

6

Similarity, Interactive Activation, and Mapping: * An Overview

*Robert L. Goldstone
Douglas L. Medin*

Goldstone, R.L., & Medin, D.L. (1994). Interactive Activation, Similarity, and Mapping. In K. Holyoak and J. Barnden (Eds.) *Advances in Connectionist and Neural Computation Theory, Vol. 2: Analogical Connections*, (pp. 321-362). Ablex : New Jersey.

1. INTRODUCTION

The act of comparison often seems to be an immediate, direct operation. Dogs and wolves appear similar simply because of the large perceptual overlap between their visual forms—they both have a head with a snout and ears, four legs, and a tail. In general, things seem similar if they share many properties. Abstract properties may also influence similarity; puppies and children seem similar because of their innocence, youth, and dependence on others. Once we find the appropriate set of property descriptions, so the argument goes, similarity assessment is a direct function of the objects' overlap/proximity on these descriptions. As such, the first step is to create representations of the to-be-compared objects in terms of their properties. Once the two property lists have been created, the similarity computation proceeds by comparing the two property lists for matching and mismatching features.

However, there is more to similarity than property listing and matching. Comparing scenes and objects with parts requires a more structured representation than a feature list, and requires a more sophisticated process than counting up matches and mismatches. Features are organized into objects; objects are organized into relations; and relations are organized into scenes. The parts of a

Dedre Gentner has played a central role in all stages of this research. Frances Kuo provided many useful comments and art work. This chapter has benefited greatly from the constructive suggestions of Jay McClelland, Keith Holyoak, Arthur Markman, Melanie Mitchell, Colleen Seifert, Ed Smith, Linda Smith, Keith Smith, Steve Sloman, and the entire Medin Lab group.

This research was supported by National Science Foundation, Grant BNS-88-12193, awarded to the second author.

scene are mapped onto the parts of the scene with which it is compared, and this process of finding corresponding parts has an important influence on the perceived similarity. The purposes of this chapter are to (a) demonstrate the process of mapping in human scene comparison, (b) present new experimental findings implicating mapping in similarity assessment, and (c) organize and interpret these results with an interactive activation model of mapping and similarity.

The organization of this chapter is as follows. First, we review the role of mapping and global consistency in both low-level visual perception and abstract analogy and then suggest that mapping and consistency also apply to similarity assessment. Next, we review current models of similarity and note that they have little to say about processes by which corresponding properties are aligned. We then describe some experiments on alignment processes associated with comparisons. We account for these results with an interactive activation model of alignment and contrast this model with a number of alternatives. Finally, we assess the role of mapping or alignment in comparisons more generally and offer some conclusions.

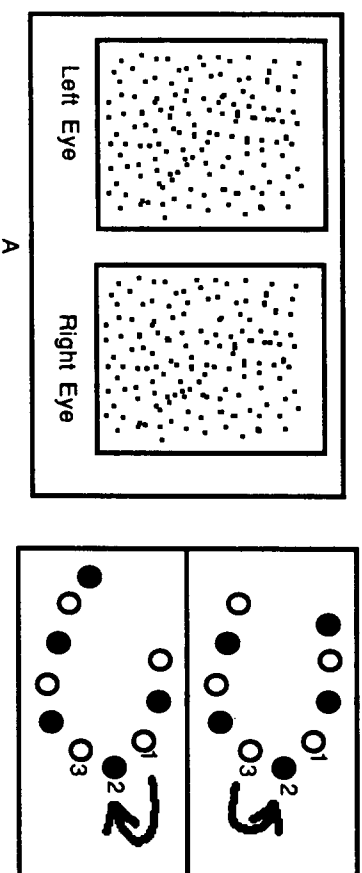
2. MAPPING AND STRUCTURE IN PERCEPTUAL AND CONCEPTUAL COMPARISONS

Similarity inhabits the broad middle ground between low-level visual perception and abstract analogy. The perception of similarity depends both on physical, concrete properties, and on knowledge-based, abstract properties. In both perceptual and conceptual comparisons, the act of mapping elements, of finding the correspondences between scene parts, plays a fundamental role.

2.1. Perceptual Mapping Processes in Comparison

In perceiving objects in depth, people combine information obtained from their two eyes. The image locations from the left eye must be placed in correspondence with the image locations from the right eye. Marr and his colleagues (Marr, 1982; Marr & Poggio, 1979) have investigated algorithms that compute the depth of an object by: (a) selecting locations from the image in one eye, (b) identifying the same location in the other eye's image, and (c) measuring the disparity between the corresponding image points. The task of identifying corresponding locations is difficult because of the "false target problem"—the problem of finding the correct location-to-location correspondences given the large number of potential mappings between the dots in the two images. Thus, in Figure 6.1A, the image that falls on the left eye is integrated with the image that falls on the right eye by creating correspondences between the images' dots.

Figure 6.1. The importance of aligning scene parts in visual perception. In A, the dots that fall on the left eye must be fused with the dots that fall on the right eye in order to form an integrated and coherent perception. In B, the white dots belonging to the first frame of an apparent motion display must be placed into correspondence with the black dots belonging to the second frame. In the top display, Dot 2 is seen as moving to become Dot 3; in the bottom display, Dot 2 maps onto Dot 1.



These mappings are formed in random-dot stereograms by constraining correspondences such that (a) black dots can only match black dots (the similarity constraint), (b) one black dot matches no more than one black dot (the uniqueness constraint), and (c) the distance between matching dots usually varies gradually (the continuity constraint). While any dot in a scene can potentially match any number of dots in the other scene, one dot's correspondences are strongly constrained by the other correspondences that are established by other scene parts. This dependence of location mappings on other location mappings is a hallmark of "cooperative" algorithms.

Cooperative algorithms also are important in apparent motion phenomena. Single frame visual displays that are presented in fairly rapid succession can yield strong subjective impressions of motion (a fact the movie industry relies on). Reviews of this phenomena are found in Kolers (1972), Ramachandran and Anits (1986), and Ullman (1979). For subjective motion to occur, people must create correspondences between the separate image frames. In Figure 6.1B, two alternating displays of two frames each are depicted. The white circles are all

displayed in one frame. The black circles appear in the second frame. (They are shaded black in the figure to show that they appear in the *second* frame.) The two frames are then alternated with each other on a computer screen once every quarter second. Subjectively, for the top apparent motion display, five dots are seen as rocking on a sideways horseshoe, with Dot 3 of the first frame corresponding to Dot 2 of the second frame. In the bottom display, the five dots subjectively move such that Dot 1 becomes Dot 2.

This small example suggests one method the mind uses to constrain motion perception—global optimization of correspondences. There is only one difference between the top and bottom display: in the top display the leftmost black dot is on the upper portion of the horseshoe, while in the bottom display the leftmost black dot is on the lower portion of the horseshoe. Although far removed from Dots 1, 2, and 3, the location of the leftmost black dot constrains the mappings of these dots. Mappings are created such that each white dot has a strong tendency to map onto one and only one black dot. Ullman (1979) and Marr (1982) argue that the subjectively perceived motion will be that motion which “maximizes the overall similarity between the frames” (Marr, 1982, p. 186). Interestingly, maximizing the overall similarity can proceed solely on the basis of local interactions between mappings.

In both the perception of depth and motion there is evidence that mappings between scenes are constructed—mappings between the left and right visual images or mappings between frames that are separated in time. In both instances, the mappings that develop are partially constrained by local affinities and by global consistency. Mappings are sought out that: (a) place similar parts in correspondence (black dots tend to map onto other black dots), and (b) place parts in correspondence that are consistent with the other correspondences that have been established.

2.2. Analogical Mapping Processes in Comparison

Establishing mappings also plays a critical role in more conceptual processes. Analogies (Gentner & Clement, 1988) are understood by creating correspondences between elements from two domains. The comprehension of the atom/solar system analogy requires setting up correspondences between the atom’s nucleus and the sun, between electrons and planets, etc.

Reminiscent of Marr’s and Ullman’s proposals, in Holyoak and Thagard’s ACME system and Gentner’s Structure Mapping Theory (SMT), comparison processes serve to: (a) place similar relations in correspondence, and (b) place relations in correspondence that are consistent with other relational correspondences. According to Gentner’s systematicity principle (Gentner, 1983) and Holyoak and Thagard’s “uniqueness” and “relational consistency” constraints (Holyoak & Thagard, 1989), elements are mapped onto each other so as to tend to yield coherent relational correspondences as opposed to isolated or inconsi-

stent correspondences. An isolated correspondence arises if there is a relational match between two domains, but the relation is not involved in other higher order relations, where a higher order relation is a relations between relations. Correspondences are inconsistent if they create many-to-one mappings or cross-mappings. Consider the three families and their telephone conversations:

John Jones = Father	Saul Smith = Son	Andy Anderson = Father
Jesse Jones = Son	Sam Smith = 2nd Son	Alan Anderson = Son
John = Not talkative	Saul = Talkative	Andy = Talkative
Jesse = Talkative	Sam = Talkative	Alan = Not Talkative

There is a *two-to-one* mapping between the Jones and the Smith family. The natural mapping would place Jesse in correspondence with both Saul and Sam since all three are sons and talkative. There is a *cross-mapping* between the Joneses and the Andersons. The term “cross-mapping” was first used by Gentner and Toupin (1986) to describe competing correspondences due to superficial and relational properties. The “fatherhood” aspect would place John in correspondence with Andy, and Jesse in correspondence with Alan, whereas the “talkativeness” aspect would set up precisely the opposite correspondences. Many-to-one mappings and cross-mappings will never be created as analogical interpretations in Falkenhainer, Forbus, and Gentner’s (1990) computational instantiation of SMT, and are subject to strong inhibitory pressures in ACME.

There is psychological evidence that people judge the goodness of an analogy by the coherence of its relational structure. Clement and Gentner (1988) show that systematic facts, facts that participate in higher order relations, are more likely to be judged as important for an analogy than are nonsystematic facts. Gentner and Toupin (1986) find that older but not younger children resolve cross-mappings by ignoring superficial/attribitional information and concentrating on highlighted relations. Gick and Holyoak (1983) present evidence that previously solved problems aid in solving analogous problems if the relational commonalities between the problems are highlighted (by presenting several problems with the same relational structure). In general, the current evidence suggests that analogies are created and evaluated by placing the elements (relations, and perhaps objects) of one domain into correspondence with the other domain. Furthermore, what correspondences are made depend on the other correspondences that have been established. In analogical reasoning, as in perceptual mapping, global consistency constrains the mappings that are created.

3. MODELS OF SIMILARITY AND MAPPING

Given that a global-consistency constraint on scene-to-scene mapping is found in both perceptual (apparent motion and stereoscopic vision) and conceptual (analogy) comparisons, it might well be expected that this constraint would also

be found in people's judgments of similarity. One reason for thinking this is that the similarity of two objects or scenes depends on both perceptual and conceptual factors. Dogs and cats are similar both because they have roughly the same body shape (four legs, body, head) and because they are domesticated pets.

Given the above framework, it is surprising that the two predominant models of similarity in cognitive psychology, multidimensional scaling (MDS), and Tversky's (1977) Contrast model, provide no allowance for mappings that are constrained by other mappings (the global-consistency constraint). In fact, there is very little in either model to suggest that the process of mapping or aligning parts of scenes/objects is part of the process of computing similarity.

3.1 Multidimensional Scaling

The purpose of multidimensional scaling (MDS) is to discover the underlying dimensions that account for a given set of proximity data, and to provide an account for how the dimensional information is combined to yield a measure of similarity. The input to MDS routines may be similarity judgments, dissimilarity judgments, confusion probabilities, correlations, joint probabilities of occurrence, or any other measure of pairwise proximity. The output from an MDS routine will be a geometrical model of the data, with each object of the data set represented as a point in an N -dimensional space. The distance between two objects' points in the space is taken to be inversely related to the objects' similarity.

The MDS modeler observes the geometric space or a rotated version of the space, and subjectively determines labels for the dimensional axes. Richardson's (1938) fundamental insight, which is the basis of contemporary use of MDS, was to begin with subjects' judgments of pair-wise object (dis)similarity, and work backward to determine the dimensions and dimension values that subjects used in making their judgments.

A study by Smith, Shoben, and Rips (1974) illustrates a classic use of MDS. They obtained similarity ratings from subjects on many pairs of birds. Submitting these pair-wise similarity ratings to MDS analysis, they hypothesized underlying features that were used for representing the birds. Assigning subjective interpretations to the geometric model's axes, the experimenters postulated that birds were represented by such features as "ferocity" and "size." In short, MDS models function both to derive possible feature dimensions and to represent similarity/proximity in terms of these feature dimensions. Even when MDS dimensions are difficult to interpret psychologically, the models can still serve a valuable function by ordering the proximity data geometrically.

3.2. Tversky's Contrast Model

A very influential model of similarity, the Contrast model, has been proposed by Amos Tversky and his associates (Gati & Tversky, 1984; Sattath & Tversky, 1987;

Tversky, 1977). In the model, the similarity of two entities is taken to be a linear contrast of the features that the entities share, minus the features possessed by one entity that are not possessed by the other. The mathematical formulation

$$\text{SIM}(A,B) = \alpha \cdot f(A \cap B) - \beta \cdot f(A-B) - \gamma \cdot f(B-A)$$

is interpreted as: The similarity of A to B is a function of the features that A and B share, minus the features that A has that B does not have, minus again the features that B has that A does not have. The greek letters are simply weighting terms that depend on the subjects' task and the stimuli. Although not inherent to the Contrast model, the further assumption is often made that the function " f " satisfies feature additivity such that $f(x)$ for any set x is expressible as the sum of the measures of all the features that belong to x .

In addition to the mathematical formulation of similarity, the Contrast model has processing principles associated with it. These principles provide rules for changing feature weights. For example, according to the diagnosticity principle (Tversky, 1977), the weight of a feature depends on its classificatory significance, the "importance or prevalence of the classifications that are based on" the feature (p. 342). According to the extension principle, the similarity of a pair of objects in an extended context will tend to be larger than the pair's similarity in a more limited context. For example, "carrot" and "cucumber" are more similar in a context that includes "jeep" than in a context that only includes vegetables.

3.3. Assumptions Common to Both Models

There are many differences between the MDS model and the Contrast model. In fact, the Contrast model was in part formulated as an alternative to the metric MDS model's strong geometric assumptions regarding symmetry (Distance(A,B) = Distance(B,A)), the triangle inequality (Distance(A,B) + Distance(B,C) ≥ Distance(A,C)), and minimality (Distance(A,B) ≥ Distance(A,A) = 0). However, the models share three assumptions with which we will take issue:

1. Entities are represented as a set of features (Contrast model) or dimension values (MDS). In the Contrast model, entities are represented by their set of features. Sets of features are compared for match and mismatch to determine the similarity of two entities. In the MDS models, entities are represented by their N -dimensional location; entities are defined by their values on each of the N dimensions.
2. Entity representations and features weightings do not depend on the actual pair-wise comparison. This is very clear in MDS: once an MDS solution space has been derived, an entity's values are set. In some MDS schemes, the weights of the dimensions can vary (Carroll & Wish, 1974; Nosofsky, 1986) and in some schemes the context can change the MDS solution (Roth

& Shoben, 1983). However, in all cases, once a context has been set up and all of the pairs have been presented, an entity's representation is fixed and the entity's features and feature weightings are defined. In the Contrast model, a feature's weighting can change in accordance with the diagnosticity or extension principles. Neither of these principles should influence feature weights after a pair-wise comparison has begun. The extension principle may weight a particular feature highly, but this change of weighting affects the representation of the entity independently of its comparison to a second item. A diagnosticity effect is either not present for pair-wise comparisons (because there is only one grouping of two items), or its effect is also prior to the actual comparison. Primarily, both the Contrast model and MDS models work by setting up feature/dimension representations of all of the entities, and then comparing the entities (measuring distances or set overlap). The actual pair-wise comparison operation uses the previously established representations; it does not establish these representations.

3. The alignment process is straightforward and simple. If an object is simply represented by a list of features, then alignment is a trivial issue and global consistency provides no constraint on creating object-to-object correspondences. Features are placed in correspondence if they are identical (one scene's "red" feature is placed in correspondence with the "red" of the scene it is compared with). In the case of MDS, the dimensions provide the alignment and values on these dimensions determine the degree or extent of matching. Although these assumptions may work for certain sets of stimuli, we shall soon see that often things are much more complicated.

The relatively unstructured representation of objects/scenes in MDS and the Contrast model stand in stark contrast to the richer first-order predicate logic and propositional representations invoked to explain analogical reasoning. The assumption that entities are represented by sets of features/dimension values makes no provision for hierarchical organization (features do not contain features as parts; features are nondecomposable primitives) or propositional organization (features do not take features as arguments).

We believe that issues of alignment are critical in visual similarity comparisons and that a major component of the comparison process involves setting up the correspondences between the objects/parts of the compared scenes. Furthermore, as with the perceptual and conceptual domains surveyed, we find that the alignment of objects is not independent of the other established correspondences—visual similarity is influenced by global consistency also. In opposition to the second assumption of comparison-independent representations, the weight that a particular feature has in a similarity calculation and the way it aligns cannot be determined before the actual similarity comparison process takes place. The process of placing scenes in correspondence changes the

weights of matching and mismatching features. Finally, the model for alignment we present argues against the idea that determining correspondences is simple or can be ignored.

As a preview of the psychological experiments to come, consider Figure 6.2, which is based on a figure from Gati and Tversky (1984). Figure 6.2 is virtually identical to Gati and Tversky's stimuli except that a third mountain has been deleted and a cloud has been added. Previous research may have ignored issues of alignment by designing stimuli that lack many-to-one mappings or cross-mappings. Gati and Tversky compare scenes like Figure 6.2A to other scenes that have features such as clouds, houses, and mountains added or deleted. By collecting similarity ratings for several variations of the same scene, they can ascertain the importance of shared and distinctive features for similarity judgments.

However, it is not always easy to say what counts as a matching feature. For example, does the fact that both the right cloud of Figure 6.2A and the left cloud of Figure 6.2B are spotted count as a *matching feature* between these scenes?

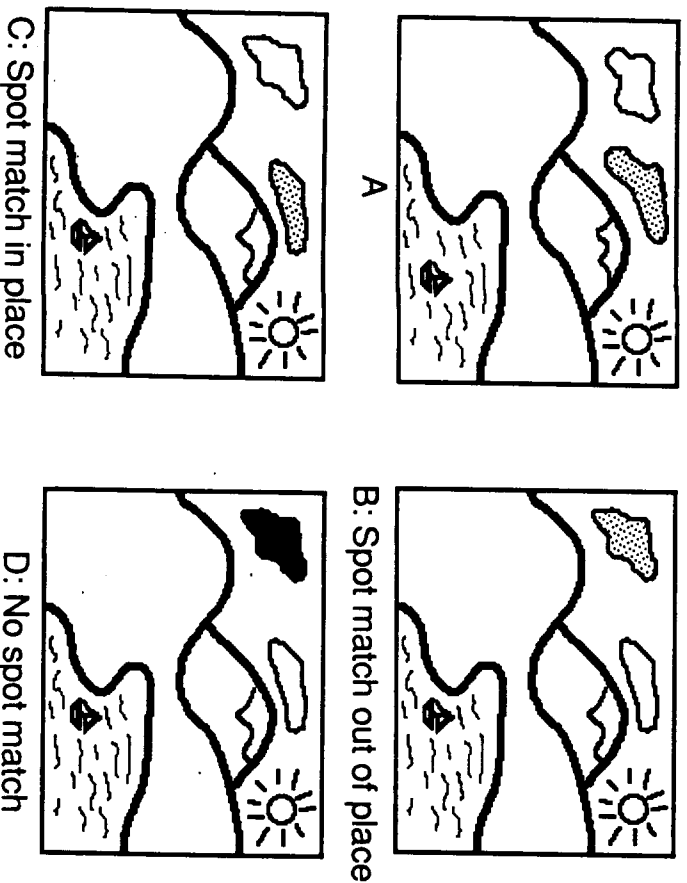


Figure 6.2.

The influence of matches in place (MIPs) and matches out of place (MOPs) on similarity. If MIPs increase similarity more than MOPs, then Scene A will be more similar to Scene C than Scene B. If MOPs increase similarity at all, then A will be more similar to B than D.

Does the feature "spots" increase the similarity of these scenes? It would seem that *how much a feature match increases similarity depends on whether it belongs to objects that correspond to each other or not*. If the right clouds of each scene were spotted, as they are in Figure 6.2A and 6.2C, then the shared spots would clearly increase similarity, but the spots would not increase similarity much if they belonged to the clouds of one scene and the sailboat of the other scene.

We will refer to feature matches between objects that are aligned with one another as *Matches in Place*, or MIPs. Feature matches between objects that are not aligned will be called *Matches out of Place*, or MOPs. We might presume that MIPs increase scene similarity more than MOPs—the fact that two scenes contain the feature "spotted" increases similarity more if the spots belongs to two clouds that correspond to each other than if the spots belong to non-corresponding objects. Thus, the similarity of Figures 6.2A and 6.2C is presumed to be higher than the similarity of Figures 6.2A and 6.2B. We might also presume that MOPs increase similarity more than no feature match at all—the fact that two noncorresponding objects are spotted results in a higher similarity than if the objects were not both spotted, all else being equal. Hypothetically, the similarity of Figures 6.2A and 6.2B is higher than the similarity of Figures 6.2A and 6.2D.

It is conceivable that a simple "feature list" representation of scenes could explain why feature matches only count if they belong to aligned objects. Instead of representing objects as containing "spotted" or "white" features, features could represent *conjunctions* of properties, such as "spotted cloud" or "right spotted cloud." If the clouds' features of Figure 6.2A are "left-white-cloud" and "right-spotted-cloud" while Figure 6.2B's features are "left-spotted-cloud" and "right-white-cloud" then the "feature list" representation could account for the inability of the scenes' matching spots to increase similarity. If these conjunctions were the encoded features, then Figure 6.2A and 6.2B have no cloud features in common at all.

Of course, it is reasonable to expect (and empirically supported, as we will see) that the matching spots increase the similarity of Figures 6.2A and 6.2B somewhat; the increase is just smaller than would be expected if the matching spots belonged to objects that were aligned. Keeping with the simple "feature list" representation of scenes, we could account for this pattern also, but only by positing both simple features such as "spotted" and conjunctive features such as "left white cloud." We need the simple features to account for the tendency of matching spots to increase similarity even when the spotted objects do not correspond to each other. We need the conjunctive features to account for the tendency for matching spots to increase similarity more if the spots belong to aligned objects than if they belong to nonaligned objects.

Problems with saving "feature list" representations by positing both simple and conjunctive features are that: (a) no account is given of how objects and object parts are aligned, (b) the number of features required grows exponentially

with the complexity of the scenes, yielding a computationally intractable representation system, and (c) the account may not be able to handle, in detail, experimental evidence concerning how MOPs and MIPs influence similarity under various task manipulations. The first problem is that an unstructured similarity model, replete with simple and conjunctive features, still does not account for the intuition that objects and object parts are placed into correspondence during similarity comparisons. Scenes are still decomposed into a "flat" list of features with no hierarchical structure or account of how some features constrain others. The second problem is that, in order to account for all hypothetical similarity comparisons, $2^k - 1$ features are required for every object in a scene, where k is the number of object components. For example, if a scene has an object with four parts/features (a, b, c, d), then 15 simple and conjunctive features ($a, b, c, d, ab, ac, ad, bc, bd, cd, abc, abd, acd, abcd$) are required just for this object to account for the scene's similarity to other related scenes. Fleshing out the third problem will be the central objective of the rest of this chapter. In the next section, we will review several experiments that investigate the influence of MIPs and MOPs on similarity ratings, indirect measures of similarity, and alignment judgments. Several of these experiments have been conducted in collaboration with Dedre Gentner. Throughout the review, the central point will be that the process of comparing scenes places scene parts in correspondence, and that this alignment process accounts for the influence of time, context, and feature distribution on scene comparison.


4. EXPERIMENTAL SUPPORT FOR ALIGNMENT IN COMPARISON

4.1. Initial Investigation of Mapping and Similarity

The first experiment was designed to corroborate our intuitions about the influence of matches in and out of place on similarity. Briefly, two scenes were shown side by side on a computer screen. Each scene was composed of two butterflies. Each butterfly varied on four dimensions: head type, tail type, body shading, and wing shading. For each pair of scenes, subjects *assigned a similarity rating* (1–9, where 9 means "HIGHLY SIMILAR") and then *indicated which butterflies corresponded to each other* between the scenes. For each butterfly of one scene, subjects indicated the butterfly of the other scene that corresponded best to it, where subjects were left to define "correspond" themselves. The two scenes were constructed so that the number of matches in and out of place were varied.

Figure 6.3 shows two of the possible comparisons that were shown to subjects. In the top comparison, the only difference between the scene on the left and the scene on the right is related to the butterflies' body shadings. Butterfly

Figure 6.3. Two possible displays from the initial experiment. In the top display, concentrating on the body shadings of the butterflies, there are 2 MOPs between the left scene and the right scene. In the bottom display, there is a single MOP along the body shading dimension.



C, while most likely corresponding to Butterfly A because they share three out of four features in common, possesses the body shading (wavy lines) of Butterfly B. Likewise, Butterfly D, though corresponding to B, has A's body shading. This is similar to the cross-mapping situation in analogy described earlier. Concentrating only on the body-shading dimension, there are two MOPs between the scenes—one between A and D, and the other between B and C. The lower comparison of Figure 6.3 has one MOP along the body-shading dimension. Again, A corresponds to C, but has the same body shading as D. Butterflies B and C have different body shadings.

In all, there were six methods for changing a dimension from one butterfly scene to another. These methods resulted in the following numbers of MIPs and MOPs along a dimension: 2 MOPs, 1 MOP, 2 MIPs, 1 MIP, 1 MOP, and 1 MIP, and no matches at all.

On half of the trials, in going from the left scene to the right scene, we change the features of one dimension, in one of six ways. On the other half of the trials, we change two dimensions, each in one of six ways.

On one-third of the trials, the butterflies that correspond to each other are placed in the same relative positions in the two scenes. On one-third of the trials, the butterflies are given new unrelated positions. The particular features that were switched, the positions of the butterflies, and the particular values that the dimensions had were all randomized.

The results, as shown in the first column of Figure 6.4, reveal an influence of both matches in and out of place on similarity. First of all, similarity ratings for 0, 1, and 2 MIPs averaged 5.5, 6.4, and 6.1, respectively. MOPs have a much smaller effect; the ratings for 0, 1, and 2 MOPs averaged 5.5, 5.5, and 5.9, respectively. The fact that MIPs increase similarity more than MOPs can also be seen by looking at scene comparisons that have the same number of total matches. For example, the similarity rating in the first column for two MIPs is 7.1. This decreases to 5.9 in scenes where there are two MOPs. (In the first column, rating differences of 0.2 are significant at $p < .05$.) Further, the similarity rating for scenes with 1 MIP (6.4) is greater than the rating for scenes with 1 MOP(5.5).

Description	Similarity		Mapping Accuracy	
	True Mappings	False Mappings	One Dimension Changed	Two Dimensions Changed
Two matches in place	7.1	6.6	91%	—
One match in place, one match out of place	6.5	5.5	90%	83%
One match in place	6.4	6.0	90%	85%
No match	5.5	4.9	89%	83%
One match out of place	5.5	4.8	86%	76%
Two matches out of place	5.9	5.3	86%	62%

Figure 6.4. Results from initial experiment investigating the influence of MIPs and MOPs on similarity.

We also find that if scenes share MOPs, feature matches that belong to butterflies that do not correspond to each other, then similarity is higher than if the scenes share no match along a dimension. In Figure 6.4, scenes with 2 MOPs are more similar than scenes with 1 MOP or no matches in common. However, there is no difference between scenes with 1 MOP and no matches at all. Interestingly, the phenomenon that $(2 \text{ MOPs} - 1 \text{ MOP}) > (1 \text{ MOP} - 0 \text{ MOPs})$ cannot be explained by an exponential relation between featural overlap and similarity (e.g., Nososky, 1986) because the finding holds even when the 0 and 1 MOPs displays are more similar overall than the 1 and 2 MOPs displays that are compared. Overall similarity is manipulated by varying the number of MIPs. In addition, positing a nonlinear relation between featural overlap and judged addition, posing a nonlinear relation between featural overlap and judged similarity does not explain why the phenomenon occurs when 2 MOPs are created by *swapping* dimension values but not otherwise. This inconstant influence of a MOP will be explained by our model in terms of cooperation and competition between developing correspondences.

A question might arise: How does the experimenter know that a feature match is really in or out of place? A MOP would be a MIP if subjects gave the opposite mapping of butterflies than was expected. In Figure 6.3, the expected mapping was to place A and C in correspondence, and B and D. Perhaps the hypothesized influence of MOPs is due to trials in which the subject gives the unexpected mapping (A is aligned with D, B with C).

To address this objection, in the first column, we only include similarity ratings for trials where the subject and experimenter are in agreement as to which butterflies correspond to one another. These trials are called "true mapping" trials because the mappings that the subject gives are optimal in the sense of maximizing the number of matches that are MIPs as opposed to MOPs. In the top half of Figure 6.3, the mapping "A goes to C, B goes to D" results in six MIPs and two MOPs. The alternative mapping of "A goes to D, B goes to C" results in six MOPs and two MIPs. Thus, the first mapping is the "true mapping" and the second mapping is the "false mapping." According to the first column in Figure 6.4, MOPs increase similarity even when they are MOPs *for the subject*.

Additional support for the hypothesis that scene alignment influences similarity can be obtained by comparing the true mapping and the false mapping trials. If the subject makes the mapping that maximizes the number of matches in place (the true mapping), then similarity is greater than if subjects make a nonoptimal mapping. Both the true and false mappings result in the same number of total (MIPs + MOPs) scene matches; the true mapping results in a greater number of MIPs relative to MOPs. Thus, the difference between true and false mapping trials provides further evidence that MIPs increase similarity more than MOPs.

The relation between similarity and mapping can also be clarified by examining the percentages of true mapping trials for the different scene types.

While the two rightmost columns of Figure 6.4 have the true mapping percentages, the rightmost column provides the most sensitive data: if only one dimension is changed, subjects do not make very many false mappings (and most of these are due to the different spatial locations of the butterflies in the scenes). If two dimensions are changed, we find that MOPs decrease mapping accuracy significantly (62% accuracy with 2 MOPs compared with 83% for 0 MOPs). There is also a much smaller, but significant influence of MIPs: the more MIPs there are, the greater the proportion of true mappings. The probability of finding the best correspondences between butterflies decreases as the number of MOPs increases and increases slightly as the number of MIPs increases.

A summary of the results from the first experiment reveals: (a) MIPs and MOPs both increase similarity, but MIPs increase similarity more, (b) if subject's give nonoptimal mappings, similarity is lower than if they give the optimal mapping, (c) MIPs simultaneously increase similarity ratings and mapping accuracy but MOPs increase similarity while decreasing mapping accuracy, and (d) the influence of a MOP depends on the other feature matches. The fourth conclusion is based on the significant difference in similarity between scenes with 2 MOPs and scenes with 1 MOP, but the lack of a difference between 1 MOP and no matches. The first two conclusions speak to our most central claim: the act of assessing similarity involves placing the parts of scenes in alignment with one another.

4.2. Replications of Basic Findings

We have conducted a number of experiments to replicate and extend the above results. The materials and procedures for the closest replication were identical to the first experiment, with one exception. In the first experiment, subjects were required to rate similarity and report the butterfly-to-butterfly correspondences. In the replication, subjects were only required to rate similarity. One criticism of the first experiment is that, by requiring subjects to perform mappings, we are in effect biasing them toward placing butterflies in correspondence. Perhaps there is no natural tendency for people to place parts of scenes into correspondence; they only do so in our experiment because they are forced to as a secondary task. In our replication, we eliminate this secondary task but we find exactly the same pattern of significant results. For example, we still find that MIPs increase similarity more than MOPs. By not requiring mapping judgments to be made, we reduce the task biases substantially, but we also lose the ability to tell for certain whether the match assumed to be "in place" for the experimenter is indeed a MIP for the subject.

Similarity ratings provide one measure of similarity with a good deal of face validity, but our research strategy has been to develop many converging

Figure 6.5.

The results from four converging measures of similarity yield similar patterns of results. All measures of similarity increase more with increasing MIPs than MOPs, but show some influence due to MOPs.





Figure 6.6.

For these stimuli, a match out of place is a matching symbol that does not belong to corresponding bird parts.

measures of similarity. If our results are general, then we would expect mapping to show an influence on similarity where similarity is operationalized in many ways. Previous research has shown that different operationalizations of similarity do not necessarily converge. Empirically, the similarity of A to B as measured by the percentage of times A is confused with B is not necessarily equivalent to the similarity rating of A to B (Gati & Tversky, 1984; Keren & Schooler, 1981). Similarity as measured by the time required to search for A in a field of Bs also differs from the similarity rating measure (Beck, 1966). However, if several different measures of similarity all show an influence of mapping on similarity then we will have stronger grounds for positing a general relationship between mapping and various conceptions of similarity.

In one replication, we display scenes similar to the scenes shown in Figure 6.3. Instead of asking subjects to say how similar the scenes are, we require subjects to respond as quickly and accurately as possible as to whether the scenes have exactly the same butterflies in them. If the same butterflies are in both scenes, subjects press a key meaning "same"; otherwise, they press the "different" key (the correct response to both displays of Figure 6.3 would be "different"). Similarity is operationalized in two ways. First, similarity is operationalized as the average time required to correctly respond that two scenes are "different." If two scenes are highly similar, then it will take a relatively long time to respond that the scenes are different. Second, similarity can be operationalized as the percentage of times that a subject incorrectly responds that two scenes are the same when the scenes are actually different. If two scenes are highly similar, people who are under pressure to respond quickly will often mistakenly respond that the scenes are the same.

The results from this "same/different" response time experiment closely follow the earlier results we obtained. The lowest correlation between any of the four measures was $r = 0.72$. The results for all four measures of similarity are shown in Figure 6.5. All measures show the same pattern of influence of MIPs and MOPs. MIPs and MOPs increase similarity, and MIPs increase similarity more. There are some differences between the measures of similarity. Most notably, the same/different response time measure of similarity is more influenced by MOPs than are the other operationalizations of similarity. Still, for all operationalizations of similarity, almost the identical ordering of the six combinations of MOPs and MIPs with respect to similarity is found.

One objection to the experiments mentioned is that they all involve comparing scenes that are likely to be not very well integrated. The scenes that were compared all were comprised of two butterflies. The butterflies were coherent and structured objects on their own, but the scenes as a whole do not seem very cohesive. Perhaps MOPs only increase similarity for scenes with little structure.

To address this criticism, we conducted an experiment with materials exemplified by Figure 6.6. Each scene is comprised of a single bird; each bird has four internal parts (head, lower wing, upper wing, body). The bird on the

left is compared to the bird on the right. Here, a MOP is defined as a symbol that is present in both birds, but occurs in different parts of the birds. A MIP is a matching symbol that is in the same part of the compared birds. In different experiments with these materials, we operationalize similarity as: (a) average similarity rating, and (b) response time to correctly respond that two birds are "different." In different scenes, we orthogonally vary the number of MIPs and MOPs from zero to two. With both operationalizations we find an influence of MIPs and MOPs. MOPs increase similarity over having no matching features, and MIPs increase similarity more than MOPs. Therefore, even with scenes that seem to form single, coherent entities with role-defined parts, there is still an influence of MOPs on similarity.

Using similar materials, we have also measured similarity in terms of effects of MIPs and MOPs on categorization. In a categorization task, subjects are presented the top-left object of Figure 6.7 as a Category 1 object, and the top-right object as belonging to Category 2. Subjects are repeatedly shown the two objects with their labels until they are well memorized. Subsequently, new "transfer" objects are presented for the subject to categorize. Subjects assigned

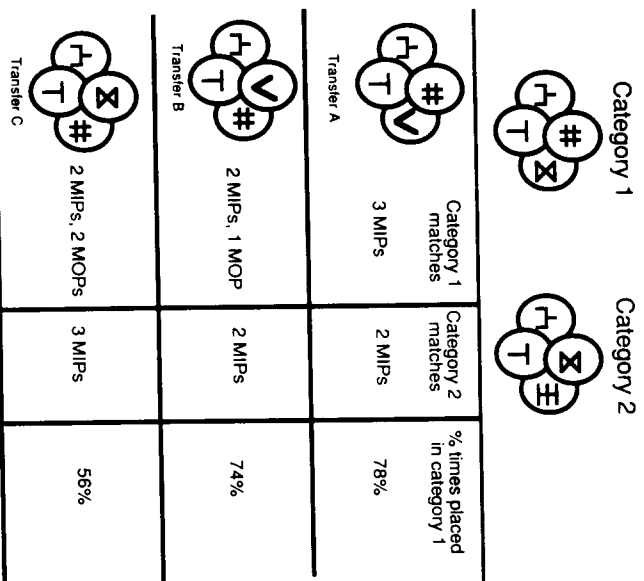


Figure 6.7.

A replication using a categorization measure of similarity. Here, similarity is defined by: X is more similar to Y than Z to the extent that X is placed in Y's category instead of Z's.

the transfer items to the category to which they seemed to belong. Subjects are not told whether they placed the novel objects in the correct category. The object labeled "Transfer A" has 3 MIPs in common with the Category 1 example, and 2 MIPs in common with the Category 2 example. The consistency with which "Transfer A" is placed in Category 1 is an indication of how much a single MIP increases the similarity of the transfer item to category members. Likewise, "Transfer B" has one more MOP in common with Category 1 than Category 2, so the subjects' tendency to categorize "Transfer B" as belonging to Category 1 reflects the importance of a MOP on similarity. "Transfer C" allows us to tell whether 1 MIP or 2 MOPs increase similarity more. If 2 MOPs increase similarity more, then this item should be placed in Category 1.

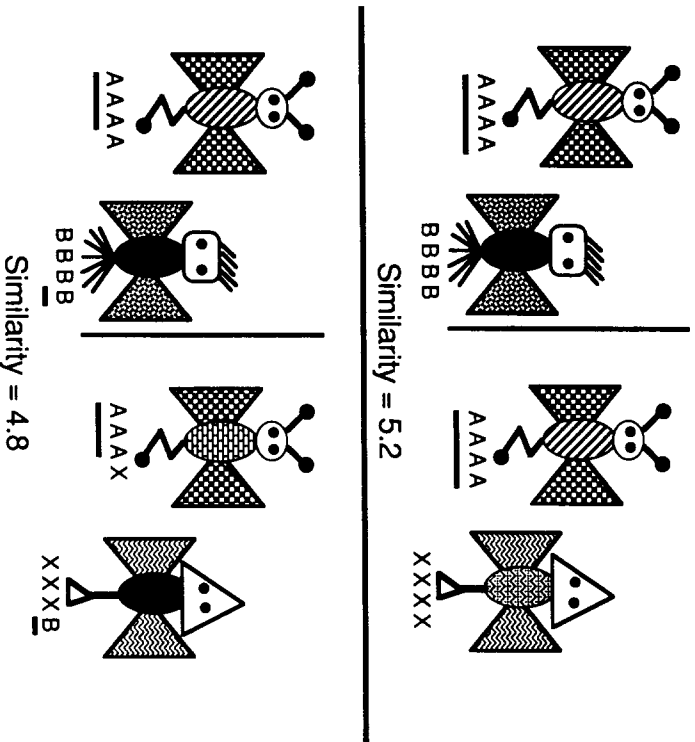
The results again show that MIPs increase similarity more than MOPs. The tendency for "Transfer A" to be placed in Category 1 is significantly greater than the tendency to place "Transfer B" in Category 1. However, MOPs still increase similarity. "Transfer B" is placed in Category 1 at a rate greater than the chance level of 50%, indicating that the MOP is influencing categorization. With this categorization paradigm, we find that 2 MOPs increase similarity slightly more than 1 MIP. Overall, we find converging results to our earlier operationalizations of similarity, using the operationalization that "the more similar a novel object is to a previously categorized object, the more likely the novel object will be placed in the categorized object's category." This result shows both the influence of aligning scene parts on categorization and a link between similarity and other basic cognitive processes.

4.3. Other Experimental Findings

Although the following findings may seem rather loosely connected, the observations will later be integrated into a single account of mapping and similarity. The empirical results will be presented first, followed later by the model's explanations of the results.

4.3.1. MIPs and Feature Distribution. A robust effect found in our experiments is that the importance of a MIP on similarity depends on how other MIPs are distributed. First, similarity increases more if feature matches are concentrated in one pair of objects. Consider the two displays in Figure 6.8. In both cases, there are four MIPs between the scenes (only true mapping trials are included). In "concentrated" displays, such as the top display, the four MIPs are concentrated in one pair of butterflies. One pair of butterflies has four matching features, and the other pair of butterflies has no matching features. The average similarity rating for "concentrated" displays is significantly greater than the rating for "distributed" displays. In the lower "distributed" scene, two butterflies have three feature matches in common, and the other two butterflies have a single matching feature. Similarity decreases still further, to 4.6, when the MIPs are evenly split between the two pairs of butterflies. In short, a MIP


Figure 6.8. MIPs that are concentrated in a single pair of objects (top) increase similarity more than MIPs that are distributed over multiple pairs of objects (bottom).



We get an analogous finding with dimensions. Namely, similarity is higher if feature matches are concentrated in a few dimensions rather than distributed across many dimensions. If four feature matches are concentrated in two dimensions (for example, wing shading and head), then similarity is higher than if the four feature matches are distributed across all four dimensions.

4.3.2. Nondiagnostic Features and Mapping Accuracy. Feature matches that cannot serve, by themselves, as cues for placing scene parts into correspondence still can increase mapping accuracy. Mapping accuracy can be measured as the percentage of time that subjects place butterflies in correspondence with each other in a manner that maximizes the number of MIPs. In the left scenes of Figure 6.9, three dimensions are nondiagnostic: wings, body, and tail. These dimensions are nondiagnostic because both butterflies in the right scene do not have any nondiagnostic features in common with the left butterflies. The ability of subjects to determine the optimal, or true mapping is poor; errors of mapping

Figure 6.9. Nondiagnostic features, if shared by two scenes, can still increase mapping accuracy. In the top display, the nondiagnostic features of the left scene are not shared by the right scene's butterflies. In the bottom display, the left scene's nondiagnostic features are shared by the right scene's butterflies, thereby decreasing the number of mapping errors made.



are made on 33% of trials.¹ A mapping error occurs if subjects respond that the top-left butterfly of the left scene corresponds to the bottom-right butterfly of the right scene.

Mapping performance greatly improves in the lower display, where the three nondiagnostic features of the left butterflies are also present in the right butterflies. Even though the nondiagnostic features provide no direct cue that the top-left butterflies correspond to each other (as do the bottom-right butterflies), the shared nondiagnostic features do increase responses on the basis of the HEAD dimension which *is* diagnostic for mapping. Subjects make their mappings on the basis of the butterfly heads more when other features between the scenes match than when they mismatch.

¹ One reason for such poor performance is that the butterflies that correspond to each other are not always placed in the same relative spatial locations as they are in Figure 6.9. If corresponding butterflies are given the same spatial locations, errors are made on only 18% of trials; this increases to 37% if the corresponding butterflies are given unrelated spatial locations, and to 41% if the spatial positions of butterflies are switched.

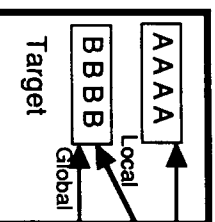
4.3.3. *The Time Course of MIPs and MOPs.* According to the feature integration theory developed by Anne Treisman and her associates (Treisman & Gelade, 1980), focused selective attention is required to bind object features to their objects/locations. If an object is displayed quickly, before attention can be directed to it, then the object's features will not be bound to the object. One result of unbound features is "illusory conjunctions"; features from different objects are combined together. The red color from one object, for example, may be perceptually conjoined with the circular shape of another object.

One might consider applying this logic to our paradigm. We might expect that MOPs would show most of their influence on similarity fairly early in processing, just as illusory conjunctions occur only before attention can bind features to objects. With time, as object-to-object correspondences become clear, only MIPs would show a large influence on similarity. The feature matches become *bound* to their correct correspondence with time. Until these correspondences are created, there would be little difference between a MIP and a MOP.

One of our "same/different" experiments supports this notion. Subjects saw two scenes and were required to report whether the scenes contained exactly the same butterflies. The error rate on "different" trials was used as a measure of similarity. Hypothetically, the more similar the two scenes are, the more likely a subject will be to incorrectly report that the scenes are the same. The butterflies in one of the scenes, the target scene, can be abstractly represented as AAAA and BBBB, signifying that the butterflies are composed of four features, and the features of one butterfly are completely different from the other butterfly's features. In the first panel of Figure 6.10 we compare the target scene to a base scene. If we require subjects to respond within 2.5 seconds, a slow deadline, subjects mistakenly think the scenes have the same butterflies 6% of the time. This error percentage increases to 27% if subjects are forced to respond within 1 second.

The important aspect of the target and base scenes is that both of the target butterflies have the most matching features with the top butterfly of the base scene. The top butterfly of the base scene, BABA, has two matches in common with both of the target scene's butterflies. Thus, if we ignore the fact that one object mapping can constrain another object mapping and just consider the locally preferred mappings, we would want to map both target butterflies onto the top butterfly of the base scene. However, if we maintain the global consistency of our mappings, then we would not permit this many-to-one mapping. The best globally consistent mapping is to map the top butterflies to each other, and the bottom butterflies to each other. MIPs are defined as feature matches between parts that are placed in "true" correspondence, and parts truly correspond if they belong to the set of consistent correspondences that results in the largest number of MIPs. As such, MIPs are matches between objects that are

Figure 6.10. **The temporal course of MIP and MOP influence.** If subjects must respond within a very short time, then locally consistent matches increase similarity (as measured by % confusions) more than globally consistent matches. If subjects are given a longer deadline, then globally consistent matches increase similarity more than locally consistent matches.

	<p>Target</p> <p>AAAA</p> <p>BBBB</p>	<p>Base</p> <p>BABA</p> <p>XXXB</p>	<p>Globally consistent match kept</p> <p>XABA</p> <p>XXXB</p>	<p>Locally preferred match kept</p> <p>BABA</p> <p>XXXX</p>
<p>% confusions with Target, Slow Deadline</p>	<p>6%</p>	<p>6%</p>	<p>3%</p>	
<p>% confusions with Target, Fast Deadline</p>	<p>27%</p>	<p>18%</p>	<p>21%</p>	

mapped in a globally consistent fashion: MOPs are matches between objects whose correspondence is not globally consistent.

In the next two scenes, we either take away one of the locally consistent matches leaving the global matches intact, or take away one of the globally consistent matches leaving the local matches intact. If we preserve the local match, then there are more confusion errors with the target display than if the global match is kept, *but only at the fast deadline*. At the slow deadline, keeping the global match increases similarity/errors more than keeping the local match. Locally consistent matches are more important than globally consistent matches for similarity *early* in processing (fast deadline). Later in processing, globally consistent matches are more important than locally consistent matches. It seems that it takes time to set up the influence that one object-to-object mapping has on another object-to-object mapping, and until this happens, error data show the influence of many-to-one mappings. At first, both butterflies of the target are mapped onto one butterfly of the other scene, but with time the influence of one mapping redirects the other mapping.

4.3.4. Sensitivity to Features of Aligned and Unaligned Objects. In addition to obtaining estimates of similarity and mapping judgments, a third source of data is the sensitivity with which subjects are able to report feature matches and mismatches. We have found that subjects are more sensitive in reporting feature matches and mismatches that occur between objects that are aligned than objects that are unaligned.

Subjects are presented scenes composed of two butterflies that are displayed on the screen for only five seconds. After the screen is erased, subjects first give a similarity rating for the two scenes. Then, two pointers appear on the screen, pointing to the previous locations of two butterflies. Subjects are told to respond as to whether the butterflies referred to by the pointers had matching values on a particular dimension (head, tail, body, or wings). Using the top display of Figure 6.3 as an illustration, the following four types of questions are asked, with the following results:

Aligned Matches: Do A and C have the same WING SHADING? The correct answer (Yes) was given on 85% of trials.

Aligned Mismatches: Do A and C have the same BODY SHADING? The correct answer (No) was given on 71% of trials.

Unaligned Matches: Do A and D have the same BODY SHADING? The correct answer (Yes) was given on 52% of trials.

Unaligned Mismatches: Do A and D have the same WING SHADING? The correct answer (No) was given on 80% of trials.

These data are based on displays where one scene is identical to another scene except along one dimension. The single changed dimension is changed by introducing one or two completely new butterfly features, or by *swapping* one or two butterfly features.

These data suggest both a response bias and a sensitivity change (for a discussion of these and other notions in signal detection theory, see Swets, Tanner, & Birdsall, 1961) due to the alignment of butterflies. The response bias is that if the butterflies correspond to one another, then subjects are more likely to respond "Yes, the features match" than if the butterflies are not aligned. The sensitivity change is reflected by the overall increase in accuracy in judgments for aligned butterflies over unaligned butterflies. A signal detection analysis indicates significantly greater sensitivity (d') for feature matches and mismatches that occur between butterflies that correspond to one another. Thus, it is not simply that people assume that all features that belong to corresponding objects match. Subjects are highly accurate at identifying *mismatching* features between corresponding objects, much more accurate than they are at reporting matching features for objects that do not correspond to one another. The act of placing objects into correspondence increases sensitivity to *all* of their feature correspondences, matching or mismatching. In order to know how likely a

person will be to detect a matching or mismatching feature between two scenes' parts, it is necessary to know whether the parts correspond.

5. AN INTERACTIVE ACTIVATION MODEL OF MAPPING AND SIMILARITY

5.1. A Brief Overview of SIAM

In attempting to provide a framework to organize the seemingly disconnected fragments of data reported, we have developed a model SIAM (Similarity as Interactive Activation and Mapping).² Originally inspired by McClelland and Rumelhart's (1981) interactive activation model of word perception, SIAM also bears many conceptual resemblances to Falkenhainer, Genter, and Forbus's (1990) SME, and many architectural and conceptual resemblances to Holyoak and Thagard's (1989) ACME system.

The primary unit of operation is the node. Nodes do only two things: (a) send activation to other nodes, and (b) receive activation from other nodes. As in ACME, nodes represent hypotheses that two entities correspond to one another in two scenes. In SIAM, there are two types of nodes: feature-to-feature nodes and object-to-object nodes.

Feature-to-feature nodes each represent a hypothesis that two features correspond to each other. There will be one node for every pair of features that belong to the same dimension; if each scene has O objects with F features each, there would be O^2F feature-to-feature nodes.³ As the activation of a feature-to-feature node increases, the two features referenced by the node will be placed in stronger correspondence. All node activations range from 0 to 1. In addition to activation, feature-to-feature nodes also have a "match value," a number between 0 and 1 that indicates how similar the two features' values on a dimension are. If two butterflies have the same type of head, then the node hypothesizing that these two heads correspond to each other would receive a match value of one. The match value decreases monotonically as the similarity

² A simplified version of SIAM is presented. A slightly different activation function and network architecture is presented in Goldstone (1991).

³ Although substantially less costly than the simple feature + conjunction feature encoding of figures (requiring 2 \times 1 features per object if there are k simple features per object), SIAM and ACME both require the postulation of a large number of nodes for complex objects. Both SIAM and ACME create nodes for every possible feature-to-feature mapping. This proliferation of nodes is both computationally costly and psychologically implausible. Mitchell and Hofstadter's (this volume) Copycat architecture permits creation of mappings "on the basis of need," mappings are created only if code is executed to create them; it remains to be seen whether the Copycat architecture can successfully mimic the complex inhibitions and excitations of mappings in SIAM and ACME without generating all potential mappings.

of the butterflies' heads decreases. Match values in SIAM play a similar role to the semantic unit in ACME. Both structures serve to place parts that are primitively similar in correspondence. In addition, match values influence similarity directly; similarity is a function of the match values, weighted by the attention paid to them.

Each object-to-object nodes represents an hypothesis that two objects correspond to one another. There will be O^2 object-to-object nodes if there are O objects in each of two scenes. As the activation of an object-to-object node increases, the two objects are said to be placed in tighter correspondence with each other.

At a broad level, SIAM works by first creating correspondences between the features of scenes. At first, SIAM has "no idea" what objects belong together. Once features begin to be placed into correspondence, SIAM begins to place objects into correspondence that are consistent with the feature correspondences. Once objects begin to be placed in correspondence, activation is fed back down to the feature (mis)matches that are consistent with the object alignments. In this way, object matches influence activation of feature matches at the same time that feature matches influence the activation of object matches.

As in ACME and McClelland and Rumelhart's original work, activation spreads in SIAM by two principles: (a) nodes that are consistent send excitatory activation to each other, and (b) nodes that are inconsistent inhibit one another. Figure 6.11 illustrates the basic varieties of excitatory and inhibitory connections in SIAM. There are four ways in which the activation from one node influences the activation of another node:

1. Feature-to-feature nodes inhibit and excite other feature-to-feature nodes. Feature correspondences that result in two-to-one mappings are inconsistent; all other correspondences are consistent. The node that places Feature 2 of Object A in correspondence with Feature 2 of C (the A2 ↔ C2 node) is inconsistent with the node that places Feature 2 of C in correspondence with Feature 2 of B (the B2 ↔ C2 node). These nodes are inconsistent because they would place two features from one scene into correspondence with a single feature of the other scene. These nodes inhibit one another. The A2 ↔ C2 node is consistent with the B2 ↔ D2 node; consequently, these nodes will excite one another.
2. Object-to-object nodes inhibit and excite other object-to-object nodes. This is analogous to the previous type of connection. Object correspondences that are inconsistent inhibit one another. The node that places A and C in correspondence inhibits the node that places B and C in correspondence (A and B cannot both map onto C) and excites the node that places B and D in correspondence.

3. Feature-to-feature nodes excite, and are excited by, object-to-object nodes. Object-to-object nodes that are consistent with feature-to-feature nodes will be excited. The node that places A into correspondence with C is excited by the node that places Feature 2 of A into correspondence of Feature 2 of C. The excitation is bidirectional; a node placing two features in correspondence will be excited by the node that places the objects composed of the features into correspondence. In other words, the A ↔ C node sends activation back down to the A2 ↔ C2 node.
4. Match values excite feature-to-feature nodes. Features are placed in correspondence to the extent that their features match. If a match value is greater than .5 (a value of 1.0 signifies two identical features), then the activation of the node that places the features in correspondence will increase. Otherwise the feature-to-feature node activation decreases.

These four activation-passing methods incorporate both weak and strong consistency. Correspondences are weakly consistent if they do not yield an inconsistent, many-to-one mapping. Lateral excitatory activation is based on




Figure 6.11.

Sample connections present in SIAM. Correspondences that are consistent excite each other. Inconsistent correspondences inhibit each other. Nodes are represented by ovals. Excitatory and inhibitory links are represented by solid lines. Dashed lines represent the object or feature mapping indicated by the node.

weak consistency. The $A2 \leftrightarrow C2$ feature-to-feature node is weakly consistent with the $B2 \leftrightarrow D2$ node, and the $A \leftrightarrow C$ object-to-object node is weakly consistent with the $B \leftrightarrow D$ node. Correspondences are strongly consistent if they are structural consequences of one another. Vertical, or hierarchical, excitatory activation is based on strong consistency. The $A2 \leftrightarrow C2$ feature-to-feature node sends excitatory activation to the $A \leftrightarrow C$ object-to-object node because if two objects' features correspond to each other, then the objects should correspond as well.

The net input to a node i is given by:

$$net_{i0} = \frac{\sum_{j=1}^n (A_{j0} W_{ij}) - \text{MIN}}{(\text{MAX} - \text{MIN})}$$

Where n is the number of afferent links to node i (including excitatory links from match values to nodes), A_{j0} is the activation of node j at time t , and W_{ij} is the weight of the link going from unit j to unit i . In the current modeling, all weights are set equal to 1.0 (for excitatory connections) or -1.0 (for inhibitory connections).⁴ Net_{i0} is the activation of node i normalized by the difference between the maximum (MAX) and minimum (MIN) activation values that i can possibly attain, given the number of inhibitory and excitatory afferents to i . Net_{i0} is constrained to lie between 0 and 1. If i has 2 inhibitory and 1 excitatory afferents, then $\text{MIN} = -2$ (if both inhibitory inputs were completely activated, and the excitatory input was zero) and $\text{MAX} = 1$. The new activation of a node at time $t+1$ is a synchronously updated function of the old activation at time t and the normalized input activation, net_{i0} , received by the node:

$$\text{if } net_{i0} > 0.5 \text{ then}$$

$$A(t+1) = A(t) + (1 - A(t)) * (net_{i0} - 0.5) * B, \text{ otherwise}$$

$$A(t+1) = A(t) * (0.5 - net_{i0}) * B$$

where B is a parameter for the rate of activation adjustment.

Once a cycle of activation adjustment has passed, similarity is computed via:

⁴ Goldstone (1991) allows for different sources of information to have different weights. Separate weights are given for the influence of match values, features, and objects on feature-to-feature nodes, and for the influence of features and objects on object-to-object nodes. When the weight associated with the influence of objects on features is a free parameter, data fits are somewhat better than in the currently presented model. Individual weight terms are not required for each feature value/dimension because the experiments randomized these variables, and the modeled data collapse over different feature value/dimension configurations.

$$\text{similarity} = \frac{\sum_{i=1}^n (\text{match value} * A_i)}{\sum_{i=1}^n A_i}$$

As such, similarity is computed as a function of the match values for each feature-to-feature node, weighted by the activation of the node. Thus, the more active a feature-to-feature node is, the more the particular matching or mismatching value shared by the features will influence similarity. If the features have the same feature value, then similarity will increase more if the feature-to-feature node's activation is high. Likewise, similarity will decrease more if the features do not have the same value if they are also placed in strong correspondence. The activation of a feature-to-feature node can be interpreted as the attention paid to a matching/mismatching feature. The greater the activation of a feature-to-feature node, the more the feature match of mismatch will influence similarity. Similarity is normalized (by dividing by $\sum A_i$) such that the minimum similarity is 0.0 and the maximum similarity is 1.0. By normalizing similarity, the similarity of two-scene displays with different numbers of objects and features can be compared⁵. It should not be assumed that similarity is algebraically computed by SIAM at the end of each cycle. Rather, this formula for similarity should be viewed as a shorthand way of characterizing the state of the network as a whole.

In our modeling, we set the following parameters to default values: $B = 1$, value for matching features = 1, and value for mismatching features = 0. The following other options of SIAM were not used for the present modeling: decay of activation, different W_{ij} terms for different types of connections, different W_{ij} terms for matching and mismatching values, different W_{ij} terms for matching and mismatching values, different feature saliences, and asymmetrical weighting of features from the two scenes. In the most general version of SIAM, these would all be free parameters.

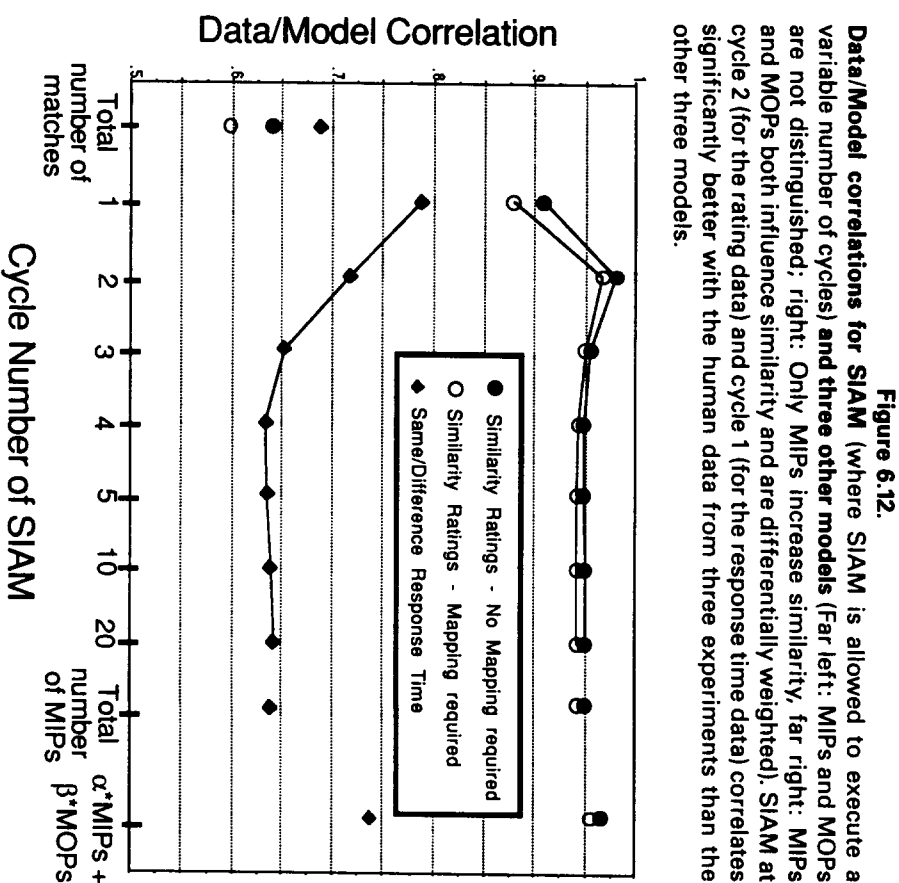
⁵ Tversky's Contrast model does not normalize for number of features. Two arguments are given for not normalizing: (a) objects with many features are judged to be both more similar and more different than objects with fewer features (e.g., Russia and United States are judged to be both more similar to and different from each other than are Ceylon and Nepal), and (b) many-featured identical objects are more similar than few-featured identical objects (e.g., an elephant is more similar to itself than a simple line is to itself). The first result is accommodated by SIAM if matching features are relatively influential in similarity judgments and mismatching features are relatively influential in dissimilarity judgments (the Contrast model also requires this assumption), and if missing features are given match values intermediate to those given for matching and mismatching features. The second claim is incompatible with SIAM's normalization, but its empirical status is uncertain, and intuitions are not completely clear that the claim is correct.

For the present modeling, the only parameter that is allowed to vary is the number of times SIAM cycles through the four methods of adjusting activations. Roughly speaking, the more cycles SIAM is allowed to complete, the more individual node activations will be influenced by global consistency. If SIAM runs for only a single cycle of activation adjustment, then MOPs and MIPs are almost equally weighted. If SIAM is allowed many cycles of activation adjustment, MIPs will become much more influential than MOPs. Objects will begin to be placed into correspondence, and once objects are placed in correspondence, they will send activation down to the features that are compatible with their alignment. MIPs will generally be compatible with the most activated object-to-object correspondences, and thus will receive the most activation. The isolated body shading MOPs of Figure 6.3 will not receive much weight in the similarity calculation because they will not receive much activation from the node that hypothesizes that A corresponds to D. The $A \leftrightarrow D$ node will not be highly activated because there are three inhibitory feature-to-feature nodes (head, tail, wings) and only one excitatory feature-to-feature node (body shading) activating it. In short, the influence of consistent feature-to-feature and object-to-object mappings becomes more pronounced as the number of iterations of activation adjustment increases. Aspects of the temporal course of human similarity assessments are modeled by varying the number of cycles SIAM executes.

SIAM is given two scene descriptions as an input. Each scene is organized into objects, and the objects are organized into features with feature values. A scene might be expressed as ((object 1 (head square) (tail zig-zag) (body-shading white) (wing-shading checkered)) (object 2 (head triangle) (tail zig-zag) (body-shading striped) (wing-shading spotted))). On the basis of the initial scene descriptions, SIAM constructs a network of feature-to-feature nodes and object-to-object nodes, and assigns initial match values (0 or 1) to feature-to-feature nodes. All node activations are initially set to 0.5. SIAM gives as output the similarity of the two scenes at every cycle of activation, the feature-to-feature correspondences (derived from the feature-to-feature node activations), the object-to-object correspondences, and the likelihood of performing a particular mapping (derived from the object-to-object activations).

5.2. Evaluation of SIAM

The results of fitting SIAM's output to subjects' data are shown in Figure 6.12. Data from three experiments are used: the original similarity rating study, the replication in which no mapping judgments are required, and the same/different response time experiment. Similarity, as measured by ratings and response time, is compared to predicted values from SIAM and three other models. SIAM is allowed to run from 1 to 20 cycles of activation adjustment. SIAM is fit



to the subjects' data with only "number of cycles" fit as a free parameter. All other parameters are set to their default values. SIAM is fit to the average similarity assessments for each functionally different display, averaging over subjects and over the particular featural instantiation of the display type. For example, in the first experiment, similarity ratings are modeled for 21 different display types, each with a different combination of MIPs and MOPs. For the similarity rating data, SIAM at two cycles correlates best with subject data (Pearson's $r = .983$ if no mappings are required, $r = .968$ if subjects are required to place butterflies in correspondence with each other).⁶ For the

⁶ More cycles of activation adjustment would be required if B were set to a value less than one. In the version of SIAM discussed by Goldstone (1991), more cycles of activation adjustment are required for the best fitting model because node activations do not asymptote to zero (for MOPs) or one (for MIPs). Even after activations have asymptoted, there is still an influence of MOPs on similarity.

response time data, SIAM fits the data best when it is only allowed one cycle of activation adjustment ($r = .78$). The different settings of the "number of cycles" parameter are reasonable because the "same/different" judgments are made much more quickly than the similarity ratings are, and speed is modeled in SIAM by reducing the number of cycles.

The three models SIAM is compared against include models where MOPs and MIPs are treated the same and both influence similarity, where only MIPs influence similarity, and where MOPs and MIPs are differentially weighted and both influence similarity. All three models yield significantly inferior fits to the subject generated data for all three experimental tasks. The worst fit is provided by the model that claims that similarity is a function of the total number of matching features, irrespective of whether the match is a MIP or a MOP. This model can be considered the "no alignment is necessary" model because alignment is necessary to decide whether a match is in place or out of place. This model is not capable of explaining our basic experimental result that the importance of a match depends on whether it belongs to objects that correspond to one another. According to this model, all matches are treated equivalently. This model is tested by running a linear regression of "total number of MIPs and MOPs" on similarity assessments. The demise of this model supports our claim that object alignment is a necessary consideration when evaluating a feature's impact on similarity.

The second "straw-person" model characterizes similarity as only a function of MIPs; MOPs do not influence similarity at all. The ability of this model to fit the rating data is quite good, reaching correlations of $r = .95$ and $r = .94$ for the two sets of rating data. Its fit of the response time data is much worse, as might be expected given the more substantial influence of MOPs in this data set. SIAM's superior fit⁷ supports the previous claim that pure "conjunctive feature" (combinations of simple features, such as "Red square") accounts of similarity are unable to explain the influence of MOPs on similarity. At the very least, we need both simple and conjunctive features.⁸

⁷ Although the differences in correlation are small, they are significant. When data/model correlations are near 1.00, very small differences are needed for statistical significance.

⁸ A model that includes both simple and conjunctive features is tested in Goldstone (1991) and found to fit the data worse than SIAM. A fatal downfall of the model is that it bases the importance of a feature match on the *similarity* of objects involved, and not on whether the objects are aligned. For example, in Figure 6.10, target object BBBB and base object BABA have two simple features, (B---and-B-) and one conjunctive feature (B-B-) in common. If either B from BABA were replaced, (B---and-B-) and one conjunctive feature match would be lost: BBBB and XXXB only have one feature both a simple and a conjunctive feature match would be lost: BBBB and XXXB only have one feature in common. As such, according to the "simple and conjunctive features" model, replacing a B in BABA will always decrease target-base similarity more than changing the B in XXXB. Conversely, SIAM predicts that, in the long run, the importance of a feature match depends on object alignment. BBBB is aligned with XXXB on the basis of global consistency. The empirical results support feature weighting based on object alignment rather than object similarity when these two factors are dissociated.

The final model assumes that both MIPs and MOPs influence similarity, and that the two types of match receive different weights. Similarity is assumed to be a linear regression of "number of MIPs" and "number of MOPs," where the weights given to MIPs and MOPs are free parameters chosen to optimize the fit to the subject data. This model captures the intuition that MIPs and MOPs both increase similarity, and that MIPs increase similarity more than MOPs.

SIAM predicts the subjects' data from the three tasks with reliably greater accuracy than even this third, more sophisticated model. Therefore, SIAM's success at modeling the subject's data is not simply due to the fact that it weights both MIPs and MOPs, and weights MIPs more.

How does SIAM differ from a model that simply weights MIPs and MOPs differentially? Put another way, *why do we need the complex interactive activation that SIAM has?* For the particular data that were modeled, there are two contributory reasons:

1. SIAM correctly predicts that 2 MOPs increases similarity over 1 MOP, whereas the difference between 1 MOP and 0 MOP is not significant. If 2 MOPs are arranged as they were in the scene's shown subjects (see the top half of Figure 6.3), then the MOPs will mutually excite one another because they are consistent with each other. Even though the MOPs will not be consistent with the optimal mapping, they will activate nodes that will send excitatory activation to each other. Two MOPs will support each other if they are consistent with the same object-to-object mappings, and they will support each other more if they occur on the same dimension (because feature-to-feature modes facilitate other consistent feature-to-feature nodes). A single MOP will not receive much weight because there will be no other feature matches that support it. Therefore, the non-linear relation between number of MOPs and judged similarity is a natural outcome of the mutual support relations in the interactive activation process.

2. SIAM correctly predicts that concentrated feature matches increase similarity more than distributed feature matches. Experiments described earlier showed that if feature matches are concentrated on a single dimension or on a single object then similarity will be relatively high. In SIAM, the reason for this is (again) that concentrated matches will mutually excite each other. If there are four feature matches in common between two objects, then the objects will be placed in strong correspondence with one another. Therefore, the object-to-object node representing these objects will be highly activated. Once activated, the object-to-object node will send activation down to the features that compose the objects. This "fed-back" activation results in a similarity assessment that strongly weights the four concentrated MIPs. Thus, the four feature matches influence similarity greatly because of the strong object correspondences they establish. Distributed MIPs will not activate any single object pair as strongly, and therefore, they will not receive as much "fed-back" activation.

There are also a number of advantages of SIAM over the other three models that have influenced its ability to account for other trends in the data.

1. SIAM provides an account for how we match objects. The best alternatives to SIAM are the model that has similarity as only a function of the number of MIPs, and the model that differentially weights MIPs and MOPs to arrive at a similarity assessment. *Both* of these models presume that there is a method for determining whether a match is a match in place or a match out of place. One model only includes MIPs, and the other model differentially weights MIPs. SIAM gives a process model for determining how objects and features are aligned and consequently gives a method for determining whether a given feature match is in or out of place. SIAM, in addition to computing similarity, also computes feature-to-feature and object-to-object alignments.

2. SIAM correctly makes time course predictions. SIAM predicts our results that a globally consistent feature match increases similarity more than a local match *late* in processing, whereas the locally consistent feature match increases similarity more than the global match *early* in processing. The more cycles SIAM executes, the more similarity is influenced by the requirement that object mappings be consistent. While two objects may strongly map onto one object at first, the nodes representing these incompatible mappings will strongly inhibit each other. If there is even a weak advantage to one of the mappings, then the other mapping will become increasingly weakened with time. SIAM also correctly predicts that MOPs increase early assessments of similarity more than late assessments, for the same reason. With time, object-to-object correspondences will begin to influence feature-to-feature correspondences, and with this influence, similarity will become selectively influenced by MIPs. The more cycles SIAM completes, the more any given node is activated so as to be consistent with all of the other nodes.

The analogical reasoning simulations that inspired SIAM, SME (Falkenhainer et al., 1990) and ACME (Holyoak & Thagard, 1989), also incorporate a local-to-global process of correspondence resolution. In SME, locally consistent entity correspondences are first computed. These correspondences are later combined into more globally consistent mappings. In SIAM, like ACME, node activations become more influenced by other network activations as more cycles of activation passing transpire.

3. SIAM predicts that nondiagnostic features, if present in two scenes, increase subjects' accuracy in making the correct butterfly-to-butterfly mappings. Mapping accuracy in SIAM is modeled by comparing the magnitudes of object-to-object node activations. If the A-to-C node activation is 0.8 and the A-to-D activation is 0.4, then the probability of a subject mapping A to C is $0.8/(0.8 + 0.4) = 67\%$. The more features (diagnostic and nondiagnostic) that two objects share, the more strongly the objects will be placed in correspondence, and consequently, the more strongly all feature matches shared by the objects will be activated, including the diagnostic feature match. If the two scenes do not agree on the nondiagnostic features, no objects will be placed in strong correspondence, and no substantial level of activation will be fed back to

the diagnostic feature. In this way, SIAM correctly predicts that even features that provide no cue about what objects correspond to each other still can increase mapping accuracy.

4. SIAM predicts that sensitivity (in the signal detection sense of the word) is higher for feature matches and mismatches that occur in aligned objects than unaligned objects. If two objects are placed in strong correspondence, then all of the matching and mismatching features of those objects are made more important for similarity assessments. If we assume that sensitivity to (mis)matching features is a monotonically increasing function of the feature-to-feature node activation, then subjects will be more sensitive to aligned feature (mis)matches. Feature-to-feature node activations are used as indications of how much *attention* a particular correspondence receives; the more a feature-to-feature node is activated, the more the feature-to-feature value influences similarity. The more attention a feature-to-feature (mis)match receives, the greater the sensitivity will be for match/mismatch questions. McClelland and Elman's (1986) interactive activation model of speech perception has received criticism for making a similar claim. Massaro (1989) argues that McClelland and Elman's TRACE model predicts non-constant phoneme sensitivities for different speech contexts (for a more recent version of TRACE that does not make this prediction see McClelland, 1991). No such sensitivity changes are empirically obtained in the spoken word stimuli that have been used. However, in the case of our butterfly scenes, we in fact *do* find sensitivity differences for feature matches depending on the alignments of objects. For our domain, the fact that an interactive activation model predicts context-dependent sensitivity changes is a point in favor of the model.

5.3. Comparisons and Future Directions

The empirical results presented here are problematic for traditional models of similarity that create independent representations of entities and compare these representations for matching and mismatching features (Contrast model), or for proximity on feature values (MDS). Even if entities are represented in terms of simple ("white") and conjunctive ("white spotted triangle") features, many of the more detailed results we obtained would not be predicted. In particular, no account would be given of: the nonlinear effect of MOPs ($2 \text{ MOPs} > 1 \text{ MOP} = 0$ MOPs), the difference between true and false mapping trials on similarity ratings, the relatively large impact of MOPs on similarity *early* in processing, the influence of nondiagnostic features on mapping accuracy, the increased similarity due to MIPs that are concentrated in few dimensions as opposed to distributed across many dimensions, and the influence of alignment on feature sensitivity.


In many ways, the closest neighbors to SIAM are models of analogy such as SME and ACME. Our empirical results were not compared with these models'

predictions in a rigorous manner. SME and ACME were not principally designed to handle similarity data, and there are several issues that arise when they are applied to our results. In particular, if SME's assessment of a pair of scenes is limited to a single best GMAP (a single set of *consistent* correspondences between two structured representations) then SME does not predict that MOPs increase similarity. For a MOP to increase similarity, mutually exclusive correspondences must be integrated into a single estimate of similarity. SME would have to allow for multiple conflicting GMAPs to simultaneously increase similarity (Janice Skorstad and Dedre Gentner are currently working on possibilities along these lines).

Likewise, MOPs do not increase similarity in ACME if similarity is measured as the numbers of "cycles to success, the number of cycles required for all of the correct individual mapping hypotheses to reach an activation level exceeding that of their nearest competitor." Instead, MOPs would decrease this measure of similarity because MOPs tend to activate globally inconsistent object-to-object alignments, and thus slow down convergence to the correct object-to-object mappings. ACME also incorporates a measure G , "a rough index of the overall fit of the emerging mapping to the constraints" imposed by the network. In some senses, G captures some notion of similarity: if the system is relatively "happy" in the correspondences it is setting up between scenes, G will tend to be high. G increases monotonically with the number of cycles of activation adjustment; as ACME settles into a solution, G necessarily increases. While G captures the notion that activations tend to strive for global consistency, it does not necessarily capture the notion of overall similarity. Human assessments of similarity do not always increase with time; in fact, with similarity measured as the number of errors on "different" trials, similarity virtually always decreases with longer display exposure or longer response deadlines. In SIAM, similarity is a function of activations and the feature match values. While a feature-to-feature activation may become increasingly influenced by object correspondences, if the features do not in fact have the same feature value then similarity will still be lowered by this mismatch. Massaro (1989) has argued that interactive activation models are nonoptimal because they distort environmental input more than is reasonable; the integrity of the information sources is compromised by other activations. In SIAM, the feature match values are not subject to modification by activations, and thus the information they provide is never lost by other activations. What changes is the weight given to the match values.

ACME and SME both support a large range of relational/propositional expressions. With SIAM, we have opted for a much simpler, comparatively impoverished model in order to account for the empirical details of a few particular data sets. We have also begun to extend SIAM in the direction of increasing relational complexity. In a more recent version of SIAM (Goldstone, 1991), in addition to feature-to-feature and object-to-object connections, there

Figure 6.13. The comparison of these sport scenes involves alignment on the basis of functional role. The bowling ball has more physical features in common with the bull's-eye, and the arrows have more physical features in common with the bowling pins. However, the intuitive mapping most likely places the bowling ball in correspondence with the arrow, and the pins in correspondence with the target.



are role-to-role nodes that place the abstract roles of one scene into correspondence with the abstract roles of the other scene. With these nodes, objects can be placed in correspondence because of their featural overlap or because they play the same role in two scenes. For example, the question can be raised: What does the bowling ball of Figure 6.13A correspond to in Figure 6.13B? Although the roundness of the ball may prompt aligning it with the bull's-eye of Figure 6.13B, the role of the ball in the abstract description "actor propelling projectile toward target" would dictate mapping the bowling ball of Figure 6.13A onto the arrow of Figure 6.13B. The role-to-role correspondence between the projectiles of the scenes would activate the object-to-object node that places the ball and arrow into correspondence (see Markman & Gentner, 1990, for experimental confirmation of the importance of roles in similarity comparisons).

One of our central claims has been that scene representations are not completely determined prior to the comparison process. We have argued that the weight that a feature match or mismatch has depends on how the two scenes align with one another. However, there is a more fundamental way in which representations may not be precomputable. We have recently obtained results suggesting that the actual features that describe scenes may not be developed independently of the comparison. In Figure 6.14, when A is compared to B, more often than not, a subject will list as a shared feature the fact that both A and B have three prongs/fingers. However, a second group of subjects, asked to compare B and C, more than half of the time list as a shared feature the fact that both B and C have four prongs/fingers. Assuming these groups of subjects have the same initial tendency to see B as three vs. four pronged, we conclude that the featural description of B depends on what it is compared with. Given the




Figure 6.14.

Subjects are asked to compare B with A and C, listing similarities and differences. When B is compared to A, B is most often given the description "three prongs." When B is compared to C, B is most often given the description "four prongs."

implausibility of B having static, precomparison descriptions of both "three prongs" and "four prongs," we conclude that the prong description for B is at least partially created only once it has been paired with A or C. Not only are feature *saliences* influenced by the process of aligning scene parts, but the features themselves are determined by the comparison process as well. SIAM as currently conceived cannot handle this finding, but the finding is consistent with the general theory that entity representations are dynamically constructed on the basis of the actual comparison being made, and are not static or completely prewired.

It might be argued that SIAM's representations are no less prewired than those used by the Contrast model or MDS. SIAM takes as input scene descriptions that are expressed as objects with feature slots and values, and these descriptions do not change with processing. However, what does change with processing is the degree to which two feature values correspond to each other. The salience of a feature (mis)match depends on the alignment of the objects involved, and the objects' alignment cannot be determined until other features and objects are compared. It is for this reason that we claim that the weight or salience of a feature in a similarity assessment is not completely determined before it enters into the comparison.

6. CONCLUSIONS

The experiments and model presented here have pointed to three conclusions:

1. The act of comparing things naturally involves aligning the parts of the things to be compared.
2. Similarity assessments are well captured by an interactive activation process between feature and object correspondences.
3. What counts as a feature match, and how much it will count, depends on the particular things being compared.

Comparison naturally involves alignment. Even when subjects are not instructed to do so, even when indirect measures of similarity are used, subjects in our experiments set up correspondences between the parts of things they compared. These correspondences influence the particular features that are attended, and the other correspondences that are created. Relative to feature matches between noncorresponding objects (MOPs), feature matches between corresponding objects (MIPs) increase similarity ratings, slow down subjects' responses to say scenes are different, increase the proportion of trials in which subjects call two different scenes the same, and influence categorization

decisions. The actual correspondences set up by a subject influence the perceived similarity of two scenes.

Similarity assessments are well captured by an interactive activation process between feature and object correspondences. SIAM is able to capture empirical details not captured even by the model of similarity that makes a conceptual distinction between MIPs and MOPs and differentially weights the two types of matching features. SIAM gives good fits to data collected by subjects, and makes correct predictions with regard to the distribution of feature matches (concentrated MIPs increases similarity more than distributed MIPs), the time course of similarity (MOPs increase similarity more early in processing than later), and sensitivity to features (sensitivity is greater for feature (mis)matches belonging to aligned objects than unaligned objects). The most important insight of SIAM is that there is a relation of mutual dependency between hierarchical levels of a scene. Object correspondences depend on feature-to-feature similarities. The more features two objects have in common, the more strongly the objects will be placed in correspondence. Reciprocally, feature correspondences depend on object-to-object similarities. The greater the similarity between two objects, the more strongly the features of the objects will be placed in correspondence. How much a feature match counts toward similarity depends on whether it matches aligned objects. Feature and object alignments mutually influence each other, and together they determine the similarity of whole scenes. This picture of mutual dependency is in contrast to traditional models of similarity which have little to say concerning how feature matches and mismatches are determined.

What counts as a feature match, and how much it will count, depends on the particular things being compared. We have argued against models of comparison that develop independent representations of the things to be compared, and then evaluate these representations for overlap and proximity of values. Representations are not created independently—the weight that a feature has in a comparison cannot be determined until the feature is brought into alignment with its counterpart in the other scene. We cannot know a priori how much a given feature such as “triangle,” if it matches, will influence similarity. The feature match may increase similarity a great deal, or not very much at all, depending on how the scene’s parts are aligned more globally.

If we begin our analysis of the comparison process by assuming “A has these features, B has these features, and the features have saliences X, Y, and Z” then we may unwisely ignore the most interesting cognitive phenomena involved in comparing things. It may turn out that the cognitive processes most in need of explanation are not those processes responsible for integrating matching and mismatching features into a single estimate of similarity, but are those processes responsible for figuring out exactly what will *count* as matching and mismatching features and how much weight to give these features.

REFERENCES

- Beck, J. (1966). Effect of orientations and of shape similarity on perceptual grouping. *Perception and Psychophysics, 1*, 300–302.
- Carroll, J. D., & Wish, M. (1974). Models and methods for three-way multidimensional scaling. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 2, pp. 57–105). San Francisco: Freeman.
- Clement, C., & Gentner, D. (1988). Systematicity as a selection constraint in analogical mapping. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 421–419). Hillsdale, NJ: Erlbaum.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1990). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41*, 1–63.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. *Cognitive Psychology, 16*, 341–370.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155–170.
- Gentner, D., & 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* (Vol. 22, pp. 307–358). New York: Academic Press.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10*(3), 277–300.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1–39.
- Goldstone, R. L. (1991). *Similarity, interactive activation, and mapping*. Unpublished doctoral dissertation, University of Michigan, Ann Arbor.
- Hall, R. P. (1989). Computational approaches to analogical reasoning: A comparative analysis. *Artificial Intelligence, 39*, 39–120.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constrain satisfaction. *Cognitive Science, 13*, 295–355.
- Keren, G., & Baggen, S. (1981). Recognition model of alphanumeric characters. *Perception and Psychophysics, 29*, 289–294.
- Kolers, P. A. (1972). *Aspects of motion perception*. New York: Pergamon Press.
- Markman, A. B., & Gentner, D. G. (1990). Analogical mapping during similarity judgments. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 38–44).
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Marr, D., & Poggio, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society of London, 204*, 301–328.
- Massaro, D. W. (1989). Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology, 21*, 398–421.
- McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology, 23*, 1–144.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*, 1–86.

- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- Ramachandran, V. S., & Anits, S. M. (1986). Perception of apparent motion. *Scientific American*, 254, 102-109.
- Richardson, M. W. (1938). Multidimensional psychophysics. *Psychological Bulletin*, 35, 659-660.
- Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.
- Satrah, S., & Tversky, A. (1987). On the relation between common and distinctive feature models. *Psychological Review*, 94, 16-22.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81, 214-241.
- Swets, J. A., Tanner, W. P., & Birdsall, T. G. (1961). Decision processes in perception. *Psychological Review*, 68, 301-340.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.
- Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.

7

Connectionist Implications for Processing Capacity Limitations in Analogies

Graeme S. Halford
William H. Wilson
Jian Guo, Ross W. Gayler, Janet Wiles,
and J.E.M. Stewart

1. CONNECTIONIST IMPLICATIONS FOR PROCESSING CAPACITY LIMITATIONS IN ANALOGIES

There is now a reasonable amount of consensus that an analogy entails a mapping from one structure, the base or source, to another structure, the target (Gentner, 1983, 1989; Holyoak & Thagard, 1989). Theories of human analogical reasoning have been reviewed by Gentner (1989), who concludes that there is basic agreement on the one-to-one mapping of elements and the carry over of predicates. Furthermore, as Palmer (1989) points out, some of the theoretical differences represent different levels of description rather than competing models. Despite this consensus about the central role of structure mapping, it really only treats the syntax of analogies, and there are also important pragmatic factors, as has been pointed out by Holland, Holyoak, Nisbett, and Thagard (1986) and Holyoak and Thagard (1989). However in this chapter we are primarily concerned with the problem of how to model the structure mapping or syntactic component of analogical reasoning in terms of parallel distributed processing (PDP) architectures.

According to Gentner (1983), attributes are not normally mapped in analogies, and only certain relations are mapped, the selection being based on systematicity, or the degree to which relations enter into a coherent structure.

* We are grateful to Murray Maybery for very stimulating discussion of some of the issues addressed in this chapter.