that similarity is hopelessly flexible. It may be true that from some perspective, any two items can be made to seem similar (Murphy & Medin, 12985). However, research on culture, expertise, and development gives reason to believe that there are systematic constraints placed on naturally adopted perspectives, and that true shifts in perspective are infrequent. Koreans have difficulty perceiving the dissimilarity of /J/ and /Z/ speech sounds, and Americans have difficulty conceiving of the similarity between a branch and a broken toe, objects that are conceptually similar to Hopi Indians (Whorf, 1941).

IV. Implications.

Let us briefly review our agenda. Nelson Goodman has argued that the similarity of A to B is meaningless unless one can specific the <u>respects</u> in which they are similar. He correctly pointed out that whatever fixes respects is doing the explanatory work and that the abstract notion of similarity is empty and only gives the illusion of being coherent and useful. Our theses is that similarity involves multiple predicates and that similarity statements of the form "A is similar to B" are really shorthand for "A is similar to B in respects C according to comparison process D, relative to some standard E mapped onto judgments by some function F for some purpose G." Furthermore the predicates of similarity form a critical research agenda because understanding similarity entails understanding its predicates.

At a very superficial level our experiments show that similarity is even more flexible than Goodman suggested. For example, we noted that in some cases that the features some entity "has" depend on what is is compared with. At the same time, however, similarity is systematically fixed in context. The entities entering into a comparison jointly constrain each other's interpretation and jointly determine the outcome of a similarity comparison. The predicates of similarity work together to allow one person to understand what another person means when they make a similarity statement, even when they do not explicitly specify respects. The comparison processes are dynamic but they are also lawful.

Our results are especially compatible with research on analogy. Indeed our concern with alignment and structure is motivated by the work of Gentner and others. Although early work on analogy treated the representation of the base and target of the analogy as fixed, more recently ideas about dynamic re-representation and comparison-dependent inferences have been advanced (e.g. Gentner, 1989; Gentner and Rattermann, 1991; Hofstader and Mitchell, in press). No doubt further research will pursue these parallels.

At the risk of oversimplification and presumptiveness, we might say that greater attention must be paid to two aspects of similarity: structure and process. Although the idea of treating entities as comprised of lists of features has led to some powerful and innovative ideas, we believe that feature lists have inherent limitations. Our work suggests that attributes and relations are psychologically distinct and other evidence shows that both attributes and relations may involve levels of abstraction (in some cases being hierarchically organized). These organizational properties require a sophisticated alignment process that has been ignored by current theories of similarity. Once we allow for structurally richer representations than fixed feature sets, we also must devise more intricate processes to deploy these structures. Sophisticated processes may not be required to match lists of features, but they will certainly be required when entities are composed of internally-complex parts that must be aligned, descriptions at multiple levels of abstraction, and context-sensitive characterizations.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. Memory & Cognition, 10, 82-93.

Barsalou, L.W. (1983). Ad hoc categories. Memory & Cognition, 11, 211-227.

Chi, M.T.H; Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Clement, C. A. & Gentner, D. (1991). Systematicity as a selectional constraint in analogical mapping. Cognitive Science, 15, 89-132.

Falkenhainer, B. (1990). Analogical interpretation in context. Proceedings of the Twelfth Annual Conference of the Cognitive Science Society. Hillsdale,NJ: Erlbaum.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989/90). The structure-mapping engine: Algorithm and examples. Artificial Intelligence, 41, 1-63.

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. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), Similarity and analogical reasoning New York: Cambridge University Press.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Brynes (Eds.). Perspectives on thought and language: Interrelations in development. London: Cambridge University Press.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. Cognitive Science, 10, 277-300.

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. Psychological Review, 97, 3-18.

Goldstone, R.L. (1991). Similarity, Interactive Activation, and Mapping. Unpublished Doctoral Dissertation. University of Michigan.

Goldstone, R.L. & Medin, D. L. (in press). Similarity,Interactive Activation, and Mapping: An overview. In K. Holyoak and J Barnden (Eds.) Advances in Connectionist and Neural-Computation Theory, Vol. 2: Connectionist Approaches to Analogy,Metaphor, and Case Based Reasoning. Norwood,NJ: Ablex Publishing.

Goldstone, R. L., Gentner, D., & Medin, D. L. (1989). Relations relating relations. Proceedings of the Eleventh Annual Conference of the Cognitive Science Society] , Ann Arbor, MI. Hillsdale, NJ: Erlbaum.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. Cognitive Psychology, 23 , 222-262.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), Problems and projects. New York: The Bobbs-Merrill Co.

Hofstader, D. & Mitchell, M. (in press). The copycat project: An overview. In K. Holyoak and J. Barnden (Eds.), Advances in Connectionist and Neural Computation Theory, Vol. 2: Connectionist Approaches to Analogy, Matephor, and Case-based Reasoning. Norwood, NJ: Ablex Publishing.

James, W. (1890). The Principles of Psychology. New York:Holt.

Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. Psychological Review , 85, 445-463.

Liberman, A.M., Harris, K.S., Hoffman, H.S., & Griffith, B.C. (1957). The discrimination of speech sounds within and across phoneme boundaries. Journal of Experimental Psychology, 54, 358-368.

Markman, A. B., & Gentner, D. (1990). Analogical mapping during similarity judgments. Proceedings of the Twelfth Annual Conference of the Cognitive Science Society, Cambridge, MA. Hillsdale, NJ: Erlbaum.

Markman, A. B. & Gentner, D. (1991). Communalities, differences, and the alignment of conceptual frames during similarity judgments. Proceeding of the Thirteenth Annual Conference of the Cognitive Science Society. Hillsdale,NJ: Erlbaum.

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

Nosofsky, R.M. (1989). Further tests of an exemplar-similarity approach to relating identification and classification. Perception & Psychophysics, 45, 279-290.

Nosofsky, R. M. (1992). Similarity scaling and cognitive process-models. Annual Review of Psychology, 43, 25-53.

Ortony, A. (1979). Beyond literal similarity. Psychological Review, 86, 161-180.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. Psychological Review, 97, 185-200.

Palmer, S. E. (1978). Structural aspects of visual similarity. Memory & Cognition, 6, 91-97.

Parducci, A. (1965). Category judgment: A range-frequency model. Psychological Review, 72, 407-418.

Roth, E.M. & Shoben, E.J. (1983). The effect of context on the structure of categories. Cognitive Psychology, 15, 346-378.

Schwartz, N. & Strack, F. (in press). Context effects in attitude surveys: Applying cognitive theory to social research. In W. Stroebe and M. Hewstone (Eds.), European Review of Social Psychology, Vol. 2.
Chichester, England: Wiley.

Shanon, B. (1988). On similarity of features. New Ideas in Psychology, 6, (3), 307-321.

Sjoberg, L. (1972). A cognitive theory of similarity. Goteborg-Psychological Reports, 2(10).

Smith, L.B. & Heise, D. (in press). Perceptual similarity and conceptual structure. In B. Burns (Ed.), Percepts, Concepts, and Categories: The Representation and Processing of Information. New York: Elsevier.

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

Tversky, A. & Gati, I. (1978). Studies of similarity. In E. Rosch and B.B. Lloyd (Eds.), Cognition and Categorization. Hillsdale, NJ: Erlbaum.

Whorf, B. L. (1941). Languages and logic. Technology Review (M.I.T.), 43, 250-252, 266, 268, 272.

Wisniewski, E. J., & Medin, D. L. (1991). Harpoons and longsticks: The interaction of theory and similarity in rule induction. In D. Fisher & M. Pazzani (Eds.), ComputationalApproaches to Concept Formation. San Mateo, CA: Morgan Kaufman.

Table 1.
Seventeen word pairs used in Asymmetry

X Y	Similarity of X to Y	Similarity of Y to X
England, United States	12.84	11.40
China, Korea	12.21	11.35
Russia, Poland	11.21	10.39
Wisconsin, Michigan	11.26	11.65
Prunes, Apples	9.84	8.77
Blimps, Cars	5.00	6.12
Spanish, English,	8.16	7.94
Skateboards, Bicycles	10.37	11.24
Albert Einstein, Benjamin Franklin	11.16	11.88
Squirrels, Mice	10.84	11.94
Cows, Dogs	8.37	7.47
Gorillas, Elephants	8.42	7.50
Physics, Philosophy	5.53	5.65
Doctors, Engineers	9.79	10.94
Chocolate bars, Popcorn	7.11	7.47
Frisbees, Boomerangs	11.16	12.94
Pencils, Crayons	12.68	15.13

Figure Captions

Figure 1. Representative stimuli from experiment on comparison-dependent features. A and C's interpretations are "carried over" to B.

Figure 2. Sample stimuli from the Markman and Gentner (1990) experiments. On superficial grounds the woman in (a) should be matched with the woman in (b). With respect to roles, however, the squirrel in (b) is the recipient and should be matched with the woman in (a) who is also the recipient.

Figure 3. Sample stimuli from the Goldstone, Medin, and Gentner (1991) experiments. As one moves from A and C to B and D an attributional match with T is removed (shading) and a relational match is added (same shading). Similarly, in going from A and B to C and D an attributional match is (triangle) deleted and a relational match is (same shape figures) is added.

Representative Stimuli from Set 1: A and C's interpretations are "carried over" to B

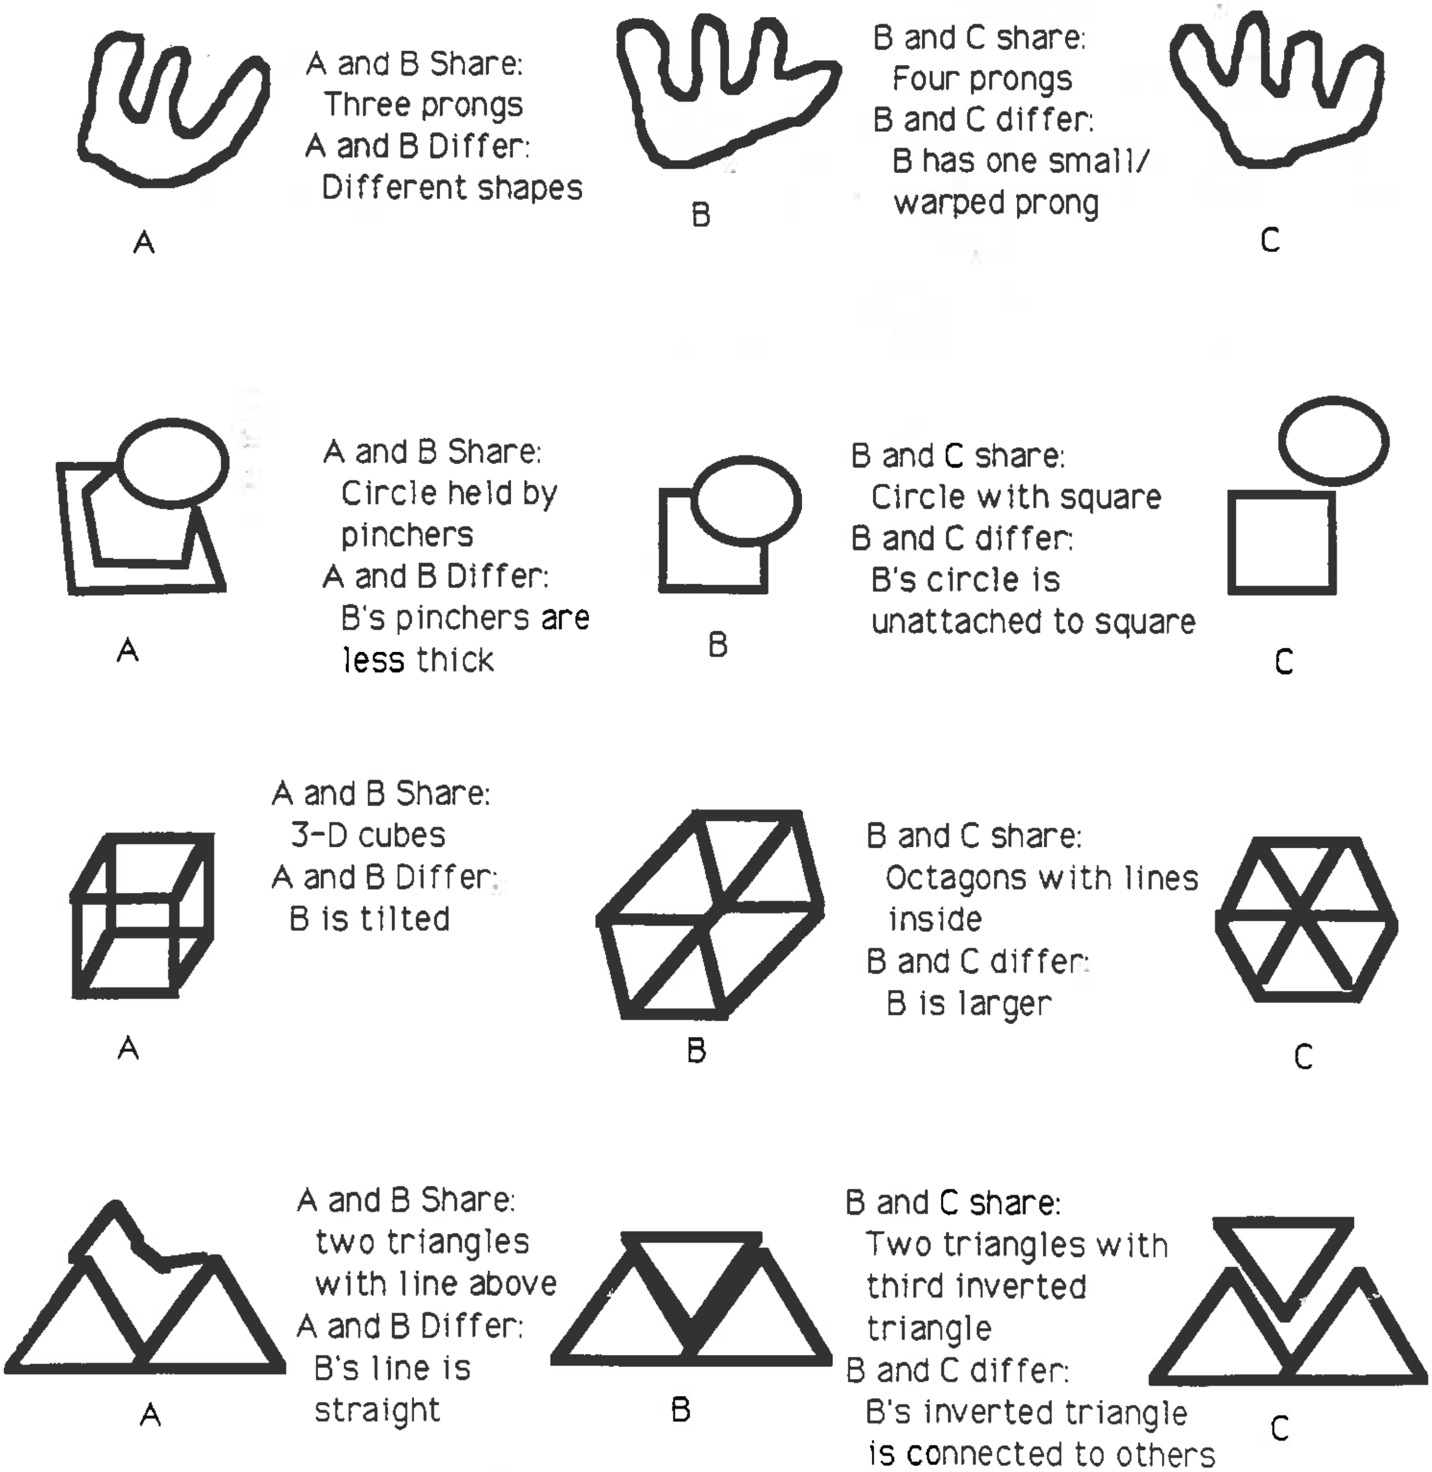

Figure 1

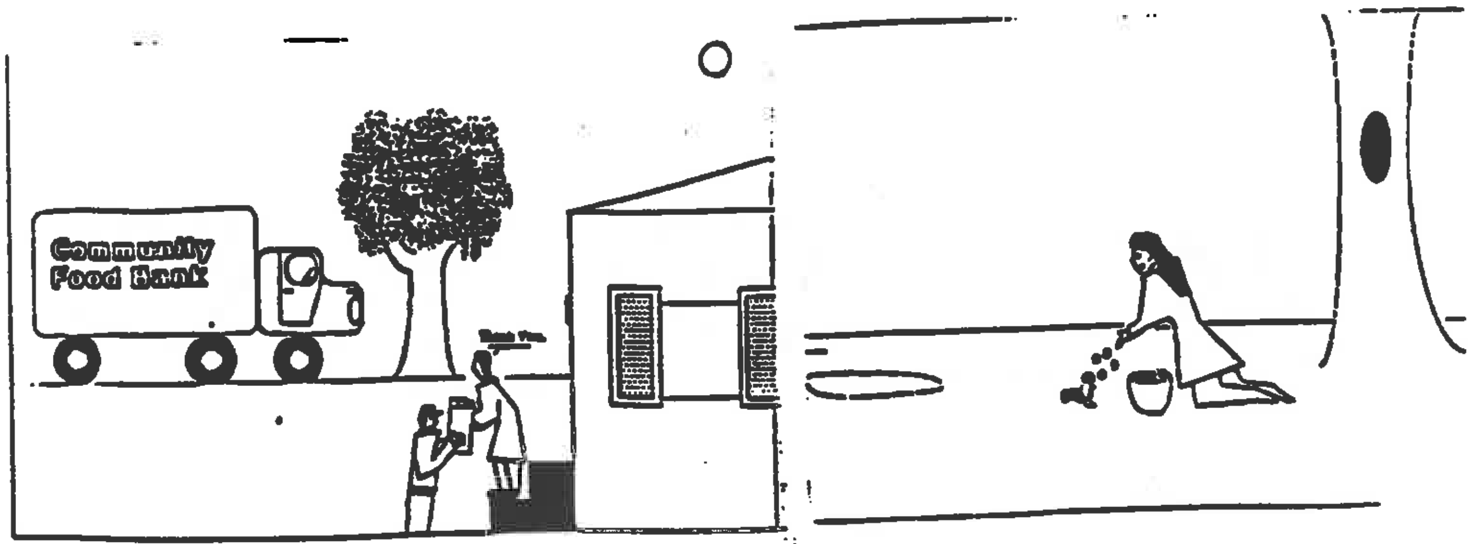

(a) (b)

Figure 2

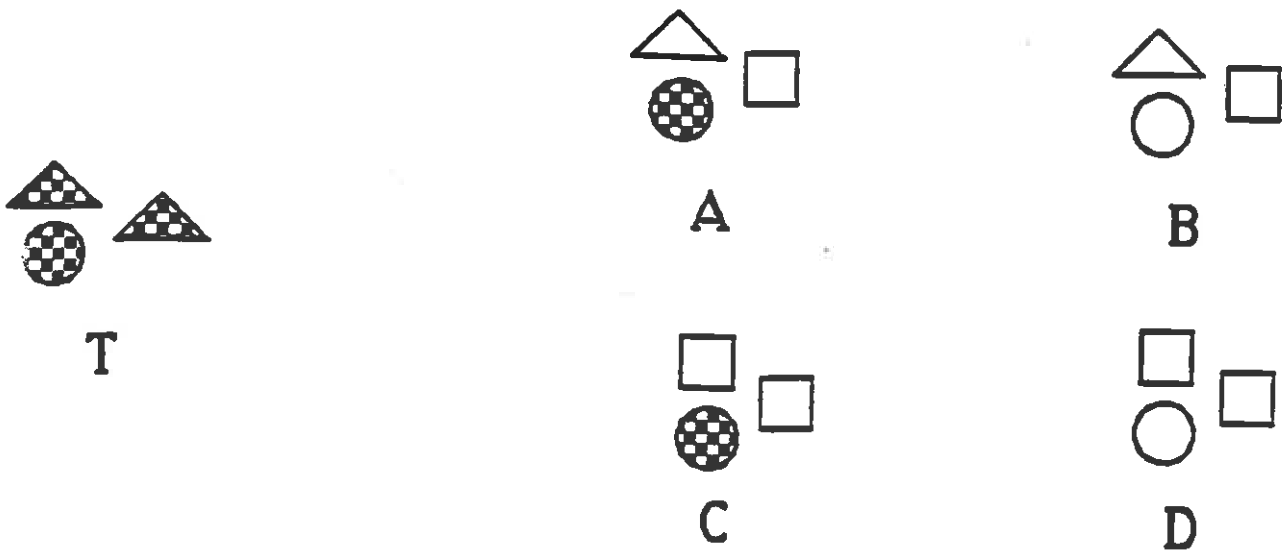

Figure 3