COGNITIVE PSYCHOLOGY 25, 431–467 (1993)

Structural Alignment during Similarity Comparisons

ARTHUR B. MARKMAN

University of Illinois

AND

DEDRE GENTNER

Northwestern University

Similarity comparisons are a basic component of cognition, and there are many elegant models of this process. None of these models describe comparisons of structured representations, although mounting evidence suggests that mental representations are well characterized by structured hierarchical systems of relations. We propose that structured representations can be compared via *structural alignment*, a process derived from models of analogical reasoning. The general prediction of structural alignment is that similarity comparisons lead subjects to attend to the matching relational structure in a pair of items. This prediction is illustrated with a computational simulation that also suggests that the strength of the relational focus is diminished when the relational match is impoverished, or when competing interpretations lead to rich object matches. These claims are tested in four experiments using the one-shot mapping paradigm, which places object similarity and relational similarity in opposition. The results support the hypothesis that similarity involves the alignment of structured representations.

INTRODUCTION

Similarity is a central component of models of a variety of cognitive processes. Categorization models assume that new exemplars are classified based on their similarity to some prototype, abstraction, or previous

Address correspondence and reprint requests to Arthur B. Markman, Department of Psychology, Columbia University, 406 Schermerhorn Hall, New York, NY 10027. This research was supported by NSF Grant BNS-87-20301 given to the second author, and a University Fellowship given to the first author by the University of Illinois. Research on the SME simulation was supported by ONR Grant N00014-89-J1272 given to the second author. We thank Neal Cohen, Gary Dell, Cindy Fisher, Ken Forbus, Doug Medin, Greg Murphy, Steve Palmer, and Brian Ross for extensive comments on earlier versions of this work. We also thank Connie Boronat, Mutsumi Imai, Phyllis Koenig, Laura Kotovsky, Mary Jo Rattermann, Ed Wisniewski, Phillip Wolff, and the rest of the Similarity and Analogy group for many stimulating and inspiring discussions of these issues. Thanks also to Lawrence Taylor for 12 great seasons. We also thank Rob Goldstone who collaborated with the second author on pilot research that laid the groundwork for these studies.

0010-0285/93 $9.00

exemplar (i.e., Rosch, 1975; Medin & Schaffer, 1978; Nosofsky, 1986). Theories of problem solving propose that new problems are solved using previous similar problems as examples (Ross, 1987, 1989; Holyoak & Koh, 1987; Novick, 1988). Norm theory presupposes that new situations are categorized based on prototypical instances constructed from prior similar episodes (Kahneman & Miller, 1986). Theories of transfer propose that new skills will be easier to learn to the extent that they are similar to things previously learned (Thorndike & Woodworth, 1901; Singley & Anderson, 1989). Because of its importance, similarity itself has been the focus of a large body of research.

This work has led to a number of elegant and sophisticated models. One in particular is Tversky's (1977) contrast model, a feature-based approach formalizing the fundamental insight that the similarity of a pair increases with its commonalities and decreases with its differences. This influential model has provided a basis for inquiries into the nature of psychological similarity and has been incorporated into models of other cognitive processes. According to this model, objects are represented by features drawn from a universal set. The *common features* of a pair are simply those elements in the intersection of the feature sets. The *distinctive features* of the pair are the elements in the feature sets of only one object (the contrast model assumes that the distinctive features of each object are kept in separate sets). By this model, the psychological similarity of a pair increases with the size of the set of common features and decreases with the size of the sets of distinctive features.

However, there is a striking lacuna in this work. While models of similarity have focused on featural representations, mounting research in cognitive psychology indicates that for many purposes, psychological representations are better characterized by structured hierarchical systems. For example, in an effort to explain the global perceptual properties described by the Gestalt psychologists (e.g., Goldmeier, 1972), Palmer (1977) suggested that perceptual representations consist of configural relations between perceptual units (like line segments). These relational groups may be hierarchically organized by additional relations. Similar representations have been presented by Biederman (1987) in a model of object recognition and Hinton (1979) and Tye (1991) to explain mental imagery phenomena. Even connectionist models have begun to examine methods for storing structured representations (Smolensky, 1990; Pollack, 1990).

Palmer (1977) explored the role of structure in perceptual representations by asking subjects to "parse" a set of configurations constructed from line segments into their "natural parts" and found that they often divided figures into the predicted types of relational groupings. Similar evidence has been obtained by Lockhead and King (1977), who found that subjects were able to discriminate ((from () more easily than (from), because the first pair, but not the second, gave rise to emergent properties. Additional evidence for the importance of emergent properties comes from studies by Pomerantz, Sager, and Stoever (1977) and Treisman and Paterson (1984).

The importance of structure has also been examined in studies exploring the role of *relations* and *attributes* in similarity. Relations can be conceived as predicates that link two or more *arguments*, which may be objects, object descriptions (called *attributes*), or other relations. Goldstone, Medin, and Gentner (1991) found that subjects can be differentially sensitive to attributes and relations. They asked subjects to select which of two comparison objects was similar to the target. The comparison figures shared either mostly attribute similarity with the target, or mostly relational similarity with the target. Their results indicate that increasing the relational similarity of two scenes has a greater impact on their perceived similarity if the commonalities of the scenes are primarily relational than if the commonalities are primarily attributional. Similarly, increasing the attributional similarity of two scenes has a greater impact on similarity if the commonalities of the scenes are primarily attributional than if the commonalities are primarily relational. To account for these data, Goldstone et al. proposed the MAX model, which assumes that matching attributes and matching relations form two "pools." During a similarity judgment, elements in the larger pool are given more weight than elements in the smaller pool.

Other work has demonstrated that subjects often find relational matches more compelling than attribute matches. For example, Rattermann and Gentner (1987; Gentner, Rattermann, & Forbus, in press) and Schumacher and Gentner (1987) found that stimuli containing only relational commonalities were often considered to be more similar than stimuli containing only object commonalities. In addition, Goldstone, Gentner, and Medin (1989) found that stimuli containing many consistent (i.e., same polarity) relations (e.g., bigger and darker) were thought to be more similar than stimuli containing an equal number of inconsistent (i.e., opposite polarity) relations (e.g., smaller and darker).

STRUCTURE AND COMPARISON

This evidence suggests that the process of determining the similarity of a pair must be sensitive to relations in stimulus representations. This paper attempts to reconcile the insights gained from featural models of similarity with the importance of structure in similarity. In order to model the comparison of structured representations, a mechanism must be developed to handle the connections between elements in representations.

We will illustrate this point with an example. Suppose the task were to compare assertions about two different planets:

cause (**greater** (mass(Sun), mass(Jupiter)), **revolve**(Jupiter,Sun))[1] (1)
and
cause (**greater** (mass(Sun), mass (Mars)), **revolve**(Mars, Sun)), (2)
where similar objects share some attributes. This pair is literally similar because it contains both relational commonalities and object similarities. Thus, Jupiter corresponds to Mars both because, as planets, they share many physical similarities and because they play a common role in the matching relational structure. In this pair, object and relational commonalities are correlated, so the separate impact of relational alignment on similarity cannot be assessed.

However, object and relational similarity can be separated by designing comparisons in which one of the objects is cross-mapped (i.e., the object appears in both relational structures, but plays a different role in each (Gentner and Toupin, 1986)). For example, in the comparison of

cause (**greater** (mass(Sun), mass (Jupiter)), **revolve**(Jupiter,Sun)) (1)
and
cause (**greater** (mass(Jupiter), mass (Io)), **revolve**(Io,Jupiter)), (3)
Jupiter is cross-mapped. If we attend to relational similarities, the Sun in (1) is placed in correspondence with Jupiter in (3), because things revolve around both of them. However, if we attend to attribute similarities, then Jupiter in (1) is placed in correspondence with Jupiter in (3) because they are the same planet. A simple combination of these solutions is not sufficient, because the final match between (1) and (3) must reflect that focusing on different types of similarity leads objects to be placed in correspondence in different ways.

We propose that this difficulty can be captured by viewing comparison as structural alignment. Structural alignment is the underlying process associated with many models of analogical reasoning (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 1989; Greiner, 1988; Holyoak & Thagard, 1989; Keane, 1990; Winston, 1982; cf. Hall, 1989, and Kedar-Cabelli, 1985, for general reviews). According to Gentner (1983, 1989), the match between two structured representations must be *structurally consistent:* that is, it must conform to the *one-to-one-mapping* and *connectivity* constraints. One-to-one-mapping means that for any given match between representations, each element in one representation will map to at most one element in the other representation. Connectivity

[1] This assertion can be read as: The fact that the mass of the sun is **greater** than the mass of Jupiter **causes** Jupiter to **revolve** around the Sun.

mandates that if a match is made between predicates, the arguments of those predicates must match as well.

In many cases, more than one structurally consistent match between two representations is possible. However, only a single interpretation will be used as the basis for finding the commonalities and differences when determining the similarity of a pair. The preferred match is generally the one that is most *systematic*. Systematicity requires that deeply connected relational matches be preferred to matches that preserve only scattered, unconnected relations (even holding the number of relations constant).

The central claim of this research is that *similarity comparisons involve a process of structural alignment*. For example, carrying out the comparison of (1) and (3) would tend to result in placing Jupiter in (1) and Io in (3) in correspondence, because they play similar roles in the matching relational structure. This relational match would be made in spite of the initially appealing match between Jupiter in (1) and Jupiter in (3). On this account, the likelihood of a structural alignment between two stimuli is greater after a similarity comparison than before.

We tested this prediction using the same logic as in the example above. The idea was to create stimuli with an attractive local mapping and test whether carrying out similarity comparisons would lead subjects to neglect this mapping in favor of one based on structural alignment. We presented subjects with pairs of scenes containing cross-mappings and asked them to perform a *one-shot mapping*. In the one-shot mapping task, an experimenter points to the cross-mapped object in one scene and asks the subject to select the object in the other scene that goes with that object. The key manipulation was whether subjects were first asked to assess the similarity of the two scenes.

The materials were deliberately constructed to make the object match salient. For example, Fig. 1 presents a sample stimulus from the first experiment. The resemblance between the two women is immediately apparent. Therefore, we would expect subjects who perform a simple one-shot mapping to align the two women based on their perceptual similarity. However, it can also be seen that the women in these scenes are cross-mapped. The woman in Fig. 1a is *receiving* food from the man, while the woman in Fig. 1b is *giving* food to the squirrel. If similarity comparisons involve structural alignment, then subjects who first rate the similarity of the pairs should align objects based on their position in the matching relational structure in the one-shot mapping task. In this example, subjects would place the woman in Fig. 1a in correspondence with the squirrel in Fig. 1b.

In addition to these two conditions, we included a convergent task to assess subjects' ability to make relational mappings for the stimulus pairs. For this purpose, subjects were asked to make three object mappings for

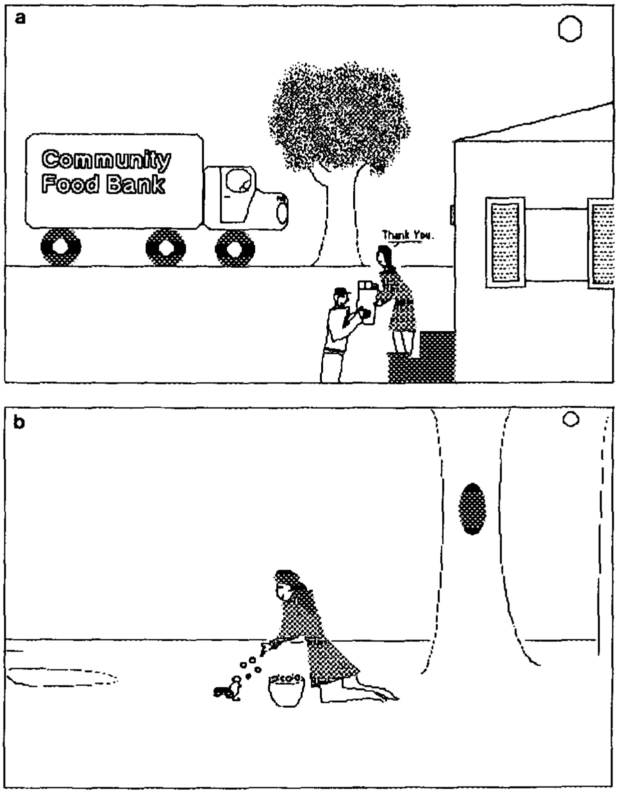

FIG. 1. Sample pair of causal scenes containing a cross-mapping. The woman in the top scene is receiving food, while the woman in the bottom scene is giving food away.

each pair of scenes. This more comprehensive mapping task, which requires that subjects consider several simultaneous object correspondences, should force them to align the objects based on the matching relational structure. In the current example, while the women in Figs. 1a and 1b can be matched on the basis of perceptual similarities, there is no perceptual match in Fig. 1b for the man in Fig. 1a. The only plausible match for this object is the woman in Fig. 1b who is also giving away food. However, we assume that subjects have a bias against mapping an object in one scene with more than one object in the other scene. Hence, they should rarely place *both* the woman and man in Fig. 1a in correspondence with the woman in Fig. 1b. These predictions will be tested in Experiment 1.

EXPERIMENT 1

Experiment 1 uses the one-shot mapping technique to test the prediction that similarity comparisons promote structural alignment. For this purpose, eight pairs of causal scenes were drawn. Each pair contained a cross-mapped object. A sample stimulus pair is presented in Fig. 1. With this stimulus set, subjects were run in one of the three experimental conditions described above. Subjects in the *1map* condition performed a one-shot mapping on each scene and later rated the similarity of all the scenes. Subjects in the *Similarity-first* condition (Sim → 1map) first rated the similarity of the scenes and then performed the one-shot mapping. Finally, subjects in the *3map* condition performed three mappings for each pair of scenes and later rated their similarity.

The predictions are, first, subjects in the Sim → 1map condition should map on the basis of the matching relational structure. This result is predicted on the grounds that carrying out a similarity comparison will cause subjects to align the scene representations, thereby increasing their sensitivity to the common relations. Second, subjects in the 1map condition should map on the basis of local object similarities, because simple one-shot mapping does not require a global comparison of the scenes. Third, subjects in the 3map condition should make many relational responses, because, as described above, the common relations provide a natural basis for multiple consistent mappings. We make no ordinal predictions about the level of relational responding in the Sim → 1map condition relative to the 3map condition. We simply expect both conditions to exhibit a higher level of relational responding than the 1map condition.

In addition to the 3map condition, one other control is needed. Subjects in the Sim → 1map condition are exposed to the pictures prior to the mapping task, while subjects in the 1map and 3map conditions are not. Thus, if subjects in the Sim → 1map condition make more relational responses than subjects in the 1map condition, this difference could be explained by their greater degree of familiarity with the stimuli. A *familiarity control* condition is needed to guard against this possibility. Subjects in this condition are shown all of the pictures individually for 5 s—roughly the amount of time subjects see the pictures when rating their similarity—and are told to study them for a later memory test. Following the study stage, they perform the one-shot mapping task. Subjects in this condition should make fewer relational responses than subjects in the Sim → 1map condition.

Method

Subjects. Subjects were 48 undergraduates (12/condition) at the University of Illinois who received $4.00 or course credit in introductory psychology classes for their participation.

Design. Subjects were randomly assigned to one of the three between-subject Mapping conditions or the familiarity control. Four random orders of stimulus presentation were used in each condition. The performance of subjects in the Familiarity Control condition will be compared only to the performance of subjects in the Sim → 1map condition.

Stimuli. The stimuli were eight pairs of pictures portraying causal scenes (like the pair in Fig. 1). Each pair of scenes contained a cross-mapping. In the context of this experiment, a cross-mapping was operationalized as a pair of perceptually similar objects that played different roles in the matching relational structure of the two scenes. In half of the pairs, the perceptually similar objects occupied roughly the same spatial position in both scenes (as with the two women in Fig. 1). In the other half, the objects playing the same role occupied the same spatial position in both scenes. Further, in half of the scenes, the event path moved from left to right (e.g., the giver was on the left and the receiver was on the right) and in the other half the event path moved from right to left. Finally, in half of the scenes the cross-mapped objects were shown in the same left-right orientation and in half of the scenes the objects were flipped horizontally to face in the opposite direction.[2] A summary of the stimulus set is presented in Table 1.

Procedure. Subjects were run one at a time. They were seated at a table with an experimenter seated beside them. Subjects participated in only one experimental condition. The experimenter had no knowledge of the hypothesis being tested.

Subjects in the one-shot mapping (1map) condition were shown each of the base/target pairs in turn. The experimenter pointed to the cross-mapped object and asked the subject to point to the object in the other picture that went with it. The subject's response was recorded and the next pair of pictures was presented. After completing the mapping task, subjects rated the similarity of each pair on a scale from 1 to 9. Ratings were made orally.

Subjects in the three-mapping (3map) condition were also shown each of the base/target pairs. Before the first trial, subjects were told that they would be making three object mappings. Then, the experimenter pointed (one at a time) to three of the objects making up the central relational structure of one scene, and asked the subject to point to the object in the other scene that went with each object. The cross-mapped object was always tested first so that subjects' first responses in the 3map condition would be comparable to subjects' first responses in the other three conditions. After completing the mapping task, these subjects were also asked to rate the similarity of all of the pairs of pictures.

Subjects in the similarity-first task (Sim → 1map) were given a pair of scenes and told to rate their similarity on a scale from 1 to 9. After the subject completed the rating for a pair the experimenter pointed to the cross-mapped object and asked the subject to point to the object in the other picture that went with it. The responses were recorded and the next pair was presented.

In the familiarity control condition, subjects were shown the set of pictures one at a time and told to study them carefully for a later memory test. Subjects saw each picture for 5 s, roughly the amount of time subjects in the Sim → 1map condition saw the pictures while making similarity comparisons. After examining the entire set of pictures, subjects in the control condition performed the one-shot mapping task.

Results and Discussion

Comparison and mapping. For each pair, subjects made an *object mapping* if they responded on the basis of perceptual similarity (e.g., matching the two women), a *relational mapping* if they responded on the basis of the common relational structure (e.g., matching the woman with the squirrel) or a *spurious mapping* if any other choice was made (e.g., matching the woman to the tree). Almost all of the responses were either

[2] Although we were able to control these factors somewhat, they were not fully counterbalanced. However, a careful analysis of presentation factors is carried out in Experiment 3.

TABLE 1
Summary of Causal Stimuli Used in Experiment 1

Stimulus	Relation	Relation in Same Direction	Number of Objects	Cross-mapped Object in Same Orientation	Number of Cross-mapped Objects	Cross-mapped Object	Object for Multiple Mappings	
1	Reaching	No	4	Yes	1	Cookie jar	Person	Chair
2	Repair	Yes	3	No	1	Robot arm	Wrench	Car
3	Towing	No	3	No	1	Car	Truck	Hook
4	Pitching	No	4	No	1	Pitcher	Batter	Ball
5	Giving	Yes	4	Yes	1	Woman	Man	Bag of food
6	Execution	Yes	3	No	1	Pirate	Sharks	Rebel
7	Break-in	No	4	Yes	1	Box	Bandit	Knife
8	Eclipse	Yes	3	Yes	2	Earth	Moon	Sun

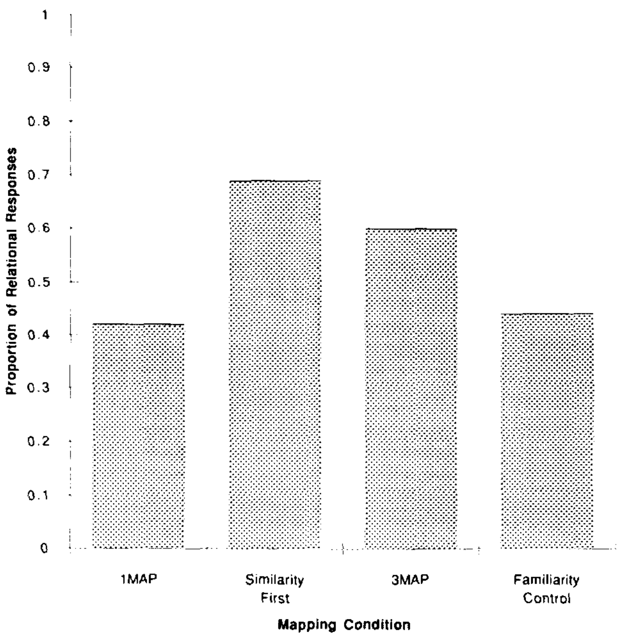

FIG. 2. Proportion of relational responses in each Mapping condition of Experiment 1 with the causal scenes.

object mappings or relational mappings. Spurious mappings accounted for less than 2% of all choices in all of the experiments we performed and will not be considered further.

Figure 2 depicts the proportion of relational responses by subjects in each condition of the experiment. A one-way ANOVA was performed on the three mapping conditions excluding the familiarity control. This analysis reveals a significant difference in the level of relational responding across the conditions, $F(2,33) = 3.83$, $p < .05$. As predicted, subjects in the Sim → 1map condition were more likely to map on the basis of the relational structure ($m = 0.69$)[3] than subjects in the 1map condition ($m = 0.42$), $F(1,33) = 7.30$, $p < .05$.[4] In addition, subjects were marginally more likely to map on the basis of the relations in the 3map ($m = 0.60$) condition than subjects in the 1map condition $F(1,33) = 3.50$, $.05 < p < .10$. Furthermore, subjects were also more likely to map on the basis of the relations in the Sim → 1map condition than they were if they simply saw the scenes singly for 5 s ($m = 0.44$), $t(22) = 2.08$, $p < .05$.

[3] All condition means in this paper will be reported as mean proportions.

[4] All planned comparisons in this paper used the Bonferroni procedure unless otherwise stated.

This result is consistent with the predictions of the structural-alignment view of similarity. Subjects who rated the similarity of a pair prior to making a one-shot mapping made more relational mappings than subjects who simply performed the one-shot mapping task. Since subjects in the 3map control also made more relational responses than subjects in the 1map task, we have additional evidence supporting the interpretation that subjects in the Similarity-first condition mapped on the basis of the matching relational structure. Finally, subjects in the familiarity control made few relational responses, indicating that mere familiarity with the pictures is not the factor underlying the difference in relational responding between subjects in the similarity first and 1map conditions.

As a check on the consistency of these results, we can examine subjects' mappings for individual stimulus pairs. Six of the eight pairs (75%) were given more relational responses in the Sim → 1map condition than in the 1map condition, $p > .10$ by sign test. The number of relational responses was equal in both conditions for the remaining two pairs. In addition, for five of the eight pairs (68%) subjects made more relational responses in the 3map condition than in the 1map condition, $p > .10$ by sign test. For two of the pairs an equal number of relational responses was given in the 3map and 1map conditions, while for one pair, more relational responses were given in the 1map condition than in the 3map condition. The general pattern of results examined by individual pairs is the same as that found in the subject analysis, but this pattern is not as strong.

We also determined the mean rated similarity for all stimuli in each condition. A one-way ANOVA showed no significant differences between conditions, $F(2,285) = 1.67$, $p > .10$. This result suggests that subjects in different conditions did not vary substantially in their feelings about the overall similarity of the stimuli.

Familiarity, Ratings, and Mapping

Before embracing an interpretation of the results in terms of the promotion of relational matches by similarity comparisons, we must consider an alternative explanation for these findings. Perhaps subjects in the Similarity-first condition attended to the stimuli more carefully than subjects in the 1map condition. The familiarity control condition ensured that the increase in relational responding for subjects in the Sim → 1map condition was not simply a result of having seen the pictures prior to the one-shot mapping. However, subjects in the familiarity control only viewed the pictures for a memory task, while subjects in the Similarity-first condition used the pictures to perform an explicit task. Perhaps it was this greater focus on the stimuli that led to the increase in relational responding for the Similarity-first subjects.

In order to rule out this possibility and to replicate the findings of

Experiment 1, we repeated the one-shot mapping task in Experiment 1b. As before, one group of subjects participated in the 1map condition, and a second group of subjects participated in the Similarity-first condition. However, a third group rated the artistic merit of the pairs of scenes prior to performing the one-shot mapping task. We reasoned that artistic merit ratings should lead to sustained attention to the materials, but they were not expected to involve a comparison. Therefore, if comparison is the crucial element in promoting structural alignment, then making artistic merit ratings will not lead to an increase in the level of relational responding relative to the 1map condition.

EXPERIMENT 1b

Method

Subjects. Subjects were 72 students from Northwestern University (24/condition). These subjects received course credit in introductory psychology for their participation.

Stimuli. Stimuli were the same eight causal scenes from Experiment 1. In this study, each pair of pictures was placed on a page of a booklet one above the other. In the mapping booklet, an arrow was placed above the cross-mapped object. In the similarity rating booklet, a line with the numbers 1 through 9 was placed at the bottom of the page. The words Low similarity appeared below the 1, and the words High similarity appeared below the 9. The artistry rating booklet was identical to the similarity rating booklet except that the words Low artistry and High artistry appeared beneath the 1 an 9, respectively.

Design. There were three between-subjects Mapping Conditions in this study: 1map, Sim → 1map and Art → 1map.

Procedure. Instructions on the first page of the one-shot mapping booklet asked subjects to draw a line from the object with the arrow over it in the top scene to the object in the bottom scene that best goes with that object. Instructions on the first page of the similarity booklet asked subjects to rate the similarity of the scenes on the scale provided. Finally, instructions on the first page of the artistry rating booklet told subjects that the different pictures were drawn by different artists. Subjects were asked to rate how artistic they felt the pictures were on the scale provided.

Subjects in the 1map condition received the one-shot mapping booklet followed by a rating booklet for another task.[5] Subjects in the Sim → 1map condition received the similarity rating booklet followed by the one-shot mapping booklet. Finally, subjects in the Art → 1map condition received the artistry rating booklet followed by the one-shot mapping booklet. The task took between 10 and 15 min to complete.

Results and Discussion

The results of this study are similar to the results of Experiment 1. A one-way ANOVA on these data indicates that the proportion of relational responses differs between conditions, $F(2,69) = 6.11$, $p < .005$. Deeper analysis reveals that subjects made more relational responses in the Sim

[5] The three conditions reported here were run as part of another study testing a different point than the one made here. Hence, subjects in the 1map condition actually made difference ratings on the pictures following the one-shot mapping task, rather than similarity.

→ 1map condition ($m = 0.69$) than in the 1map condition ($m = 0.47$), $F(1,69) = 6.00$, $p < .05$ or the Art → 1map condition ($m = 0.39$), $F(1,69) = 11.44$, $p < 05$. Item analyses support this result. More relational responses were made to all eight items in the Similarity-first condition than either the 1map or the Artistry-first conditions.

This study further clarifies the results of Experiment 1. First, because subjects in the Similarity-first condition made more relational responses than subjects in the 1map condition, this study provides further evidence for our focal hypothesis that similarity comparisons promote the relational alignment of scene representations. Second, this effect was obtained for all eight items suggesting that the finding is not limited to particular items. Third, the finding that subjects made fewer relational responses in the Art → 1map condition than in the Sim → 1map suggests that there is something special about *comparison* that causes subjects to shift from a match based on object similarities to a match based on relational similarities.

SME as a Process Model of Comparison

The results of the first two studies support the general prediction that similarity comparisons promote structural alignment, which, in turn, focuses subjects on the matching relational structure of the scenes. We now present a computational simulation of the alignment process in order to explain how the preference for relations in these stimuli might arise. The simulation will also reveal a set of more general predictions of structural alignment that we will test in further experiments.

The structure-mapping engine (SME; Falkenhainer et al., 1987, 1989) can be used to simulate subjects' performance in the Similarity-first condition. This computational model embodies the processing assumptions of structural alignment outlined above. The simulation is helpful because it makes clear that structural alignment does not always predict that similarity comparisons lead to a preference for common relations. Rather, it provides specific mechanisms for determining how a pair is similar and for deciding which match is best when multiple correspondences exist. Finally, it demonstrates that structural alignment is a computationally viable process.

SME takes two propositional representations composed of *entities, attributes, functions,* and *relations* and builds globally consistent matches by starting with local similarities. Entities correspond to the objects in a domain. Attributes are unary predicates that provide descriptive information about entities. Functions are predicates that map onto objects or values other than truth values and can be used to represent dimensions. Finally, relations are multiplace predicates that link two or more elements. Some relations link other relations or propositions, resulting in higher-order relational structure.

SME can be run in various modes. Here we describe Literal Similarity mode in which attributes, relations and functions participate in the mapping process. SME's other modes include Analogy (relations only) and Mere Appearance (attributes only). In the first step of the match, all identical predicates are placed in correspondence without regard to overall consistency (e.g., many-to-one mappings are tolerated). These local matches then coalesce into larger subsystems. Structural consistency is enforced within these clusters by checking whether matching predicates have matching arguments, while ensuring that each object in one representation maps to at most one object in the other. Further, new correspondences are proposed between non-identical elements based on structural consistency. If two relations have been placed in correspondence, then SME attempts to place their arguments in correspondence. Finally, these pieces of matching structure are grouped into maximal mutually consistent sets to form global interpretations or GMAPs. For each GMAP, candidate inferences are generated by completing the common structure. These inferences are useful for models of analogical transfer, but it is an open question whether these inferences play a role in determining subjective similarity.

We used an algorithm that generates all possible alternative interpretations of the match between two representations. Competing GMAPS resulting from this process are evaluated using a cascade-like algorithm: starting at the highest level relational matches in the GMAP, matching predicates pass evidence down to their arguments, and they to their arguments (Forbus and Gentner, 1989). This evaluation procedure implements a preference for deep systematic structures over shallow structures (even those containing the same number of predicate matches).

In order to simulate the Similarity-first condition, the pair of scenes from Fig. 1 was encoded into propositional form and submitted to SME. In creating scene representations, we assumed that subjects encode information about both the attributes of the objects and the relations between the objects. In addition, we assumed that perceptually similar objects are encoded by identical attributes, and that relational similarities are represented as identical relational structures. We postulated that subjects encode both the central relational structure of the scene as well as other relations peripheral to that central structure. Finally, we assumed that more higher-order relational structure is devoted to the central causal structure than is used to encode other relations in the scene. Figure 3 depicts the representations used to encode the pictures in Fig. 1.

Two of the interpretations generated by SME are shown in Fig. 4. Figure 4a presents the top rated GMAP (evaluation score 12.94), which

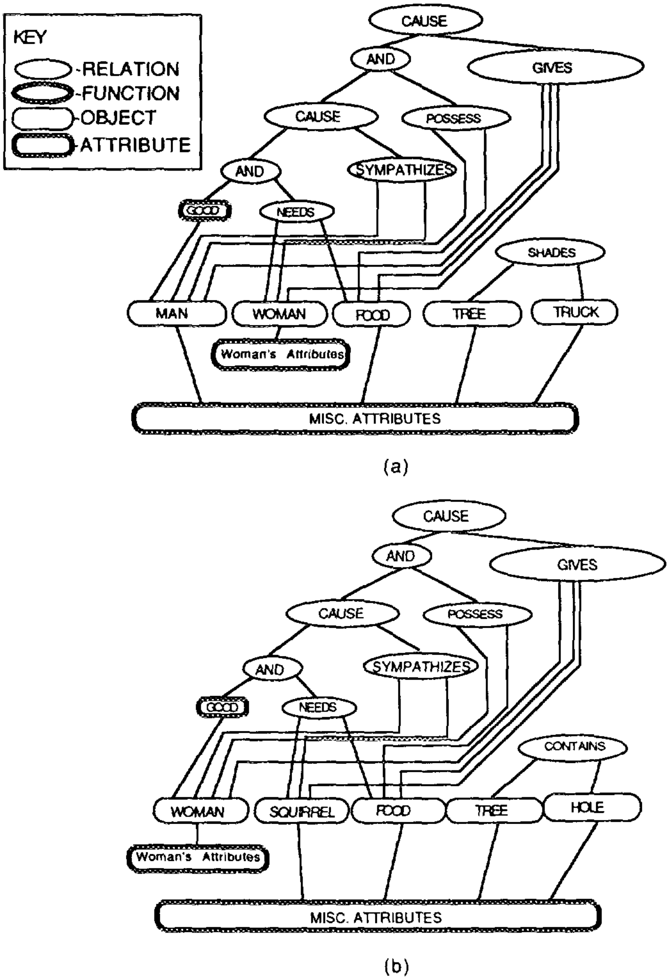

FIG. 3. Relational structures given to SME in a simulation of Experiment 1. These structures represent the scenes in Fig. 1.

was the interpretation based on matching the causal relational structure of the two scenes. In this interpretation, the man giving food is placed in correspondence with the woman giving food, and the woman receiving food is placed in correspondence with the squirrel receiving food (see Fig. 1). Figure 4b shows the most preferred object-based interpretation. Here, the matching attributes of the two women have been allowed to dominate the mapping, so the woman receiving food is placed in correspondence with the woman giving food. Very little of the common relational structure is preserved in this interpretation. This GMAP (evaluation score =

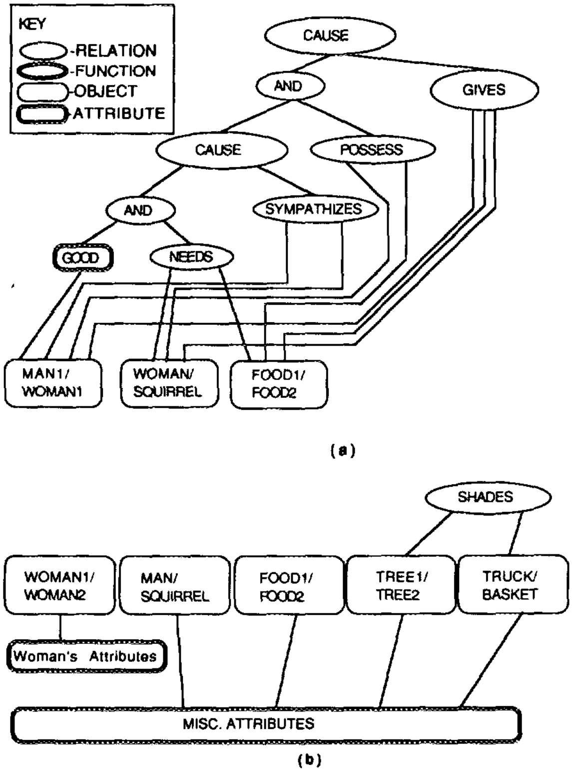

FIG. 4. Two interpretations arising from application of SME to the relational structures in Fig. 3. The relation-based GMAP (a) receives a higher evaluation score than the object-based GMAP (b).

11.28) received a lower evaluation score than the relational interpretation, but a higher score than all other possible interpretations (which ranged from 5.68 to 7.96).

The simulation results are consistent with the findings of Experiments 1 and 1b. SME recognized that the objects could be placed in correspondence based on attribute similarities or relational similarities, but preferred the interpretation based on the matching relational structure. For comparisons of deep causal structures, the preferred mapping is typically based on the relational commonalities between scenes. However, the presence of an object-based interpretation suggests that the object match would be preferred if it received a higher evaluation score than the relational match. An examination of the procedure for evaluating competing matches suggests some factors that determine when the object match is preferred to the relational match. Specifically, the evaluation of a match depends on the amount and depth of the matching relational structure and the extent of the object match (i.e., the number of common object attributes supporting a given object correspondence). Thus, an evaluation score can be increased either by increasing the extent of the matching relational structure (primarily the depth), or by adding more common attributes to corresponding objects.

This pattern of tradeoffs reveals a more general set of testable predictions of the structural alignment view of similarity. In the case of literal similarity, attribute and relational similarities lead to the same object correspondences, so that increasing either relational depth or object richness improves the match. However, when an object is cross-mapped, interpretations based on attribute similarity and those based on relational similarity are placed in opposition. This competition must be resolved to form the interpretation. In general, the relational match is selected as the preferred interpretation for a similarity comparison. However, all else being equal, decreasing the depth and coherence of a relational match should decrease subjects' preference for an interpretation based on relational commonalities. Similarly, increasing the extent of the object similarities should increase subjects' preference for an interpretation based on object similarities. Experiments 2 and 3 test these predictions.

In designing these studies, it was obviously crucial to select materials that would permit variations of relational structure or object richness. Therefore, we turned to perceptual stimuli (that is, pairs of pictures whose similarity resides in common perceptual relations like *above*, or *larger-than*). This change had the further advantage of allowing us to test the generality of the phenomenon. It could be argued that the causal scenes used in Experiments 1 and 1b are not representative of typical similarity comparisons. They embody rich causal schemas in which the objects each play identifiable and distinct relational roles. Thus, the results of the first two experiments may simply have confirmed that structural alignment is used to compare analogically similar pairs. Stimuli embodying perceptual relations are more like those traditionally used in studies of similarity, thus allowing a broader test of our hypotheses.

EXPERIMENT 2

Experiment 2 examines whether the relational focus observed in the first two experiments depends on the depth of the matching relational structure as suggested by the SME simulation. To create the necessary contrasts, we required perceptual scenes in which object similarity and relational similarity were placed in opposition. Thus, this study required stimulus pairs depicting relational structure in such a way that the objects would have clear relational correspondences within that structure. In addition, the stimuli had to be amenable to manipulations of the amount of matching relational structure. Finally, in order to provide a competitor to the relational match, we required objects that were perceptually rich. These constraints were satisfied by stimuli like those in Fig. 5, which were drawn to resemble patterned rugs.

The relational structure in these stimuli was provided by multiple sym-

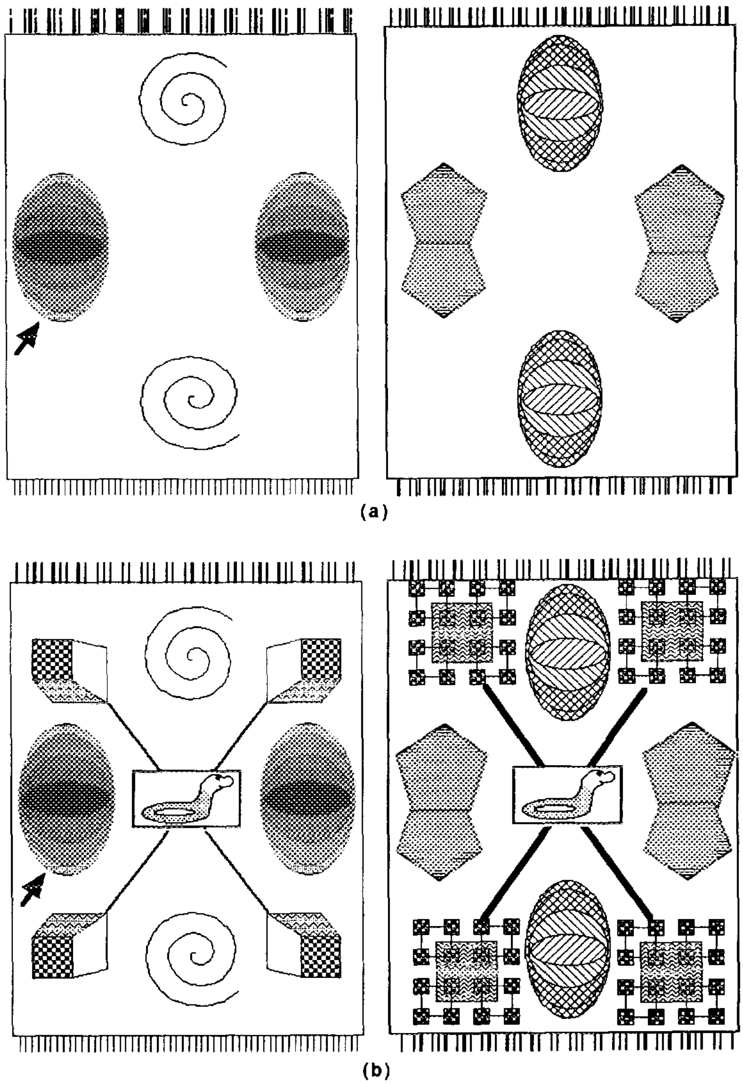

FIG. 5. Sample perceptual stimuli resembling rugs used in Experiment 2. One pair of Low and High Systematicity relational structures are pictured here.

metries. Kubovy (in preparation) suggests that symmetry relations have a powerful organizational effect on the perceptual system. Kubovy refers to the particular symmetries used in these stimuli as *reflection isometries*. That is, each object in a rug had a mirror image on the other side of the rug. Objects on the left were mirror images of objects on the right, and objects on the top were mirror images of objects on the bottom. Thus, the relational match to an object was the object in the same relative spatial position in the other rug.[6] The depth and coherence of the matching relational structure was varied by adding details that created further isometries, such as diagonal axes. The Low Systematicity pairs contained four objects, while High Systematicity items contained at least nine objects (as in Fig. 5).

In order to ensure that each object in one rug had a clear correspondence in the other, we placed "fringes" at the narrow ends of the "rugs." In addition, each pair of High Systematicity stimuli had an identical central object that was vertically symmetrical. These features were designed to give the rugs a clear top and bottom. Because the central objects in each pair were identical, they were expected to facilitate the relational alignment of the rugs. Finally, one pair of symmetric objects was cross-mapped in each pair of rugs. Cross-mappings were operationalized as perceptually similar pairs that occupied different relational roles (and hence spatial positions) in the patterns of the rugs.

As before, the most important predictions concern the contrast between the 1map and Sim → 1map conditions. For the deep, coherent relational match of the High systematicity stimuli, we should observe an elevation in relational responding for subjects who rate the similarity of the scenes before mapping over subjects who simply make one-shot mappings. However, for the impoverished relational match in the Low systematicity stimuli, the elevation in relational responding should be smaller, than for the High systematicity stimuli. This pattern of results would support the view that structural alignment generally leads to a relational focus, although this preference can be mitigated when the relational match is poor.

Method

Subjects. Subjects in this experiment were 48 undergraduates (12/condition) from Northwestern University. Subjects received $1.00 for their participation in this study.

---

[6] For the Low systematicity stimuli (in Fig. 5a), the relational correspondences were intentionally made to be less clear. For these stimuli, either the object in the same spatial position or the corresponding isometric object was counted as a relational response, although subjects selecting the relational response generally selected the object in the same relative spatial location.

Design. Factors in this experiment were Relational Systematicity (Low, High) and Mapping Condition (1map, Sim → 1map). Both factors were run between subjects.

Stimuli. Six sets of stimulus pairs were created for this experiment. Figure 5 shows a sample set. Rugs depicting shallow relations (the *Low Systematicity* figures, shown in Fig. 5a) contained four objects that were placed within an oblong box with "fringes" at the top and bottom. One object was placed along the border of the rug at the center of the top and bottom as well as at the center of the border of the rug on the right and left. The objects on the top and bottom and left and right were mirror images. One pair of objects in each pair of rugs was cross-mapped. This cross-mapping was achieved by making the top and bottom objects in one rug in a pair perceptually similar to the left and right objects in the other rug. Six sets of High Systematicity items (pictured in Fig. 5b) were also constructed by adding objects to the six Low Systematicity pairs just described. Each item in a pair contained a central object that was identical within each pair, but different between pairs. In addition, one object was placed in each corner of the rug. In two of the pairs, the corner objects were attached to the center object in some way, while in the other four pictures the corner objects were separate. The corner objects were symmetric within the objects in both adjoining corners (so that all four corner objects were the same).

Two booklets were made, each containing an instruction sheet on the first page. Each booklet contained all six pairs, with each pair of rugs placed side-by-side on a sheet of paper. In the mapping booklet, the rug on the left side of each page had an arrow pointing at one of the cross-mapped objects. In the similarity booklet, a similarity scale ranging from 1 to 9 was placed below each pair.

Procedure. Subjects participated in either the 1map or Sim → 1map condition. These conditions differed only in the order the one-shot mapping and similarity rating tasks were performed. Subjects in the 1map condition performed one-shot mapping followed by similarity rating, while subjects in the Sim → 1map condition did the reverse. Subjects in the one-shot mapping task were presented with the rug pairs with the arrows pointing to one of the cross-mapped objects. Subjects were asked to select the object in the other rug that went with that object. In the similarity judgment task, subjects were asked to rate the similarity of the rugs on the 9-point scale provided. The entire task took between 5 and 10 min.

Results

The proportion of relational responses by subjects in each Mapping and Systematicity condition are shown in Fig. 6. As predicted there was a significant main effect of Mapping condition, $F(1,44) = 4.28$, $p < .05$, reflecting that subjects in the Sim → 1map ($m = 0.45$) condition made more relational responses overall than subjects in the 1map condition ($m = 0.24$). The main effect of Systematicity was not significant, $F(1,44) = 0.47$, $p > .10$. However, this effect must be interpreted in terms of a significant interaction between Mapping condition and Systematicity, $F(1,44) = 6.85$, $p < .05$. A deeper analysis reveals that, as predicted, significantly more relational responses were made in the Similarity-first condition than in the 1map condition for the High systematicity stimuli, $t(22) = 3.46$, $p < .01$ ($m = 0.61$ (Sim → 1map), $m = 0.14$ (1map)). Surprisingly, for the Low systematicity stimuli, there were actually fewer relational responses in the Sim → 1map condition (m = 0.28) than in the 1map condition ($m = 0.33$) although this difference was not significant, $t(22) = 0.37$, $p > .10$. In fact, the level of relational responding in the

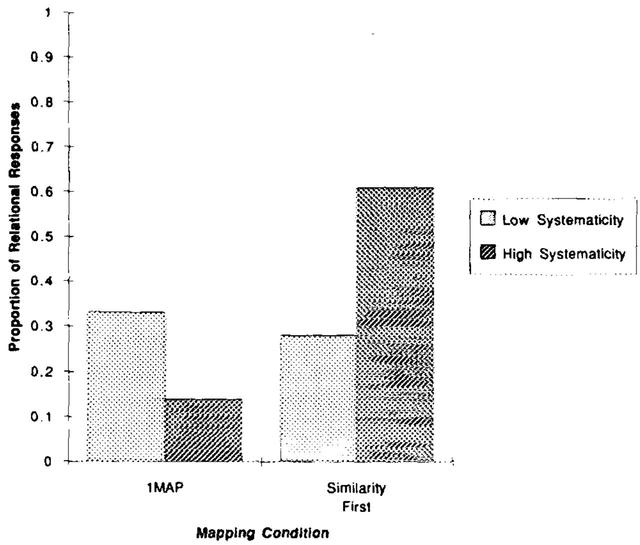

FIG. 6. Proportion of relational responses in each Mapping and Systematicity condition of Experiment 2.

1map Low systematicity condition was higher than expected, nonsignificantly exceeding that of the 1map High systematicity condition (m = 0.14), $t(22) = 1.39$, $p > .10$.

Item analyses are in the direction of the predicted interaction between Systematicity and Mapping condition. For the Low Systematicity stimuli, subjects in the Sim → 1map condition made more relational responses than subjects in the 1map condition for 2/6 (33%) of the stimuli and the same number of relational responses in both conditions for one of the items. In contrast, subjects made more relational responses in the Sim → 1map condition than in the 1map condition for all six High systematicity stimuli.

Discussion

These data provide additional support for the structural alignment view. As before, subjects who simply rated the similarity of the stimulus pairs made significantly more relational responses than subjects who performed the one-shot mapping task without comparing the items first. This relational advantage was obtained for the High Systematicity stimuli that had many isometries to constrain the relative positions of the objects. For the Low systematicity items, where the relational match was designed to provide a poor competitor to the object match, similarity comparisons did not lead to an increase in relational responding. These results also extend the relational advantage found for causal stimuli to perceptual materials. This finding suggests that structural alignment provides a plausible account for similarity comparisons of both causal and perceptual relations.

One puzzling finding was that the baseline level of relational responding was higher for the Low systematicity than the High systematicity stimuli. Although this difference was not statistically significant, we have replicated it in pilot studies with the same materials. One possibility is that the method we used to increase relational depth—namely, adding strategically placed objects to the Low systematicity stimuli—had the side effect of making the pictures appear more crowded at first glance. This perceptual complex may have made subjects in the 1map condition particularly likely to attend to the salient object similarity. Yet when subjects are asked to compare the pairs for similarity, they perform a structural alignment that highlights the common structure of the pattern and leads to a relational response. The idea that comparisons accentuate structure that is not obvious in an item in isolation has also been raised by Medin, Goldstone and Gentner (1993). They asked subjects to compare pairs of figures, one of which contained some ambiguity. Their subjects appeared to create new features for the ambiguous items in an effort to maximize the structural commonalities of each pair.

EXPERIMENT 3

Experiment 3 further tests the predictions of structural alignment highlighted by the simulation. In Experiment 2, we demonstrated a relational advantage for the High systematicity stimuli, but not for the Low systematicity stimuli. According to the structural alignment view, this result was obtained because the impoverished relational match in the Low systematicity stimuli could not compete with the rich cross-mapped object match. This explanation leaves open two possibilities. The first is the one we have suggested: that the degree of relational responding depends on the dynamic evaluation of competing interpretations. The second possibility is that the level of relational responding is controlled primarily by the degree of relational match, with good relational matches yielding many relational responses and poor relational matches yielding few relational responses. According to the first account, if the degree of relational match is kept constant, an increase in the salience of the object similarities should decrease the level of relational responding following a similarity comparison. According to the second, there should be no effect of object richness.

The present study also allows us to address a potential problem in interpreting the results of Experiment 2. The relational roles in these materials were perfectly correlated with the relative spatial locations within a global pattern. Perhaps relative spatial position has special properties not representative of relational structure in general. Indeed, it could even be argued that position does not reflect multiple interconnecting relations, but is simply an attribute of a design. To remove this concern, the present experiment uses a different set of perceptual relations than the ones used in Experiment 2.

The stimuli consist of arrays of three objects that vary in height or color saturation. The arrays depict either the *monotonic increase* relation or the *symmetry* relation (Gentner, Markman, Rattermann, & Kotovsky, 1990). Monotonic increase is operationalized as three objects increasing along some ordered dimension such as size or color saturation from left to right, or from right to left (i.e., 134 or 542). Symmetry is operationalized as a central object flanked by identical objects (i.e., 131 or 323). Sample stimulus pairs depicting the symmetry and monotonic increase relations are shown in Fig. 7. The relational roles in these items are determined by the relative height or darkness of a particular object in the array, rather than its spatial location within some design. One set of these relations is constructed from sparse geometric forms like circles and squares. For these stimuli, we expect that more relational responses will be made in the similarity-first condition than in the simple one-shot mapping task. A second set is constructed from rich objects like palm trees and globes. This increase in the matching attributes of the cross-mapped objects is expected to increase subjects' preference for the object match in the similarity-first condition.

Method

Subjects. Subjects were 50 undergraduates at the University of Illinois, who received course credit in introductory psychology for their participation. Two subjects were eliminated from the study for failing to follow instructions, leaving a total of forty-eight (eight/condition).

Design. There are three Mapping conditions (1map, Sim → 1map and 3map) and two levels of Richness (Sparse and Rich). Both factors were between-subjects.

Stimuli. The stimuli were 16 pairs of pictures of perceptual relations, like those shown in Fig. 7. Symmetry relations were operationalized as an array of objects containing a central object flanked by identical objects. We will use the following naming convention: if the outer objects are smaller or lighter than the central object, the relation is of *positive polarity*. If the outer objects are larger or darker than the central object, then the relation is of *negative polarity*. A pair of stimuli displaying the symmetry relation can have the same polarity (i.e., positive/positive or negative/negative) or opposite polarity (i.e., positive/negative or negative/positive).

Stimuli displaying the monotonic increase relation were similarly organized. The relation was operationalized as a series of three objects increasing, either in height or in color saturation, from left to right or right to left. The relation is of *positive polarity* if the increase takes place from left to right and of *negative polarity* if the increase takes place from right to left. As for pairs displaying symmetry, pairs displaying the monotonic increase relation can be of the same polarity or of opposite polarity.

Half the pairs depicted symmetry relations and half depicted monotonic increase rela-

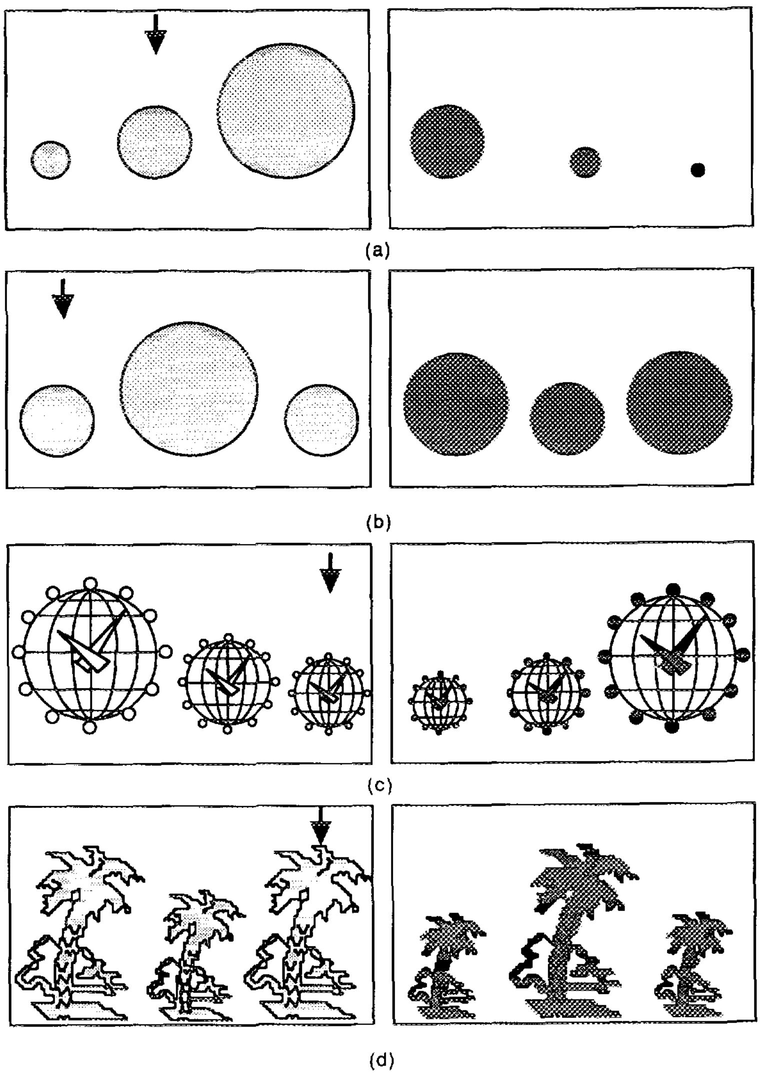

FIG. 7. Sample monotonic increase (a and c) and symmetry stimuli (b and d), like those used in Experiment 4. Sample Sparse (a and b) and Rich (c and d) stimuli.

tions. Half the pairs were of the same polarity and half were of opposite polarity. Finally, in half the pairs, the objects varied in size, while in the other half, the objects varied in saturation. These three presentation factors were fully counterbalanced. There were two examples of each combination of the three factors in the set of 16 pairs. A summary of the structure of the stimulus set is presented in Table 2.

The stimuli were presented on a Macintosh II color computer screen. In the sparse condition, the stimuli consisted of either circles or squares displaying either symmetry or

TABLE 2
Structure of Geometric Stimuli Used in Experiment 3

Stimulis	Stimulus dimensions		
	Relation	Polarity	Varies in
1	Monotonicity	Same	Color
2	Monotonicity	Same	Size
3	Monotonicity	Same	Color
4	Monotonicity	Same	Size
5	Monotonicity	Opposite	Color
6	Monotonicity	Opposite	Size
7	Monotonicity	Opposite	Color
8	Monotonicity	Opposite	Size
9	Symmetry	Same	Color
10	Symmetry	Same	Size
11	Symmetry	Same	Color
12	Symmetry	Same	Size
13	Symmetry	Opposite	Color
14	Symmetry	Opposite	Size
15	Symmetry	Opposite	Color
16	Symmetry	Opposite	Size

monotonic increase. Six easily discriminable shades of grey were generated by using equal levels of red, green and blue at 15, 30, 45, 60, 75, and 100% saturation. For size variation, five easily discriminable object sizes were used, (radii of 3, 6, 12, 24, and 48 pixels).

The Rich objects were constructed similarly, except that the objects (house, globe, investigator, scale, palm-tree and light bulb) were taken from a public-domain set of clip art for the Macintosh Computer. Five distinct sizes of these objects were made. In addition, six saturation levels of six hues were used (black, red, green, blue, turquoise, and yellow). Hue was arbitrarily paired with shape between items. Within a given array, all items were the same shape and hue. In addition, within an array, all items were identical except for the dimension along which the symmetry or monotonic increase relation occurred. Within a stimulus pair, all items were of the same shape and hue.

Each pair of scenes contained a cross-mapping, operationalized as a pair of perceptually similar objects in the two scenes that played different roles in the relational structure in each scene. In order to create this perceptual similarity, the objects were identical along the dimension varying in the relation size or saturation), but slightly different along the dimension irrelevant to the relational structure.

Procedure. Subjects were run one at a time using a Macintosh II computer. The program for running the experiment was written in the cT language (Sherwood & Sherwood, 1988). In the mapping portion of the study, subjects were told that they would see two scenes. An arrow would appear over one of the objects in the top scene and they were to select the object that best "goes with" that object by moving a cursor on the screen with the mouse and clicking on the preferred object. In order to get subjects comfortable with this use of the mouse to choose objects, they were given practice trials in which scene would appear on the screen, and they were told to select either the right, middle or left object.

On the computer, the basic one-shot mapping task took the following form. At the beginning of each trial, a small (20 pixel × 10 pixel) box appeared at the top center of the screen. Subjects were told to move the cursor inside of the box and click. This was done to ensure that the cursor was always centered at the start of a trial. When the subject clicked in the box, it disappeared and a fixation point appeared at the center of the screen for 500 ms. Then the two scenes appeared, presented one above the other. An arrow pointed at the cross-mapped object in the top scene. The trial ended when the subject selected one of the objects in the bottom scene. The subject's choice was recorded. If subjects clicked on an object in the top scene or outside of the scene boxes they were told to select again.

The three-mappings task was identical to the one-shot mapping task, except that subjects were told initially that they would be asked to perform a mapping for all three objects in a scene. After the first mapping (for the cross-mapped object), the first arrow disappeared and a second arrow was placed over a different object in the top scene. When the subject selected another object in the bottom scene, the second arrow disappeared and a third arrow was drawn over the last object and subjects made a third mapping. All three responses were recorded.

In the similarity part of the experiment, subjects were told to rate the similarity of the pairs of scenes they would see on a scale from 1 (low) to 9 (high). The pairs were presented one at a time and the subjects were asked to type in their rating from the computer keyboard. Subjects' similarity ratings for each pair were recorded.

Subjects in the 1map condition performed the one-shot mapping task followed by the similarity ratings task. Subjects in the Sim → 1map condition performed a block of similarity ratings before performing the one-shot mapping task. Finally, subjects in the 3map condition performed the three-mappings task followed by the similarity ratings task.

Results

The proportion of relational responses made by subjects in each Mapping and Richness condition is shown in Fig. 8. As predicted, there was a significant main effect of Mapping condition $F(2,42) = 6.40, p < .005$. The main effect of Richness was not significant $F(1,42) = 1.11, p > .10$. Also as predicted, there was a significant Richness × Mapping interaction $F(2,42) = 4.07, p < .05$. This interaction reflects that more relational responses were made to Sparse items in the Sim → 1map condition (m =

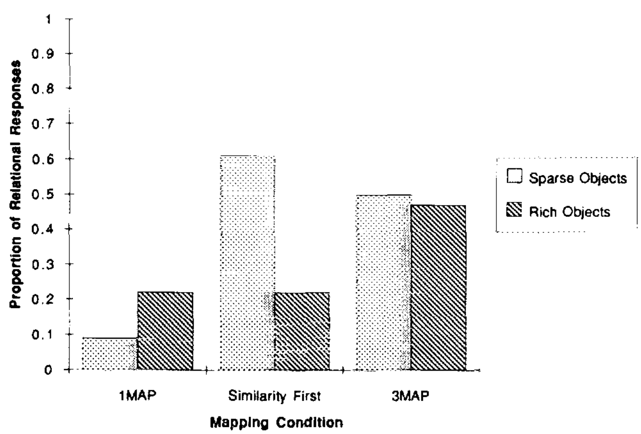

FIG. 8. Proportion of relational responses in each Mapping and Richness condition of Experiment 3.

0.61) than in the 1map condition (m = 0.09), $t(14) = 3.44, p < .01$, but the same number of relational responses was made to the Rich objects in the Sim → 1map condition (m = 0.22) and 1map condition (m = 0.22). Thus, as predicted by structural alignment, the presence of very rich objects decreased subjects' propensity to select the relational mapping.

An item analysis provides corroboration for the observed interaction. More relational responses were made in the Sim → 1map condition than in the 1map condition for all sixteen Sparse stimuli (100%), $p < .001$ by sign test. However, only 5/16 (31%) of the Rich items received more relational responses in the Sim → 1map condition than in the 1map condition, $p > .10$ by sign test. In addition, more relational responses were made in the 3map condition than in the 1map condition for all sixteen items (100%) for both the Sparse and Rich stimuli, $p < .001$ by sign test.

Discussion

The results of this study supported the predictions of the structural alignment view. Specifically, similarity comparisons led to an increase in relational responding for the sparse stimuli. Thus, when a simple object match was placed in opposition to a simple relational match, the relational match was preferred. In contrast, when very rich object similarities were placed in opposition to a simple relational structure, the object match was preferred. Thus, as predicted, when the object and relational similarity were placed in opposition, increases in the number of matching attributes decreased subjects' preference for the relational match. This study was the complement of Experiment 2 which demonstrated that decreases in the extent of the matching relational structure also decreased subjects' preference for the relational match.

Perceptual relations and relative position. One of the motivations for the use of the monotonic increase and symmetry relations in this study was that relative spatial location and relational role were perfectly correlated in Experiment 2. The stimuli in this experiment allow an explicit test of the hypothesis that subjects in the Sim → 1map condition were responding solely on the basis of relative spatial position. In these materials, relative spatial location and relational structure are negatively correlated for the monotonic increase relations of opposite polarity (items 5 through 8). Thus, if subjects select the relational mapping for these stimuli, we can reject the possibility that their mappings simply reflect a preference for relative spatial position.

For these stimuli, subjects made a greater proportion of relational mappings in the similarity first condition (m = 0.50) than in the 1map condition (m = 0.09). Furthermore, the proportion of relational responses in the Sim → 1map condition for these items (m = 0.50) is comparable to the proportion of relational responses to the rest of the stimuli in the similarity-first condition ($m = 0.65$), $t(14) = 1.13, p > .10$. A regression analysis was also performed on these data. We attempted to predict the proportion of relational responses from a dummy variable with the value 1 for stimuli for which relative position lead to a relational match for -1 for stimuli for which relative position and relational match were negatively correlated. The regression coefficient in this analysis was not significant (standardized coefficient = 0.29), $t(14) = 1.13, p > .10$. Finally, subjects who did not make a relational response tended to map based on object similarities, not positional similarities. These results support the claim that subjects were mapping on the basis of the matching relational structure, not spatial position.

Pairwise similarity and relational similarity. We have argued that the elevation in relational responding following similarity comparisons in the sparse condition resulted from structural alignment. However, an alternate explanation is that subjects simply attempted to minimize the pairwise dissimilarity between objects placed in correspondence, while still respecting the one-to-one mapping constraint. On this account, relational alignment is irrelevant, even though the results appeared to support relational alignment. For example, if the objects in a pair of monotonic-increase stimuli are placed in the same correspondences as are dictated by aligning the relational structure, the pairwise similarities are maximized. Each pair of objects differs slightly, but any other way of aligning the scenes would place two highly dissimilar objects in correspondence.

Fortunately, this possibility can be evaluated by examining symmetry relations of opposite polarity (items 13 through 16), where the relational mapping requires dissimilar objects to be placed in correspondence. Examination of subjects' performance on these items indicates that their level of relational responding ($m = 0.53$) is about the same as the level observed for all other stimuli ($m = 0.61$), $t(14) = 0.75, p > .10$. A similar regression analysis was performed for this analysis, with a dummy variable coded 1 for stimuli where global similarity and relational similarity were correlated, and -1 when they were not. Once again, the coefficient was not significant (standardized coefficient = 0.20), $t(14) = 0.75$, $p > .10$. Furthermore, more relational responses were made to all four symmetry relations of opposite polarity in the Sim → 1map condition than in the 1map condition. Thus, it appears that the increase in relational responding in the similarity-first condition for Sparse stimuli was not due merely to minimizing overall object dissimilarity.

Analysis of presentation factors. These perceptual stimuli allowed us to control the type of relation, polarity and dimension along which the items differed (see Table 2). Analysis of the Sim → 1map condition with sparse stimuli allows us to assess these factors in a task which we assume promotes relational mapping. The type of relation did not appear to matter:

subjects made the same proportion of relational responses to stimuli depicting the monotonic increase and symmetry relations (m = 0.61 for both relations). Consistent with our intuitions, there was a nonsignificant tendency for subjects to make more relational responses to stimuli of the same polarity (m = 0.67) than to stimuli of opposite polarity (m = 0.54), $t(14) = 1.71, p > .10$. Kotovsky and Gentner (1990) found the same result for 4-, 6-, and 8-year-olds. Furthermore, Goldstone and Gentner in an unpublished pilot study also found an advantage for same polarity stimuli. There was also a nonsignificant tendency for subjects to make more relational responses to stimuli varying in size (m = 0.70) than to stimuli varying in color (m = 0.51), $t(14) = 1.74, p > .10$. This result meshes with findings by Smith and Sera (1992) that the size dimension has clear "more" and "less" directions for adults, while the darkness dimension does not.[7]

GENERAL DISCUSSION

We have proposed that similarity comparisons are carried out by a process of structural alignment. According to this hypothesis, when object similarity and relational similarity are placed in opposition, as they were in the one-shot mapping task, objects are generally placed in correspondence based on their roles within the matching relational structure. However, the structural alignment view suggests two factors that will decrease the strength of this relational focus: decreasing the depth and coherence of the relational match and increasing the number of matching attributes of the cross-mapped objects.

The experiments presented here tested these predictions explicitly. In the one-shot mapping task, subjects were presented with stimuli containing cross-mappings and were asked to select the object in one scene that went with the cross-mapped object in the other. All of these stimuli were explicitly designed so that subjects' natural tendency was to select the similar object. However, as summarized in Table 3, when subjects rated the similarity of the scenes prior to performing the one-shot mapping, they often responded relationally. In particular, they made many more relational responses than subjects who simply performed one-shot mappings without prior similarity comparisons. This relational advantage (over the four cases where it is predicted) ranged from a gain of 47% in Experiment 1b to a gain of 578% for the sparse stimuli in Experiment 3.

However, a fundamental claim is that structural alignment is a process whereby relational focus arises as competing global interpretations emerge from local correspondences. In these experiments, the competi-

[7] Goldstone, Gentner and Medin (1989) obtained a similar result.

TABLE 3
Summary of Results across Experiments

Experiment	Degree of Spontaneous Alignment[a]	Alignment in Sim → Imap Condition[b]	Percentage Gain from Imap to Sim → Imap
1	0.42	0.69	64
1b	0.47	0.69	47
2:Low Syst (c)	*0.33*	*0.28*	*−1*
2: High Syst	0.14	0.61	335
3: Sparse	0.09	0.61	578
3: Rich	*0.22*	*0.22*	*0*

[a] Level of relational responses in Imap condition.
[b] Number of relational responses in Sim → Imap condition.
[c] Italics indicate conditions not predicted to show a relational advantage.

tion was between an interpretation based on object similarities and an interpretation based on relational similarities. In Experiments 2 and 3, some of the stimuli were explicitly designed to reverse the normal advantage for relations (these conditions are shown in italics in Table 3). In the Low systematicity condition of Experiment 2, the relational match was designed to be poor, thereby decreasing the strength of the relational match relative to the object match. Conversely, in the Rich condition of Experiment 3, extremely rich objects were used to construct simple relations, thereby increasing the strength of the object match relative to the relational match. In both of these cases, no relational advantage was observed. Taken together, these results suggest that, while there may be other processes that can arrive at a relational match or an attribute match for a given pair, general similarity comparisons involve structural alignment.

Relationship between Rated Similarity and Relationality

One often replicated finding is that stimulus pairs with primarily relational commonalities are often judged to be more similar than stimulus pairs with primarily object commonalities (Gentner, Rattermann, & Forbus, in press; Goldstone, Medin, & Gentner, 1991; Rattermann & Gentner, 1987; Schumacher & Gentner, 1986). For example, in a forced-choice task, Goldstone et al. (1991) found that subjects were much more likely to select a figure that was relationally similar to a target than a figure that had attribute similarities with a target. Similarly, Gentner and Clement (1988) found that subjects rated metaphors as more apt when they produced relational interpretations than when they produced attribute interpretations. We examined our data for corroboration of this phenomenon.

First, we calculated the correlation between the number of relational responses and the mean rated similarity of each item. On the basis of prior research, we expected that similarity ratings would be positively correlated with the number of relational responses given to an item. We anticipated that this finding would be most pronounced in the Sim → 1map condition since it is there that subjects' similarity judgments should most influence the mapping task. The correlation between rated similarity and number of relational responses was calculated for all conditions in all experiments. These data are presented in Table 4. As expected, these correlations were positive in 11 of 14 conditions, although only 3 of 13 reached significance. Furthermore, as expected, all six correlations in the Sim → 1map condition were positive. In general, subjects gave higher similarity ratings to items when they placed the objects in correspondence based on the relational matches than when they mapped based on attribute matches.

The relationship between similarity and relationality may be examined more directly by separating items given relational responses from items given object responses and finding the mean rated similarity given to the items in each group. This analysis, presented in Table 5, has the advantage that it is less sensitive than correlations to the small number of stimuli used in some of the studies. As anticipated, in 12/14 (86%) conditions, subjects gave higher mean similarity ratings to items to which they gave relational responses than to items to which they gave object responses. Seven out of 14 of these differences were statistically significant, including five of the six similarity-first conditions. Interestingly, the only case where the difference between rated similarity following object and relational responses in the similarity-first condition was not significant was for the Low systematicity rugs in Experiment 2, which were designed to have a poor relational match. This overall pattern of similarity

TABLE 4
Correlation between Mean Rated Similarity and Number of Relational Responses for All Conditions of All Experiments

		Mapping Condition		
Experiment	*df*	1map	Sim → 1map	3map
1	6	−0.09	0.51	0.54
1b	6	xx	0.47	xx
2: Low syst	4	−0.48	0.38	xx
2: High syst	4	0.08	0.79*	xx
3: Sparse	14	−0.15	0.60*	0.46*
3: Rich	14	0.13	0.22	0.22

Note. "xx" Denotes conditions not run.
* $p < .05$, one-tailed.

TABLE 5
Mean Similarities Given Object Responses and Relational Responses

	Mapping Condition					
	1map		Sim → 1map		3map	
Experiment	Object Response	Relational Response	Object Response	Relational Response	Object Response	Relational Response
1	5.13	4.98	3.90	5.15*	4.44	5.88*
1b	xx	xx	4.58	5.14*	xx	xx
2: Low syst	5.08	5.79	5.23	5.80	xx	xx
2: High syst	4.26	5.00	3.50	4.21*	xx	xx
3: Sparse	4.60	5.38	4.00	5.04*	4.89	6.09*
3: Rich	5.01	5.58	5.48	6.39*	5.91	5.86

Note. "xx" Denotes conditions not run.
* $p < .05$, Independent samples t test, 1 tailed.

ratings mirrors the findings described above that subjects find relational matches more appealing than attribute matches.

Computational Models of Similarity

The SME simulation was used to delineate the predictions of structural alignment, and to suggest a particular process by which structural comparisons may be carried out. In the simulation, initial local matches coalesce into global matches. When there are competing interpretations (as with the object-based and relation-based interpretations in these studies), the interpretation with the highest evaluation is selected. This choice depends on the sheer number of matching predicates in each interpretation (attributes, functions or relations), but also depends strongly on the depth and connectivity of the matching structure. All else being equal, relational matches are preferred over attribute matches because of their greater depth.[8]

Of course, the competition between object similarity and relational similarity that we created here represents an extreme case. In most natural settings, object similarity and relational similarity are correlated, so that object and relational matches can be part of the same interpretation. Nonetheless, the phenomenon of competition among potential interpretations is quite general. Even in literal similarity comparisons there are typically multiple competing interpretations that can be built from the same set of local matches. For example, when comparing Chicago and

[8] Other factors may help determine the preferred match. For example, salience of object similarities, factual correctness of inferences arising from the match and relevance of the match to the current task may all be important.

New York, we could focus on commerce in which case Chicago's Commodities Exchange would correspond to New York's World Trade Center. However, we could also focus on tourist spots, in which case Chicago's Sears Tower would correspond to New York's World Trade Center. Thus, the processing principles used to select between competing matches in similarity comparisons have widespread application.

Taking SME as a possible cognitive model of the comparison process, let us compare how well its predictions are borne out relative to other computational models of comparison. SME differs from other competing models in its adherence to structural consistency—one-to-one mapping and connectivity—as strict constraints on mapping. Holyoak and Thagard's (1989) ACME loosens the one-to-one mapping constraint so that it is only a pressure on alignment. For this reason, cross-mappings pose a particular difficulty for this model. As Holyoak and Thagard discuss, when ACME is given analogies with potential cross-mappings, it often settles on matches that are not one-to-one. The object matches in these interpretations are based on both object similarity and relational similarity. However, the increase in relational responding following similarity comparisons in these studies (even in the presence of tempting cross-mappings) is evidence for the psychological importance of the one-to-one mapping constraint. Further evidence that subjects prefer structurally consistent interpretations comes from the finding that rated similarity was higher when subjects chose relational correspondences than when they chose object correspondences. If subjects did not have a preference for structural consistency, they could have chosen the tempting object match even if they saw the relational alignment.

The results of Experiment 3 also pose problems for Bakker and Halford's (1988) MATCHMAKER program. According to their model, when two representations are compared, objects are first placed in correspondence if they play the same role in a matching relational structure. Then remaining objects are aligned based on object similarities. However, subjects in Experiment 3 made many object mappings, even in the presence of a good relational match. This finding suggests that object and relational similarities simultaneously determine the best match, rather than being applied sequentially.

SME appears to provide a plausible process model for human similarity comparisons, but there are some areas where the model should be modified. For example, early versions of SME generated all possible interpretations. Clearly, this approach is not psychologically plausible. Forbus and Oblinger (1990) have developed an algorithm to allow SME to generate a single match. The algorithm has been tested computationally on 66 analogies, and 94% of the time it produced the highest rated GMAP produced by SME when it generated all matches exhaustively. Further work must be done to test the detailed predictions of this algorithm for human performance, but so far it appears promising.

Other process models of analogical mapping do not generate all possible matches. For example, Holyoak and Thagard's (1989) ACME and Goldstone and Medin's (in press) SIAM, generate a single best interpretation. While this approach seems more reasonable than an exhaustive search, the data from our studies suggest that subjects consider both an object match and a relational match. Thus, more work must be done to determine how a small number of plausible interpretations can be generated. It remains an open question as to how many different interpretations compete during similarity comparisons.

Finally, the output of SME is deterministic, while responses by subjects are variable. We think the best explanation for this variability is that subjects vary in their encoding of the stimuli. This hypothesis is consistent with Barsalou's (1989) finding that the same person will retrieve different information about a concept at different times. Another explanation for this variability is that competing interpretations may be given some probability of being used based on the evaluation score they receive. A third possibility is that some resource, such as attention or effort, varies nondeterministically and affects the mapping outcome. For example, relational matches may be difficult to determine, and thus may be more likely to be calculated when attentional resources are high than when there is competition for attentional resources. Hofstadter and Mitchell (1988, in preparation) have addressed variability in analogy directly. Their COPYCAT system places constraints on the kinds of allowable correspondences and then searches for analogies using a *parallel terraced scan* which tries many possible matches simultaneously, searching each match at a depth roughly corresponding to its promise. Further research might examine the psychological plausibility of such dynamic allocation of attention.

CONCLUSIONS

The mechanism that determines psychological similarity is a natural and seemingly effortless process that can operate across a wide range of stimulus types. Among the variety of information used by alignment are object attributes, relations between objects, general domain theories and current context. This flexibility is both a blessing and a curse. On the positive side, similarity can be included as a basic component in a myriad of other cognitive processes. On the negative side, similarity becomes difficult to characterize, leaving some, like Nelson Goodman, to consider it "a pretender, an imposter, [and] a quack." (Goodman, 1971, p. 437).

The structural alignment view allows us to integrate a range of findings into a single model. By considering similarity as a process sensitive to a variety of inputs, many conflicts can be resolved (see Medin, Goldstone, & Gentner, 1993). The competition between different global interpretations means that the same process model predicts relation-based interpretations in some instances and object-based interpretations in others. Moreover, this determination depends not only on the relative numbers and individual saliences of relations and object attributes but also on the structure of the matching system. The structural alignment view offers a theoretical basis for examining the calculation of similarity using structured representations.

REFERENCES

Bakker, P. E., & Halford, G. S. (1988). *A basic computational theory of structure-mapping in analogy and transitive inference* (Tech. Report No. 88/1). Centre for Human Information Processing and Problem Solving.

Barsalou, L. W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. New York: Cambridge Univ. Press.

Biederman, I. (1987). Recognition-by-components: A theory of human-image understanding. *Psychological Review*, 94(2), 115–147.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1987). The structure mapping engine: Algorithm and examples (Tech. Report No. UIUCDCS-R-87-1361). University of Illinois.

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

Forbus, K. D., & Gentner, D. (1989). Structural evaluation of analogies: What counts? *Proceedings of the 11th Annual Conference of the Cognitive Science Society* (pp. 341–348). Ann Arbor, MI.

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the 12th Annual Conference of the Cognitive Science Society* (p. 61–68). Boston, MA.

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155–170.

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

Gentner, D., & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation, Volume 22*. Academic Press, Inc.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. London: Cambridge University Press.

Gentner, D., Markman, A. B., Rattermann, M. J., & Kotovsky, L. (1990). Similarity is like analogy. *Proceedings of the Second Annual Meeting of the Midwestern Artificial Intelligence and Cognitive Science Society*. IL: Carbondale.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (in press). The roles of similarity in transfer: Determinants of similarity-based reminding and inference. *Cognitive Psychology*.

Goldmeier, E. (1937/1972). Similarity in visually perceived forms [Monograph 29]. *Psychological Issues*, 8(1).

Goldstone, R. L., Gentner, D., & Medin, D. L. (1989). Relations relating relations. *Proceedings of the 11th Annual Meeting of the Cognitive Science Society*. MI: Ann Arbor.

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.

Goldstone, R. L., & Medin, D. L. (in press). Similarity, interactive activation and mapping. In K. J. Holyoak & J. A. Barnden (Eds.), *Symbolic and connectionist approaches to analogy, Vol. 2: Analogical connections*. Norwood, NJ: Ablex.

Goodman, N. (1972). *Problems and Projects*. Indianapolis: Bobbs-Merrill.

Greiner, R. (1988). Learning by understanding analogies. *Artificial Intelligence*, **35**, 81–125.

Hall, R. P. (1989). Computational approaches to analogical reasoning: A comparative analysis. *Artificial Intelligence*, **39**, 29–120.

Hinton, G. E. (1979). Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science*, **3**, 231–250.

Hofstadter, D. R., & Mitchell, M. (1988). Concepts, analogies, and creativity. *Proceedings of the Canadian Artificial Intelligence Conference*.

Hofstadter, D. R., & Mitchell, M. (in press). An overview of the copycat project. In K. J. Holyoak & J. A. Barnden (Eds.), *Symbolic and connectionist approaches to analogy, Vol. 2: Analogical connections*. Norwood, NJ: Ablex.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory and Cognition*, **15**(4), 332–340.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, **13**, 295–355.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, **93**(2), 136–153.

Keane, M. T. (1990). Incremental analogizing: Theory and model. In Gilhooly, Keane, Logie, & Erdos (Eds.), *Lines of thinking, Vol. 1*. London: Wiley.

Kedar-Cabelli, S. (1985). Analogy: From a unified perspective (Technical Report ML-TR-3). Rutgers University, Department of Computer Science, Laboratory for Computer Science Research.

Kotovsky, L., & Gentner, D. (1990). Pack light: You will go farther. *Proceedings of the Second Annual Meeting of the Midwestern Artificial Intelligence and Cognitive Science Society*. IL: Carbondale.

Kubovy, M. (in preparation). Symmetry as similarity: The phenomenology of repetition in one- and two-dimensional periodic patterns.

Lockhead, G. R., & King, M. C. (1977). Classifying integral stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, **3**(3), 436–443.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification. *Psychological Review*, **85**(3), 207–238.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, **100**(2), 254–278.

Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization Relationship. *Journal of Experimental Psychology: General*, **115**(1), 39–57.

Novick, L. R. (1988). Analogical transfer, problem similarity and expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **14**(3), 510–520.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, **9**, 441–474.

Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence*, **46**(1–2), 77–106.

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance*, **3**(3), 422–435.

Rattermann, M. J., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. *Proceedings of the 9th Annual Conference of the Cognitive Science Society* (pp. 23–25). Seattle, WA.

Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology*, 7, 532–547.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13(4), 629–639.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier examples. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15(3), 456–468.

Schumacher, R., & Gentner, D. (1987). *Similarity-based remindings: The effects of similarity and interitem difference*. Paper presented at the Annual Meeting of the Midwestern Psychological Association, Chicago, IL.

Sherwood, B. A., & Sherwood, J. N. (1988). The cT language and its uses: A modern programming tool (Report No. 88-28). Pittsburgh, PA: Center for Design of Educational Computing Technical, Carnegie Mellon University.

Singley, M. K., & Anderson, J. R. (1989). *The Transfer of Cognitive Skill*. Cambridge, MA: Harvard Univ. Press.

Smith, L. B., & Sera, M. D. (1992). A developmental analysis of the polar structure of dimensions. *Cognitive Psychology*, 24(1), 99–142.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1–2), 159–216.

Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review*, 8, 247–261.

Treisman, A., & Paterson, R. (1984). Emergent features, attention and object perception. *Journal of Experimental Psychology: Human Perception and Performance*, 10(1): 12–31.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327–352.

Tye, M. (1991). *The Imagery Debate*. Cambridge, MA: The MIT Press.

Ullman, S. (1984). Visual routines. *Cognition*, 18, 97–159.

Winston, P. H. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence*, 19, 321–350.

(Accepted November 25, 1992)