

Similarity
Robert L. Goldstone
Ji Yun Son
Indiana University

Correspondence Address: Robert Goldstone
Department of Psychology
Indiana University
Bloomington, IN. 47408

Other Correspondences: rgoldsto@indiana.edu
(812) 855-4853

Introduction

Human assessments of similarity are fundamental to cognition because similarities in the world are revealing. The world is an orderly enough place that similar objects and events tend to behave similarly. This fact of the world is not just a fortunate coincidence. It is because objects are similar that they will tend to behave similarly in most respects. It is because crocodiles and alligators are similar in their external form, internal biology, behavior, diet, and customary environment that one can often successfully generalize from what one knows of one to the other. As Quine (1969) observed, “Similarity, is fundamental for learning, knowledge and thought, for only our sense of similarity allows us to order things into kinds so that these can function as stimulus meanings. Reasonable expectation depends on the similarity of circumstances and on our tendency to expect that similar causes will have similar effects (p. 114).” Similarity thus plays a crucial role in making predictions because similar things usually behave similarly.

From this perspective, psychological assessments of similarity are valuable to the extent that they provide grounds for predicting as many important aspects of our world as possible (Holland, Holyoak, Nisbett, & Thagard, 1986; Dunbar & Fugelsang, Ch. 28). Appreciating the similarity between crocodiles and alligators is helpful because information learned about one is generally true of the other. If we learned an arbitrary fact about crocodiles, such as they are very sensitive to the cold, then we would probably be safe in inferring that this fact is also true of alligators. As the similarity between A and B increases, so does the probability of correctly inferring that B has X upon knowing that A has X (Tenenbaum, 1999). This relation assumes that we have no special knowledge related to property X. Empirically, Heit and Rubinstein (1994) have shown that if we do know about the property, then this knowledge, rather than a one-size-fits-all similarity, is used to guide our inferences. For example, if people are asked to make an inference about an anatomical property, then anatomical similarities have more influence than behavioral similarities. Boars are anatomically but not behaviorally similar to pigs, and this difference successfully predicts that people are likely to make anatomical but not behavioral inferences from pigs to boars. The logical extreme of this line of reasoning (Goodman, 1972; Quine, 1977) is that if one has complete knowledge about the reasons why an object has a property, then general similarity is no longer relevant to generalizations. The knowledge itself completely guides whether the generalization is appropriate. Moonbeams and melons are not very similar generally

speaking, but if one is told that the Moonbeams have the property that the word begins with Melanie's favorite letter, then one can generalize this property to melons with very high confidence.

By contrasting the cases of crocodiles, boars, and moonbeams, we can specify the benefits and limitations of similarity. We tend to rely on similarity to generate inferences and categorize objects into kinds when we do not know exactly what properties are relevant, or when we cannot easily separate an object into separate properties. Similarity is an excellent example of a domain-general source of information. Even when we do not have specific knowledge about a domain, we can use similarity as a default method to reason about it. The contravening limitation of this domain generality is that when specific knowledge is available, then a generic assessment of similarity is no longer as relevant (Keil, 1989; Murphy, 2002; Murphy & Medin, 1985; Rips, 1989; Rips & Collins, 1993). Artificial laboratory experiments where subjects are asked to categorize unfamiliar stimuli into novel categories invented by the experimenter are situations where similarity is clearly important because subjects little else to use (Estes, 1994; Nosofsky, 1984; 1986). However, similarity is also important in many real-world situations because our knowledge does not run as deep as we think it does (Rozenblit & Keil, 2002), and because a general sense of similarity often has an influence even when more specific knowledge ought to overrule it (Allen & Brooks, 1991; Smith & Sloman, 1994).

Another argument for the importance of similarity in cognition is simply that it plays a significant role in psychological accounts of problem solving, memory, prediction, and categorization. If a problem is similar to a previously solved problem, then the solution to the old problem may be applied to the new problem (Holyoak & Koh, 1987; Ross, 1987, 1989). If a cue is similar enough to a stored memory, the memory may be retrieved (Raaijmakers & Shiffrin, 1981). If an event is similar enough to a previously experienced event, the stored event's outcome may be offered as a candidate prediction for the current event (Sloman, 1993; Tenenbaum & Griffiths, 2001). If an unknown object is similar enough to a known object, then the known object's category label may be applied to the unknown object (Nosofsky, 1986). The act of comparing events, objects, and scenes, and establishing similarities between them is of critical importance for the cognitive processes we depend upon.

The utility of similarity for grounding our concepts has been rediscovered in all of the fields comprising cognitive science (see Medin & Rips, Ch. 2). Exemplar (Estes, 1994; Kruschke, 1992; Lamberts, 2000; Medin & Schaffer, 1978; Nosofsky, 1986), instance-based (Aha, 1992), view-based (Tarr & Gauthier, 1998), case-based (Schank, 1982), nearest

neighbor (Ripley, 1996), configural cue (Gluck & Bower, 1990), and vector quantization (Kohonen, 1995) models all share the underlying strategy of giving responses learned from similar, previously presented patterns to novel patterns. Thus, a model can respond to repetitions of these patterns; it can also give responses to novel patterns that are likely to be correct by sampling responses to old patterns, weighted by their similarity to the novel pattern. Consistent with these models, psychological evidence suggests that people show good transfer to new stimuli in perceptual tasks to the extent that the new stimuli resemble previously learned stimuli (Kolers & Roediger, 1984; Palmeri, 1997). Another common feature of these approaches is that they represent patterns in a relatively raw, unprocessed form. This parallels the constraint described above on the applicability of similarity. Both raw representations and generic similarity assessments are most useful as a default strategy when one does not know exactly what properties of a stimulus are important. One's best bet is to follow the principle of least commitment (Marr, 1982) and keep mental descriptions in a relatively raw form to preserve information that may be needed at a later point.

Another reason for studying similarity is that it provides an elegant diagnostic tool for examining the structure of our mental entities and the processes that operate on them. For example, one way to tell that a physicist has progressed beyond the novice stage is that they see deep similarities between problems that require calculation of force even though the problems are superficially dissimilar (Chi, Feltovich, & Glaser, 1981; see Chi & Ohlsson, Ch. 11). Given that psychologists have no microscope with direct access to people's representations of their knowledge, appraisals of similarity provide a powerful, if indirect, lens onto representation/process assemblies (see also Doumas & Hummel, Ch. 20).

A final reason to study similarity is that it occupies an important ground between perceptual constraints and higher-level knowledge system functions. Similarity is grounded by perceptual functions. A tone of 200 hz and a tone of 202 hz sound similar (Shepard, 1987), and the similarity is cognitively impenetrable (Pylyshn, 1985) enough that there is little that can be done to alter this perceived similarity. On the other hand, similarity is also highly flexible and dependent on knowledge and purpose. By focusing on patterns of motion and relations, even electrons and planets can be made to seem similar (Gentner, 1983; Holyoak & Thagard, 1989; see Holyoak, Ch. 4). A complete account of similarity will make contact both with Fodor's (1983) isolated and modularized perceptual input devices and the "central system" where everything a person knows may be relevant.

A Survey of Major Approaches to Similarity

There have been a number of formal treatments that simultaneously provide theoretical accounts of similarity and describe how it can be empirically measured (Hahn, 2003). These models have had a profound practical impact in statistics, automatic pattern recognition by machines, data mining, and marketing (e.g. online stores can provide “people similar to you liked the following other items...”). Our brief survey is organized in terms of the following models: geometric, feature-based, alignment-based, and transformational.

Geometric models and multidimensional scaling

Geometric models of similarity have been among the most influential approaches to analyzing similarity (Carroll & Wish, 1974; Torgerson, 1958, 1965). These approaches are exemplified by nonmetric multidimensional scaling (MDS) models (Shepard, 1962a, 1962b). MDS models represent similarity relations between entities in terms of a geometric model that consists of a set of points embedded in a dimensionally organized metric space. The input to MDS routines may be similarity judgments, dissimilarity judgments, confusion matrices, correlation coefficients, joint probabilities, or any other measure of pairwise proximity. The output of an MDS routine is a geometric model of the data, with each object of the data set represented as a point in an n-dimensional space. The similarity between a pair of objects is taken to be inversely related to the distance between two objects' points in the space. In MDS, the distance between points i and j is typically computed by:

$$dissimilarity(i, j) = \left[\sum_{k=1}^n |X_{ik} - X_{jk}|^r \right]^{\frac{1}{r}} \quad (\text{Eq. 1})$$

where n is the number of dimensions, X_{ik} is the value of dimension k for item i, and r is a parameter that allows different spatial metrics to be used. With $r=2$, a standard Euclidean notion of distance is invoked, whereby the distance between two points is the length of the straight line connecting the points. If $r=1$, then distance involves a city-block metric where the distance between two points is the sum of their distances on each dimension (“short-cut” diagonal paths are not allowed to directly connect points differing on more than one dimension). A Euclidean metric often provides a better fit to empirical data when the stimuli being compared are composed of integral, perceptually fused dimensions such as the

brightness and saturation of a color. Conversely, a city-block metric is often appropriate for psychologically separated dimensions such as brightness and size (Attneave, 1950).

Richardson's (1938) fundamental insight, which is the basis of contemporary use of MDS, was to begin with subjects' judgments of pair-wise object dissimilarity, and work backward to determine the dimensions and dimension values that subjects used in making their judgments. MDS algorithms proceed by placing entities in an N-dimensional space such that the distances between the entities accurately reflects the empirically observed similarities. For example, if we asked people to rate the similarities [on a scale from 1 (low similarity) to 10 (high similarity)] of Russia, Cuba, and Jamaica, we might find:

Similarity (Russia, Cuba) = 7

Similarity (Russia, Jamaica) = 1

Similarity (Cuba, Jamaica) = 8

An MDS algorithm would try to position the three countries in a space such that countries that are rated as being highly similar are very close to each other in the space. With nonmetric scaling techniques only ordinal similarity relations are preserved. The interpoint distances suggested by the similarity ratings may not be simultaneously satisfiable in a given dimensional space. If we limit ourselves to a single dimension (we place the countries on a "number line"), then we cannot simultaneously place Russia near Cuba (similarity = 7) and place Russia far away from Jamaica (similarity = 1). In MDS terms, the "stress" of the one-dimensional solution would be high. We could increase the dimensionality of our solution and position the points in two-dimensional space. A perfect reconstruction of any set of proximities among a set of N objects can be obtained if a high enough dimensionality (specifically, N-1 dimensions) is used.

One of the main applications of MDS is to determine the underlying dimensions comprising the set of compared objects. Once the points are positioned in a way that faithfully mirrors the subjectively obtained similarities, it is often possible to give interpretations to the axes, or to rotations of the axes. In the example above, dimensions may correspond to "political affiliation" and "climate." Russia and Cuba would have similar values on the former dimension; Jamaica and Cuba would have similar values on the latter dimension. 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 suggested that birds were represented in terms of their values on dimensions such as “ferocity” and “size.” It is important to note that the proper psychological interpretation of a geometric representation of objects is not necessarily in terms of its Cartesian axes. In some domains, such as musical pitches, the best interpretation of objects may be in terms of their polar coordinates of angle and length (Shepard, 1982). Recent work has extended geometric representations still further, representing patterns of similarities by generalized, nonlinear manifolds (Tenenbaum, De Silva, & Lanford, 2000).

MDS is also used to create a compressed representation that conveys relational similarities among a set of items. A set of N items requires $N(N-1)/2$ numbers to express all pairwise distances among the items, assuming that any object has a distance of 0 to itself and distances are symmetric. However, if an MDS solution fits the distance data well, it can allow these same distances to be reconstructed using only ND numbers, where D is the number of dimensions of the MDS solution. This compression may be psychologically very useful. One of the main goals of psychological representation is to create efficient codes for representing a set of objects. Compressed representations can facilitate encoding, memory, and processing. Shimon Edelman (1999) has proposed that both people and machines efficiently code their world by creating geometric spaces for objects with much lower dimensionality than the objects’ physical description (see also Gardenfors, 2000).

A third use of MDS is to create quantitative representations that can be used in mathematical and computational models of cognitive processes. Numeric representations, namely coordinates in a psychological space, can be derived for stories, pictures, sounds, words, or any other stimuli for which one can obtain subjective similarity data. Once constructed, these numeric representations can be used to predict people’s categorization accuracy, memory performance, or learning speed. MDS models have been successful in expressing cognitive structures in stimulus domains as far removed as animals (Smith, Shoben, & Rips, 1974), Rorschach ink blots (Osterholm, Woods, & Le Unes, 1985), chess positions (Horgan, Millis, & Meimeyer, 1989), and air flight scenarios (Schvaneveldt, 1985). Many objects, situations, and concepts seem to be psychologically structured in terms of dimensions, and a geometric interpretation of the dimensional organization captures a substantial amount of that structure.

Featural Models

In 1977, Amos Tversky brought into prominence what would become the main contender to geometric models of similarity in psychology. The reason given for proposing a feature-based model was that subjective assessments of similarity did not always satisfy the assumptions of geometric models of similarity.

Problems with the Standard Geometric Model. Three assumptions of standard geometric models of similarity are:

Minimality: $D(A,B) \geq D(A,A)=0$

Symmetry: $D(A,B)=D(B,A)$

The Triangle Inequality: $D(A,B)+D(B,C) \geq D(A,C)$

where $D(A,B)$ is interpreted as the dissimilarity between items A and B.

According to the minimality assumption, all objects are equally (dis)similar to themselves. Some violations of this assumption are found (Nickerson, 1972) when confusion rates or RT measures of similarity are used. First, not all letters are equally similar to themselves. For example, in Podgorny and Garner (1979) if the letter S is shown twice on a screen, subjects are faster to correctly say that the two tokens are similar (i.e. they come from the same similarity-defined cluster) than if the twice-shown letter is W. By the reaction time measure of similarity, the letter S is more similar to itself than the letter W is to itself. Even more troublesome for the minimality assumption, two different letters may be more similar to each other than a particular letter is to itself. The letter C is more similar to the letter O than W is to itself, as measured by inter-letter confusions. In Gilmore, Hersh, Caramazza, and Griffin (1979), the letter M is more often recognized as an H (probability =.391) than as an M (probability =.180). This is problematic for geometric representations because the distance between a point and itself should be zero.

According to the symmetry assumption, (dis)similarity should not be affected by the ordering of items because the distance from Point A to B is equal to the distance from B to A. Contrary to this presumed symmetry, similarity is asymmetric on occasion (Tversky, 1977). In one of Tversky's examples, North Korea is judged to be more similar to Red China than Red China is to North Korea. Often, a non-prominent item is more similar to a prominent item than vice versa. This is consistent with the result that people judge their friends to be more similar to themselves than they themselves are to their friends (Holyoak & Gordon, 1983), under the assumption that a person is highly prominent to him/herself. More recently, Polk et al. (2002) found that when the frequency of colors was experimentally manipulated,

rare colors are judged to be more similar to common colors than common colors are to rare colors.

According to the triangle inequality assumption, the distance/dissimilarity between two points A and B cannot be more than the distance between A and a third point C plus the distance between C and B. Geometrically speaking, a straight line connecting two points is the shortest path between the points. Tversky and Gati (1982) find violations of this assumption when it is combined with an assumption of segmental additivity ($D(A,B)+D(B,C)=D(A,C)$ if A, B, and C lie on a straight line). Consider three items in multidimensional space, A, B, and C, falling on a straight line such that B is between A and C. Also consider a fourth point, E, that forms a right triangle when combined with A and C. The triangle inequality assumption *cum* segmental additivity predicts that

$$D(A,E) \geq D(A,B) \text{ and } D(E,C) \geq D(B,C)$$

or

$$D(A,E) \geq D(B,C) \text{ and } D(E,C) \geq D(A,B)$$

Systematic violations of this prediction are found such that the path going through the corner point E is shorter than the path going through the center point B. For example, if the items are instantiated as:

A= White, three inches

B= Pink, four inches

C= Red, five inches

E= Red, three inches

Insert Figure 1 about here

then people's dissimilarity ratings indicate that $D(A,E) < D(A,B)$ and $D(E,C) < D(B,C)$. Such an effect can be modeled by geometric models of similarity if r in Eq. 1 is given a value less than 1. However, if r is less than one, then dissimilarity does not satisfy a power metric, often times considered a minimal assumption for geometric solutions to be interpretable. The two assumptions of a power metric are 1) distances along straight lines are additive, and 2) the shortest path between points is a straight line.

Other potential problems with geometric models of similarity are 1) they strictly limit the number of nearest neighbors an item can have (Tversky & Hutchinson, 1986), 2) multidimensional scaling techniques have difficulty describing items that vary on a large number of features (Krumhansl, 1978), and 3) standard MDS techniques do not predict that adding common features to items increases their similarity (Tversky & Gati, 1982). On the

first point: MDS models consisting of two dimensions cannot predict that item X is the closest item to 100 other items. There would be no way of placing those 100 items in two dimensions such that X is closer to all of them than any other item is. For human data, a superordinate term (e.g. fruit) is often the nearest neighbor of many of its exemplars (apples, bananas, etc.) as measured by similarity ratings. On the second point: although there is no logical reason why geometric models cannot represent items of any number of dimensions (as long as the number of dimensions is less than number of items minus one), geometric models tend to yield the most satisfactory and interpretable solutions in low-dimensional space. MDS solutions involving more than six dimensions are rare. On the third point: the addition of the same feature to a pair of items increases their rated similarity (Gati & Tversky, 1984), but this is incompatible with simple MDS models. If adding a shared feature corresponds to adding a dimension in which the two items under consideration have the same value then there will be no change to the items' dissimilarity because the geometric distance between the points remain the same. MDS models that incorporate the dimensionality of the space could predict the influence of shared features on similarity, but such a model would no longer relate similarity directly to an inverse function of interitem distance.

One research strategy has been to augment geometrical models of similarity in ways that solve these problems. One solution, suggested by Carol Krumhansl (1978), has been to model dissimilarity in terms of both interitem distance in a multidimensional space and spatial density in the neighborhoods of the compared items. The more items there are in the vicinity of an item, the greater the spatial density of the item is. Items are more dissimilar if they have many items surrounding them (their spatial density is high) than if they have few neighboring items. By including spatial density in an MDS analysis, violations of minimality, symmetry, and the triangle inequality can potentially be accounted for, as well as some of the influence of context on similarity. However, the empirical validity of the spatial density hypothesis is in some doubt (Corter, 1987, 1988; Krumhansl, 1988; Tversky & Gati, 1982).

Robert Nosofsky (1991) has suggested another potential way to save MDS models from some of the above criticisms. He introduces individual bias parameters in addition to the inter-item relation term. Similarity is modeled in terms of inter-item distance and in terms of biases toward particular items. Biases toward items may be due to attention, salience, knowledge, and frequency of items. This revision handles asymmetric similarity results and the result that a single item may be the most similar item to many other items, but does not directly address several of the other objections.

The Contrast Model. In light of the above potential problems for geometric representations, Tversky (1977) proposed to characterize similarity in terms of a feature-matching process based on weighting common and distinctive features. In this model, entities are represented as a collection of features and similarity is computed by:

$$S(A,B) = qf(A \ll B) - af(A-B) - bf(B-A)$$

The similarity of A to B is expressed as a linear combination of the measure of the common and distinctive features. The term $(A \ll B)$ represents the features that items A and B have in common. $(A-B)$ represents the features that A has but B does not. $(B-A)$ represents the features of B that are not in A. q , a , and b are weights for the common and distinctive components. Common features as compared to distinctive features, are given relatively more weight for verbal as opposed to pictorial stimuli (Gati & Tversky, 1984), for cohesive as opposed to non-cohesive stimuli (Ritov, Gati, & Tversky, 1990), for similarity as opposed to difference judgments (Tversky, 1977), and for entities with a large number of distinctive as opposed to common features (Gati & Tversky, 1984). There are no restrictions on what may constitute a feature. A feature may be any property, characteristic or aspect of a stimulus. Features may be concrete or abstract (i.e. “symmetric” or “beautiful”).

The Contrast Model predicts asymmetric similarity because a is not constrained to equal b and $f(A - B)$ may not equal $f(B - A)$. North Korea is predicted to be more similar to Red China than vice versa if Red China has more salient distinctive features than North Korea, and a is greater than b . The Contrast Model can also account for non-mirroring between similarity and difference judgments. The common features term $(A \ll B)$ is hypothesized to receive more weight in similarity than difference judgments; the distinctive features term receives relatively more weight in difference judgments. As a result, certain pairs of stimuli may be perceived as simultaneously being more similar to and more different from each other, compared to other pairs (Tversky, 1977). Sixty-seven percent of a group of subjects selected West Germany and East Germany as more similar to each other than Ceylon and Nepal. Seventy percent of subjects also selected West Germany and East Germany as more different from each other than Ceylon and Nepal. According to Tversky, East and West Germany have more common and more distinctive features than Ceylon and Nepal. Medin, Goldstone, and Gentner (1993) present additional evidence for non-mirroring between similarity and difference, exemplified in Figure 2. When two scenes share a relatively large number of relational commonalities (e.g. Scenes T and B both have 3 objects that have the same pattern) but also a large number of differences on specific attributes (e.g. none of the

patterns in Scene T match any of the patterns in B), then the scenes tend to be judged as simultaneously very similar and very different.

A number of models are similar to the Contrast model in basing similarity on features and in using some combination of the (A«B), (A-B), and (B-A) components. Sjoberg (1972) proposes that similarity is defined as $f(A«B)/f(A»B)$. Eisler and Ekman (1959) claim that similarity is proportional to $f(A«B)/(f(A)+f(B))$. Bush and Mosteller (1951) defines similarity as $f(A«B)/f(A)$. These three models can all be considered specializations of the general equation $f(A«B)/[f(A«B)+af(A-B)+bf(B-A)]$. As such, they differ from the Contrast model by applying a ratio function as opposed to a linear contrast of common and distinctive features.

The fundamental premise of the Contrast Model, that entities can be described in terms of constituent features, is a powerful idea in cognitive psychology. Featural analyses have proliferated in domains of speech perception (Jakobson, Fant, & Halle, 1963), pattern recognition (Neisser, 1967; Treisman, 1986), perception physiology (Hubel & Wiesel, 1968), semantic content (Katz & Fodor, 1963), and categorization (Medin & Shaffer, 1978). Neural network representations are often based on features, with entities being broken down into a vector of ones and zeros, where each bit refers to a feature or “micro-feature.” Similarity plays a crucial role in many connectionist theories of generalization, concept formation, and learning. The notion of dissimilarity used in these systems is typically the fairly simple function “Hamming distance.” The Hamming distance between two strings is simply their city-block distance; that is, it is their (A - B) + (B - A) term. “1 0 0 1 1” and “1 1 1 1 1” would have a Hamming distance of 2 because they differ on two bits. Occasionally, more sophisticated measures of similarity in neural networks normalize dissimilarities by string length. Normalized Hamming distance functions can be expressed by $[(A-B) + (B-A)]/[f(A«B)]$.

Similarities Between Geometric and Feature-based Models. While MDS and featural models are often analyzed in terms of their differences, they also share a number of similarities. Recent progress has been made on combining both representations into a single model, using Bayesian statistics to determine whether a given source of variation is more efficiently represented as a feature or dimension (Navarro & Lee, in press). Tversky and Gati (1982) described methods of translating continuous dimensions into featural representations. Dimensions that are sensibly described as being more or less (e.g. loud is more sound than soft, bright is more light than dim, and large is more size than small) can be represented by sequences of nested feature sets. That is, the features of B include a subset of A’s features

whenever B is louder, brighter, or larger than A. Alternatively, for qualitative attributes like shape or hue (red is not subjectively “more” than blue), dimensions can be represented by chains of features such that if B is between A and C on the dimension, then $(A \ll B) \dots (A \ll C)$ and $(B \ll C) \dots (A \ll C)$. For example, if orange lies between red and yellow on the hue dimension, then this can be featurally represented if orange and red share features that orange and yellow do not share.

An important attribute of MDS models is that they create postulated representations, namely dimensions, that explain the systematicities present in a set of similarity data. This is a classic use of abductive reasoning; dimensional representations are hypothesized that, if they were to exist, would give rise to the obtained similarity data. Other computational techniques share with MDS the goal of discovering the underlying descriptions for items of interest, but create featural rather than dimensional representations. Hierarchical Cluster Analysis, like MDS, takes pairwise proximity data as input. Rather than output a geometric space with objects as points, Hierarchical Cluster Analysis outputs an inverted-tree diagram, with items at the root-level connected with branches. The smaller the branching distance between two items, the more similar they are. Just as the dimensional axes of MDS solutions are given subjective interpretations, the branches are also given interpretations. For example, in Shepard's (1972) analysis of speech sounds, one branch is interpreted as voiced phonemes while another branch contains the unvoiced phonemes. In additive cluster analysis (Shepard & Arable, 1979) similarity data is transformed into a set of overlapping item clusters. Items that are highly similar will tend to belong to the same clusters. Each cluster can be considered as a feature. Recent progress has been made on efficient and mathematically principled models that find such featural representations for large databases (Lee, 2002a, 2002b; Tenenbaum, 1996).

Another commonality between geometric and featural representations, one that motivates the next major class of similarity models that we consider, is that both use relatively unstructured representations. Entities are structured as sets of features or dimensions with no relations between these attributes. Entities such as stories, sentences, natural objects, words, scientific theories, landscapes, and faces are not simply a “grab bag” of attributes. Two kinds of structure seem particularly important: propositional and hierarchical. A proposition is an assertion about the relation between informational entities (Palmer, 1975). For example, relations in a visual domain might include *Above*, *Near*, *Right*, *Inside*, and *Larger-than* that take informational entities as arguments. The informational entities might include features

such as square, and values on dimensions such as 3 inches. Propositions are defined as the smallest unit of knowledge that can stand as a separate assertion and have a truth value. The order of the arguments in the predicate is critical. For example, *Above (Triangle, Circle)* does not represent the same fact as *Above (Circle, Triangle)*. Hierarchical representations involve entities that are embedded in one another. Hierarchical representations are required to represent the fact that X is part of Y or that X is a kind of Y. For example, in Collins and Quillian's (1969) propositional networks, labeled links ("Is-a" links) stand for the hierarchical relation between *Canary* and *Bird*.

[Insert Figure 2 about here.]

Some quick fixes to geometric and featural accounts of similarity are possible, but fall short of a truly general capacity to handle structured inputs. Hierarchical clustering does create trees of features, but there is no guarantee that there are relationships, such as Is-a or Part-of, between the subtrees. However, structure might exist in terms of features that represent conjunctions of properties. For example, using the materials in Figure 2, 20 undergraduates were shown triads consisting of A, B, and T, and were asked to say whether Scene A or B was more similar to T. The strong tendency to choose A over B in the first panel suggests that the feature "square" influences similarity. Other choices indicated that subjects also based similarity judgments on the spatial locations and shadings of objects as well as their shapes.

However, it is not sufficient to represent the left-most object of T as {Left, Square, Black} and base similarity on the number of shared and distinctive features. In the second panel, A is again judged to be more similar to T than is B. Both A and B have the features "Black" and "Square." The only difference is that for A and T, but not B, the "Black" and "Square" features belong to the same object. This is only compatible with feature set representations if we include the possibility of conjunctive features in addition to simple features such as "Black" and "Square" (Gluck, 1991; Hayes-Roth & Hayes-Roth, 1977). By including the conjunctive feature "Black-Square," possessed by both T and A, we can explain, using feature sets, why T is more similar to A than B. The third panel demonstrates the need for a "Black-Left" feature, and other data indicates a need for a "Square-Left" feature. Altogether, if we wish to explain the similarity judgments that people make we need a feature

set representation that includes six features (three simple and three complex) to represent the square of T.

However, there are two objects in T, bringing the total number of features required to at least two times the six features required for one object. The number of features required increases still further if we include feature-triplets such as "Left-Black-Square." In general, if there are O objects in a scene, and each object has F features, then there will be OF simple features. There will be O conjunctive features that combine two simple features (i.e. pair-wise conjunctive features). If we limit ourselves to simple and pairwise features to explain the pattern of similarity judgments in Figure 2, we still will require $OF(F+1)/2$ features per scene, or $OF(F+1)$ features for two scenes that are compared to one another.

Thus, featural approaches to similarity require a fairly large number of features to represent scenes that are organized into parts. Similar problems exist for dimensional accounts of similarity. The situation for these models becomes much worse when we consider that similarity is also influenced by relations between features such as "Black to the left of white" and "square to the left of white." Considering only binary relations, there are O^2F^2R -OFR relations within a scene that contains O objects, F features per object, and R different types of relations between features. Although more sophisticated objections have been raised about these approaches by John Hummel and colleagues (Hummel, 2000, 2001; Hummel & Biederman, 1992; Hummel & Holyoak, 1997, 2003; Holyoak & Hummel, 2000; Dumas & Hummel, Ch. 20), at the very least, geometric and featural models apparently require an implausibly large number of attributes to account for the similarity relations between structured, multi-part scenes.

Alignment-based Models

Partly in response to the difficulties that the previous models have in dealing with structured descriptions, a number of researchers have developed alignment-based models of similarity. In these models, comparison is not just matching features, but determining how elements correspond to, or align with, one another. Matching features are aligned to the extent that they play similar roles within their entities. For example, a car with a green wheel and a truck with a green hood both share the feature *green*, but this matching feature may not increase their similarity much because the car's wheel does not correspond to the truck's hood. Drawing inspiration from work on analogical reasoning (Gentner, 1983; Holyoak & Thagard, 1995; Holyoak, Ch. 4), in alignment-based models, matching features influence

similarity more if they belong to parts that are placed in correspondence and parts tend to be placed in correspondence if they have many features in common and are consistent with other emerging correspondences (Goldstone, 1994a; Markman & Gentner, 1993a). Alignment-based models make purely relational similarity possible (Falkenhainer, Forbus, & Gentner, 1989).

Initial evidence that similarity involves aligning scene descriptions comes from Markman and Gentner's (1993a) result that when subjects are asked to determine corresponding objects, they tend to make more structurally sound choices when they have first judged the similarity of the scenes that contain the objects. For example, in Figure 3, subjects could be asked which object in the bottom set corresponds to the left-most object in the top set. Subjects who had rated the similarity of the sets were more likely to choose the right-most object, presumably because of both objects were the smallest objects in their sets. Subjects who did not first assess similarity had a tendency to select the middle object because its size exactly matched the target object's size. These results are predicted if similarity judgments naturally entail aligning the elements of two scenes. Additional research has found that relational choices such as "smallest object in its set" tend to influence similarity judgments more than absolute attributes like "3 inches" when the overall amount of relational coherency across sets is high (Goldstone, Medin, & Gentner, 1991), the scenes are superficially sparse rather than rich (Gentner & Rattermann, 1991; Markman & Gentner, 1993a), subjects are given more time to make their judgments (Goldstone & Medin, 1994), the judges are adults rather than children (Gentner & Toupin, 1986), and abstract relations are initially correlated with concrete relations (Kotovsky & Gentner, 1996).

Formal models of alignment-based similarity have been developed to explain how feature matches that belong to well-aligned elements matter more for similarity than matches between poorly aligned elements. (Goldstone, 1994a; Love, 2000). Inspired by work in analogical reasoning (Holyoak & Thagard, 1989), Goldstone's (1994a) SIAM model is a neural network with nodes that represent hypotheses that elements across two scenes correspond to one another. SIAM works by first creating correspondences between the features of scenes. 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 put in correspondence, activation is fed back down to the feature (mis)matches that are consistent with the object alignments. In this way, object correspondences influence activation of feature correspondences at the same time that feature

correspondences influence the activation of object correspondences. Activation between nodes spreads in SIAM by two principles: 1) nodes that are consistent send excitatory activation to each other and 2) nodes that are inconsistent inhibit each another (see also Holyoak, Ch. 4). Nodes are inconsistent if they create two-to-one alignments -- if two elements from one scene would be placed into correspondence with one element of the other scene. Node activations affect similarity via the equation

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

where n is the number of nodes in the system, A_i is the activation of node i and the match value describes the physical similarity between the two features placed in correspondence according to the node i . By the equation above, the influence of a particular matching or mismatching feature across two scenes is modulated by the degree to which the features have been placed in alignment. Consistent with SIAM, 1) aligned features matches tend to increase similarity more than unaligned feature matches (Goldstone, 1994a), 2) the differential influence between aligned and unaligned feature matches increases as a function of processing time (Goldstone & Medin, 1994), 3) this same differential influences increases with the clarity of the alignments (Goldstone, 1994a), and 4) under some circumstances, adding a poorly aligned feature match can actually decrease similarity by interfering with the development of proper alignments (Goldstone, 1996).

Another empirically validated set of predictions stemming from an alignment-based approach to similarity concerns alignable and non-alignable differences (Markman & Gentner, 1993b). Non-alignable differences between two entities are attributes of one entity that have no corresponding attribute in the other entity. Alignable differences are differences that require that the elements of the entities first be placed in correspondence. When comparing a police car to an ambulance, a non-alignable difference is that police cars have weapons in them, but ambulances do not. There is no clear equivalent of weapons in the ambulance. Alignable differences include the following: police cars carry criminals to jails rather than carrying sick people to hospitals, a police car is a car while ambulances are vans, and police car drivers are policemen rather than emergency medical technicians. Consistent with the role of structural alignment in similarity comparisons, alignable differences influence similarity more than non-alignable differences do (Markman & Gentner, 1996), and are more

likely to be encoded in memory (Markman & Gentner, 1997). Alignable differences between objects also play a disproportionately large role in distinguishing between different basic-level categories (e.g. cats and dogs) that belong to the same superordinate category (e.g. animals) (Markman & Wisniewski, 1997). In short, knowing these correspondences affects not only how much a matching element increases similarity (Goldstone, 1994a), but also how much a mismatching element decreases similarity.

Thus far, much of the evidence for structural alignment in similarity has used somewhat artificial materials. Often times, the systems describe how “scenes” are compared, with the underlying implication that the elements comprising the scenes are not so tightly connected together as elements comprising objects. Still, if the structural alignment account proves to be fertile, it will be because it is applicable to naturally occurring materials. Toward this goal, researchers have considered structural accounts of similarity in language domains. The confusability of words depends on structural analyses to predict that “stop” is more confusable with “step” than “pest” (the “st” match is in the correct location with “step” but not “pest”), but more confusable with “pest” than “best” (the “p” match counts for something even when it is out of place). Substantial success has been made on the practical problem of determining the structural similarity of words (Bernstein, Demorest, & Eberhardt, 1994; Frisch, Broe, & Pierrehumbert, 1995). Structural alignment has also been implicated when comparing more complex language structures such as sentences (Bassok & Medin, 1997). Likewise, structural similarity has proven to be a useful notion in explaining consumer preferences of commercial products, explaining, for example, why new products are viewed more favorably when they improve over existing products along alignable rather than unalignable differences (Zhang & Markman, 1998). Additional research has shown that alignment-based models of similarity provide a better account of category-based induction than feature-based models (Lassaline, 1996). Still other researchers have applied structural accounts of similarity to the legal domain (Hahn & Chater, 1998; Simon & Holyoak, 2002). This area of application is promising because the U.S. legal system is based on cases and precedents, and cases are structurally rich and complex situations involving many interrelated parties. Retrieving a historical precedent, and assessing its relevance to a current case, almost certainly involves aligning representations that are more sophisticated than assumed by geometric or featural models.

Transformational Models

A final historic approach to similarity that has been recently resuscitated is that the comparison process proceeds by transforming one representation into the other. A critical step for these models is to specify what transformational operations are possible.

In an early incarnation of a transformational approach to cognition broadly construed, Garner (1974) stressed the notion of stimuli that are transformationally equivalent and are consequently possible alternatives for each other. In artificial intelligence, Shimon Ullman (1996) has argued that objects are recognized by being aligned with memorized pictorial descriptions. Once an unknown object has been aligned with all candidate models, the best match to the viewed object is selected. The alignment operations rotate, scale, translate, and topographically warp object descriptions. For rigid transformations, full alignment can be obtained by aligning three points on the object with three points on the model description. Unlike recognition strategies that require structural descriptions (e.g. Biederman, 1987, Hummel, 2000, 2001), Ullman's alignment does not require an image to be decomposed into parts.

In transformational accounts that are explicitly designed to model similarity data, similarity is usually defined in terms of transformational distance. In Wiener-Ehrlich, Bart, and Millward's (1980) generative representation system, subjects are assumed to possess an elementary set of transformations, and invoke these transformations when analyzing stimuli. Their subjects saw linear pairs of stimuli such as {ABCD,DABC} or two-dimensional stimuli such as $\begin{Bmatrix} AB & DA \\ CD & BC \end{Bmatrix}$. Subjects were required to rate the similarity of the pairs. The researchers determined transformations that accounted for each subjects' ratings from the set {rotate 90 degrees, rotate 180, rotate 270, horizontal reflection, vertical reflection, positive diagonal reflection, negative diagonal reflection}. Similarity was assumed to decrease monotonically as the number of transformations required to make one sequence identical to the other increased.

Imai (1977) makes a similar claim. The stimuli used were sequences such as XXOXXXOXXXOX where X's represent white ovals, and O's represent black ovals. The four basic transformations were mirror image (XXXXXOO->OOXXXXX), phase shift (XXXXXOO->XXXXOOX), reversal (XXXXXOO->OOOOOXX), and wave length (XXOOXXOO->XOXOXOXO). The researcher found that sequences that are two transformations removed (e.g. XXXOXXXOXXXO and OOXOOOXOOOXO require a phase shift and a reversal to be equated) are rated to be less similar than sequences that can be

made identical with one transformation. In addition, sequences that can be made identical by more than one transformation (XOXOXOXO and OXOXOXOX can be made identical by either mirror image, phase shift, or reversal transformations) are more similar than sequences that have only one identity-producing transformation.

Recent work has followed up on Imai's research and has generalized it to stimulus materials including arrangements of Lego bricks, geometric complexes, and sets of colored circles (Hahn, Chater, & Richardson, 2003). According to these researchers' account, the similarity between two entities is a function of the complexity required to transform the representation of one into the representation of the other. The simpler the transformation, the more similar they are assumed to be. The complexity of a transformation is determined in accord with Kolmogorov complexity theory (Li & Vitanyi, 1997), according to which the complexity of a representation is the length of the shortest computer program that can generate that representation. For example, the conditional Kolmogorov complexity between the sequence 1 2 3 4 5 6 7 8 and 2 3 4 5 6 7 8 9 is small, because the simple instructions add 1 to each digit and subtract 1 from each digit suffice to transform one into the other. Experiments by Hahn et al. demonstrate that once reasonable vocabularies of transformation are postulated, transformational complexity does indeed predict subjective similarity ratings.

It is useful to compare and contrast alignment-based and transformational accounts of similarity. Both approaches place scene elements into correspondence. Whereas the correspondences are explicitly stated in the structural alignment method, they are implicit in transformational alignment. The transformational account often does produce globally consistent correspondences, for example correspondences that obey the one-to-one mapping principle, but this consistency is a consequent of applying a pattern-wide transformation and is not enforced by interactions between emerging correspondences. It is revealing that transformational accounts have been applied almost exclusively to perceptual stimuli, while structural accounts are often most often applied to conceptual stimuli such as stories, proverbs, and scientific theories (there are also notable structural accounts in perception, i.e. Marr & Nishihara, 1978; Biederman, 1987; Hummel & Biederman, 1992; Hummel, 2000). Defining a set of constrained transformations is much more tenable for perceptual stimuli. The conceptual similarity between an atom and the solar system could possibly be discovered by transformations. As a start, a minaturization transformation could be applied to the solar system. However, this single transformation is not nearly sufficient; a nucleus is not simply a small sun. The transformations that would turn the solar system into an atom are not readily

forthcoming. If we allow transformations such as an “earth-becomes-electron” transformation, then we are simply re-expressing the structural alignment approach and its part-by-part alignment of relations and objects.

Some similarity phenomena that are well explained by structural alignment are not easily handled by transformations. To account for the similarity of “BCDCB” and “ABCDCBA” we could introduce the fairly abstract transformation “Add the left-most letter’s predecessor to both sides of string.” However, the pair “LMN” and “KLMNK” do not seem as similar as the earlier pair, even though the same transformation is applied. A transformation of the form “if the structure is symmetric, then add the preceding element in the series to both ends of the string” presupposes exactly the kind of analysis in defining “symmetric” and “preceding” that are the bread and butter of propositional representations and structural alignment. For this reason, one fertile research direction would be to combine alignment-based accounts’ focus on representing the internal structure within individual scenes with the constraints that transformational accounts provide for establishing psychologically plausible transformations (Hofstadter, 1997; Mitchell, 1993).

Conclusions and Further Directions

To provide a partial balance to our largely historical focus on similarity, we will conclude by raising some unanswered questions for the field. These questions are rooted in a desire to connect the study of similarity to cognition as a whole.

Is similarity flexible enough to provide useful explanations of cognition?

The study of similarity is typically justified by the argument that so many theories in cognition depend upon similarity as a theoretical construct. An account of what make problems, memories, objects, and words similar to one another often provides the backbone for our theories of problem solving, attention, perception, and cognition. As William James put it, “This sense of Sameness is the very keel and backbone of our thinking” (James, 1890/1950; p. 459).

However, others have argued that similarity is not flexible enough to provide a sufficient account, although it may be a necessary component. There have been many empirical demonstrations of apparent dissociations between similarity and other cognitive processes, most notably categorization. Researchers have argued that cognition is frequently based on theories (Murphy & Medin, 1985), rules (Smith & Sloman, 1994; Sloman, 1996), or

strategies that go beyond “mere” similarity. To take an example from Murphy and Medin (1985), consider a man jumping into a swimming pool fully clothed. This man may be categorized as drunk because we have a theory of behavior and inebriation that explains the man’s action. Murphy and Medin argue that the categorization of the man’s behavior does not depend on matching the man’s features to the category *drunk*’s features. It is highly unlikely that the category *drunk* would have such a specific feature as “jumps into pools fully clothed.” It is not the similarity between the instance and the category that determines the instance’s classification; it is the fact that our category provides a theory that explains the behavior.

Developmental psychologists have argued that even young children have inchoate theories that allow them to go beyond superficial similarities in creating categories (Carey, 1985; Gelman & Markman, 1986; Keil, 1989). For example, Carey (1985) observes that children choose a toy monkey over a worm as being more similar to a human, but that when they are told that humans have spleens, are more likely to infer that the worm has a spleen than that the toy monkey does. Thus, the categorization of objects into “spleen” and “no spleen” groups does not appear to depend on the same knowledge that guides similarity judgments. Adults show similar dissociations between similarity and categorization. In an experiment by Rips (1989), an animal that is transformed (by toxic waste) from a bird into something that looks like an insect is judged by subjects to be more similar to an insect, but is still judged to be a bird. Again, the category judgment seems to depend on biological, genetic, and historical knowledge, while the similarity judgments seems to depend more on gross visual appearance (see also Keil, 1989; Rips & Collins, 1993).

Despite the growing body of evidence that similarity appraisals do not always track categorization decisions, there are still some reasons to be sanguine about the continued explanatory relevance of similarity. Categorization itself may not be completely flexible. People are influenced by similarity despite the subjects’ intentions and the experimenters’ instructions (Smith & Sloman, 1994). Allen and Brooks (1991) gave subjects an easy rule for categorizing cartoon animals into two groups. Subjects were then transferred to the animals that looked very similar to one of the training stimuli, but belonged in a different category. These animals were categorized more slowly and less accurately than animals that were equally similar to an old animal but also belonged in the same category as the old animal. Likewise, Palmeri (1997) showed that even for the simple task of counting the number of dots, subjects’ performance is improved when a pattern of dots is similar to a previously seen

pattern with the same numerosity and worse when the pattern is similar to a previously seen pattern with different numerosity. People seem to have difficulties ignoring similarities between old and new patterns, even when they know a straightforward and perfectly accurate categorization rule.

There may be a mandatory consideration of similarity in many categorization judgments (Goldstone, 1994b), adding constraints to categorization. At the same time, similarity may be more flexible and sophisticated than commonly acknowledged (Jones & Smith, 1993) and this may also serve to bridge the gap between similarity and high-level cognition. Krumhansl (1978) argued that similarity between objects decreases when they are surrounded by many close neighbors that were also presented on previous trials (also see Wedell, 1994). Tversky (1977) obtained evidence for an *extension effect*, according to which features influence similarity judgments more when they vary within an entire set of stimuli. Items presented within a particular trial also influence similarity judgments. Perhaps the most famous example of this is Tversky's (1977) *diagnosticity effect*, according to which features that are diagnostic for relevant classifications will have disproportionate influence on similarity judgments. More recently, Medin, Goldstone, & Gentner (1993) have argued that different comparison standards are created depending on the items that are present on a particular trial. Other research has documented intransitivities in similarity judgments, situations where A is judged to be more similar to T than is B, B is more similar to T than is C, and C is more similar to T than is A (Goldstone, Medin, & Halberstadt, 1997). This kind of result also suggests that the properties used to assess the similarity of objects are determined, in part, by the compared objects themselves.

Similarity judgments not only depend on the context established by recently exposed items, simultaneously presented items, and inferred contrast sets, but also on the observer. Suzuki, Ohnishi, and Shigemasa (1992) have shown that similarity judgments depend on level of expertise and goals. Expert and novice subjects were asked to solve the Tower of Hanoi puzzle, and judge the similarity between the goal and various states. Experts' similarity ratings were based on the number of moves required to transform one position to the other. Less expert subjects tended to base their judgments on the number of shared superficial features. Similarly, Hardiman, Dufresne, and Mestre (1989) found that expert and novice physicists evaluate the similarity of physics problems differently, with experts basing similarity judgments more on general principles of physics than on superficial features (see Sjöberg, 1972 for other expert/novice differences in similarity ratings). The dependency of

similarity on observer-, task- and stimulus-defined contexts offers the promise that it is indeed flexible enough to subserve cognition.

Is similarity too flexible to provide useful explanations of cognition?

As a response to the skeptic of similarity's usefulness, the preceding two paragraphs could have the exact opposite of their intended effect. The skeptic might now feel that similarity is much too flexible to be a stable ground for cognition. In fact, Nelson Goodman (1972) has put forth exactly this claim, maintaining that the notion of similarity is either vague or unnecessary. He argued that "when to the statement that two things are similar we add a specification of the property that they have in common ... we render it [the similarity statement] superfluous" (p. 445). That is, all of the potential explanatory work is done by the "with respect to property Z" clause and not by the similarity statement. Instead of saying "This object belongs to Category A because it is similar to A items with respect to the property 'red'," we can simplify matters by removing any notion of similarity with "This object belongs to Category A because it is red."

There are reasons to resist Goodman's conclusion that "Similarity tends under analysis either to vanish entirely or to require for its explanation just what it purports to explain" (p. 446). In most cases, similarity is useful precisely because we cannot flesh out the "respect to property Z" clause with just a single property. Evidence suggests that assessments of overall similarity are natural and perhaps even "primitive." Evidence from children's perception of similarity suggests that children are particularly likely to judge similarity on the basis of many integrated properties rather than analysis into dimensions. Even dimensions that are perceptually separable are treated as fused in similarity judgments (Smith & Kehler, 1978). Children under five years of age tend to classify on the basis of overall similarity and not on the basis of a single criterial attribute (Keil, 1989; Smith, 1989). Children often have great difficulty identifying the dimension along which two objects vary, even though they can easily identify that the objects are different in some way (Kehler, 1983). Smith (1989) argued that it is relatively difficult for young children to say whether two objects are identical on a particular property, but relatively easy for them to say whether they are similar across many dimensions.

There is also evidence that adults often have an overall impression of similarity without analysis into specific properties. Ward (1983) found that adult subject who tended to group objects quickly also tended to group objects like children, by considering overall

similarity across all dimensions instead of maximal similarity on one dimension. Likewise, Smith and Kemler (1984) found that adults who were given a distracting task produced more judgments by overall similarity than subjects who were not. To the extent that similarity is determined by many properties, it is less subject to drastic context-driven changes. Furthermore, integrating multiple sources of information into a single assessment of similarity becomes particularly important. The four approaches to similarity described in the previous section all provide methods for integrating multiple properties into a single similarity judgment, and as such, go significantly beyond simply determining a single “property Z” to attend.

A final point to make about the potential over-flexibility of similarity is that although impressions of similarity can change with context and experience, automatic and “generic” assessments of similarity typically change slowly and with considerable inertia. Similarities that were once effortful and strategic become second nature to the organism. Roughly speaking, this is the process of perceiving what was once a conceptual similarity. At first, the novice mycologist explicitly uses rules for perceiving the dissimilarity between the pleasing *Agaricus Bisporus* mushroom and the deadly *Amanita Phalloides*. With time, this dissimilarity ceases to be effortful and rule-based, and becomes perceptual and phenomenologically direct. When this occurs, the similarity becomes generic and default, and can be used as the ground for new strategic similarities. In this way, our cognitive abilities gradually attain sophistication, by treating territory as level ground that once made for difficult mental climbing. A corollary of this contention is that our default impression similarity does not typically mislead us; it is explicitly designed to lead us to see relations between things that often function similarly in our world. People, with good reason, expect their default similarity assessments to provide good clues about where to uncover directed, nonapparent similarities (Medin & Ortony, 1989).

Should “similarity” even be a field of study within cognitive science?

This survey has proceeded under the convenient fiction that it is possible to tell a general story for how people compare things. One reason to doubt this is that the methods used for assessing similarity have large effects on the resulting similarity viewed. Similarity as measured by ratings is not equivalent to similarity as measured by perceptual discriminability. Although these measures correlate highly, systematic differences are found (Podgorny & Garner, 1979; Sergent & Takane, 1987). For example, Beck (1966) finds that an

upright T is rated as more similar to a tilted T than an upright L, but that it is also more likely to be perceptually grouped with the upright Ls. Previously reviewed experiments indicate the non-equivalence of assessments that use similarity versus dissimilarity ratings, categorization versus forced-choice similarity judgments, or speeded versus leisurely judgments. In everyday discourse we talk about the similarity of two things, forgetting that this assessment depends upon a particular task and circumstance.

Furthermore, it may turn out that the calculation of similarity is fundamentally different for different domains (see Medin, Lynch, & Solomon, 2000 for a thoughtful discussion of this issue). To know how to calculate the similarity of two faces, one would need to study faces specifically and the eventual account need not inform researchers interested in the similarity of words, works of music, or trees. A possible conclusion is that similarity is not a coherent notion at all. The term *similarity*, like the terms *bug* or *family values*, may not pick out a consolidated or principled set of things.

Although we sympathize with the impulse toward domain-specific accounts of similarity, we also believe in the value of studying general principles of comparison that potentially underlie many domains. Although we do not know whether general principles exist, one justification for pursuing them is the large pay-off that would result from discovering these principles if they do exist. A historically fruitful strategy, exemplified by Einstein's search for a law to unify gravitational and electromagnetic acceleration and Darwin's search for a unified law to understand the origins of humans and other animals, has been to understand differences as parametric variations within a single model. Finding differences across tasks does not necessarily point to the incoherency of similarity. An alternative perspective would use these task differences as an illuminating source of information in developing a unified account. The systematic nature of these task differences should stimulate accounts that include a formal description not only of stimulus components, but also of task components. Future success in understanding the task of comparison may depend on comparing tasks.

References

- Aha, D. W. (1992). Tolerating noisy, irrelevant and novel attributes in instance-based learning algorithms. International Journal of Man Machine Studies, 36, 267-287.
- Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. Journal of Experimental Psychology: General, 120, 3-19.
- Attneave, F. (1950). Dimensions of similarity. American Journal of Psychology, 63, 516-556.
- Bassok, M., & Medin, D. L. (1997). Birds of a feather flock together: Similarity judgments with semantically rich stimuli. Journal of Memory & Language, 36, 311-336.
- Beck, J. (1966). Effect of orientation and of shape similarity on perceptual grouping. Perception and Psychophysics, 1, 300-302.
- Bernstein, L. E., Demorest, M. E., & Eberhardt, S. P. (1994). A computational approach to analyzing sentential speech perception: Phoneme-to-phoneme stimulus/response alignment. Journal of the Acoustical Society of America, 95, 3617-3622.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. Psychological Review, 94, 115-147.
- Bush, R. R., & Mosteller, F. (1951). A model for stimulus generalization and discrimination. Psychological Review, 58, 413-423.
- Carey, S. (1985). Conceptual change in childhood. Cambridge, MA: Bradford Books.
- 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.
- Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8, 240-247
- Corter, J. E. (1987). Similarity, confusability, and the density hypothesis. Journal of Experimental Psychology: General, 116, 238-249.

- Corter, J. E. (1988). Testing the density hypothesis: Reply to Krumhansl. Journal of Experimental Psychology: General, 117, 105-106.
- Edelman, S. (1999). Representation and recognition in vision. Cambridge, MA: MIT Press.
- Eisler, H., & Ekman, G. (1959). A mechanism of subjective similarity. Acta Psychologica, 16, 1-10.
- Estes, W. K. (1994). Classification and cognition. New York: Oxford University Press
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. Artificial Intelligence, 41, 1-63.
- Fodor, J. A. (1983). The modularity of mind. Cambridge, MA: MIT Press/ Bradford Books.
- Frisch, S. A., Broe, M. B. & Pierrehumbert, J. B. (1995). The role of similarity in phonology: Explaining OCP-Place. In K. Elenius & P. Branderud (eds), Proceedings of the 13th International Conference of the Phonetic Sciences, 3, 544-547.
- Gardenfors, P. (2000). Conceptual spaces: The geometry of thought. Cambridge, MA: MIT Press.
- Garner, W. R. (1974). The processing of information and structure. New York: Wiley.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. Cognitive Psychology, 16, 341-370.
- Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. Cognition, 23, 183-209
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.
- Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes, (Eds.), Perspectives on language and thought interrelations in development. Cambridge, England: Cambridge University Press.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. Cognitive Science, 10(3), 277-300.
- Gilmore, G. C., Hersh, H., Caramazza, A., & Griffin, J. (1979). Multidimensional letter similarity derived from recognition errors. Perception & Psychophysics, 25, 425-431.
- Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. Psychological Science, 2, 50-55.

- Gluck, M.A., & Bower, G.H. (1990). Component and pattern information in adaptive networks. Journal of Experimental Psychology: General, 119, 105-109.
- Goldstone, R. L. (1994a). Similarity, interactive activation, and mapping. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 3-28.
- Goldstone, R. L. (1994b). The role of similarity in categorization: Providing a groundwork. Cognition, 52, 125-157.
- Goldstone, R. L. (1996). Alignment-based nonmonotonicities in similarity. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 988-1001.
- Goldstone, R. L., & Medin, D. L. (1994). The time course of comparison. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 29-50.
- Goldstone, R.L., Medin, D.L., & Gentner, D. (1991). Relations, Attributes, and the non-independence of features in similarity judgments. Cognitive Psychology. 222-264.
- Goldstone, R. L., Medin, D. L., & Halberstadt, J. (1997). Similarity in Context. Memory & Cognition, 25, 237-255.
- Goodman, N. (1972). Seven strictures on Similarity. In N. Goodman (Ed.), Problems and Projects. New York: The Bobbs-Merrill Co.
- Hahn, U. (2003). Similarity. In L. Nadel (Ed.) Encyclopedia of Cognitive Science. London: Macmillan.
- Hahn, U., & Chater, N. (1998). Understanding similarity: A joint project for psychology, case-based reasoning and law. Artificial Intelligence Review, 12, 393-427.
- Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. Cognition, 87, 1-32.
- Hardiman, P. T., Dufresne, R., & Mestre, J. P. (1989). The relation between problem categorization and problem solving among experts and novices. Memory & Cognition, 17, 627-638.
- Hayes-Roth, B., & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. Journal of Verbal Learning and Verbal Behavior, 16, 321-338
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 411-422.
- Hofstadter, D. (1997). Fluid concepts and creative analogies: computer models of the fundamental mechanisms of thought. New York: Basic Books.

- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). Induction: Processes of inference, learning, and discovery. Cambridge, MA: Bradford Books/MIT Press.
- Holyoak, K. J., & Gordon, P. C. (1983). Social reference points. Journal of Personality & Social Psychology, 44, 881-887.
- Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. Memory & Cognition, 15, 332-340.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, 13, 295-355.
- Holyoak, K.J., & Hummel, J.E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich and A. Markman (Eds.). Cognitive Dynamics: Conceptual Change in Humans and Machines. Hillsdale, NJ: Erlbaum.
- Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, 13, 295-355.
- Holyoak, K. J., & Thagard, P. (1995). Mental leaps: Analogy in creative thought. Cambridge, MA: MIT Press.
- Horgan, D. D., Millis, K., Neimeyer, R. A. (1989). Cognitive reorganization and the development of chess expertise. International Journal of Personal Construct Psychology, 2, 15-36.
- Hubel, D. H., & Wiesel (1968). Receptive fields and functional architecture of monkey striate cortex. Journal of Physiology, 195, 215-243.
- Hummel, J.E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich and A. Markman (Eds.). Cognitive Dynamics: Conceptual Change in Humans and Machines. Hillsdale, NJ: Erlbaum.
- Hummel, J.E. (2001). Complementary solutions to the binding problem in vision: Implications for shape perception and object recognition. Visual Cognition, 8, 489-517.
- Hummel, J.E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. Psychological Review, 99, 480-517.
- Hummel, J.E., & Holyoak, K.J. (1997). Distributed representations of structure: A theory of analogical access and mapping. Psychological Review, 104, 427-466.

- Hummel, J.E., & Holyoak, K.J. (2003). A symbolic-connectionist theory of relational inference and generalization. Psychological Review, 110, 220-263.
- Imai, S. (1977). Pattern similarity and cognitive transformations. Acta Psychologica, 41, 433-447
- James, W. (1890/1950). The principles of psychology. Dover: New York. (Original work published 1890)
- Jakobson, R. Fant, G., & Halle, M. (1963). Preliminaries to speech analysis : the distinctive features and their correlates. Cambridge, MA: MIT Press.
- Jones, S. S., & Smith, L. B. (1993). The place of perception in children's concepts. Cognitive Development, 8, 113-139.
- Katz, J. J., & Fodor, J. (1963). The structure of semantic theory. Language, 39, 170-210.
- Keil, F.C. (1989). Concepts, Kinds and Development. Cambridge, MA: Bradford Books/MIT Press.
- Kemler, D. G. (1983). Holistic and analytic modes in perceptual and cognitive development. In T. J. Tighe & B. E. Shepp (Eds.), Perception, cognition, and development: Interactional analyses. (pp. 77-101). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kohonen, T. (1995). Self-organizing maps. Berlin: Springer-Verlag.
- Kolers, P. A., & Roediger, H. L. (1984). Procedures of Mind. Journal of Verbal Learning and Verbal Behavior, 23, 425-449.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. Child Development, 67, 2797-2822.
- Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. Psychological Review, 85, 450-463.
- Krumhansl, C.L. (1988). Testing the density hypothesis: Comment on Corter. Journal of Experimental Psychology: General, 117, 101-104.
- Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. Psychological Review, 99, 22-44.
- Lamberts, K. (2000). Information-accumulation theory of speeded categorization. Psychological Review, 107, 227-260.
- Lassaline, M. E. (1996). Structural alignment in induction and similarity. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 754-770.

- Lee, M.D. (2002a). A simple method for generating additive clustering models with limited complexity. Machine Learning, 49, 39-58.
- Lee, M.D. (2002b). Generating additive clustering models with limited stochastic complexity. Journal of Classification, 19, 69-85.
- Li, M., & Vitanyi, P. (1997). An introduction to Kolmogorov complexity and its applications, (2nd ed.). New York: Springer-Verlag.
- Love, B.C. (2000) A Computational Level Theory of Similarity. Proceeding of the Cognitive Science Society, 316-321.
- Markman, A.B., & Gentner, D. (1993a). Structural alignment during similarity comparisons. Cognitive Psychology, 25, 431-467
- Markman, A.B., & Gentner, D. (1993b). Splitting the differences: A structural alignment view of similarity. Journal of Memory & Language, 32, 517-535.
- Markman, A.B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. Memory & Cognition, 24, 235-249
- Markman, A.B., & Gentner, D. (1997). The effects of alignability on memory. Psychological Science, 8, 363-367.
- Markman, A.B., & Wisniewski, E. J. (1997). Similar and different: The differentiation of basic-level categories. Journal of Experimental Psychology: Learning, Memory, & Cognition, 23, 54-70.
- Marr, D. (1982). Vision. San Francisco: Freeman.
- Marr, D., & Nishihara, H.K. (1978). Representation and recognition of three dimensional shapes. Proceedings of the Royal Society of London, Series B, 200, 269-294.
- Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. Psychological Review, 100, 254-278.
- Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? Annual Review of Psychology, 51, 121-147.
- Medin, D.L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & Ortony (Eds.). Similarity and Analogical Reasoning. Cambridge, MA: Cambridge University Press.
- Medin, D. L., & Shaffer, M. M. (1978). A context theory of classification learning. Psychological Review, 85, 207-238.
- Mitchell, M. (1993). Analogy-making as perception: a computer model. Cambridge, MA: MIT Press.

- Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. Psychological Review, *92*, 289-316
- Murphy, G. L. (2002). The big book of concepts. Cambridge, MA: MIT press.
- Neisser, U. (1967). Cognitive Psychology. New York: Appleton-Century-Crofts
- Navarro, D.J., & Lee, M.D. (in press). Combining dimensions and features in similarity-based representations. *Neural Information Processing Systems*.
- Nickerson, R. S. (1972). Binary Classification reaction time: A review of some studies of human information-processing capabilities. Psychonomic Monograph Supplements, *4* (whole No. 6). 275-317.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, *10*, 104-114.
- Nosofsky, R. M (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, *115*, 39-57.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. Cognitive Psychology, *23*, 94-140.
- Osterholm, K., Woods, D. J., Le Unes, A. (1985). Multidimensional scaling of Rorschach inkblots: Relationships with structured self-report. Personality & Individual Differences, *6*, 77-82.
- Palmer, S. E. (1975). Visual perception and world knowledge. In D. A. Norman & D. E. Rumelhart (Eds.), Explorations in cognition. San Francisco: W. H. Freeman.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. Journal of Experimental Psychology: Learning, Memory, & Cognition, *23*, 324-354.
- Podgorny P., & Garner, W. R. (1979). Reaction time as a measure of inter- intraobject visual similarity: Letters of the alphabet. Perception & Psychophysics, *26*
- Polk, T. A., Behensky, C., Gonzalez, R., & Smith, E. E. (2002). Rating the similarity of simple perceptual stimuli: Asymmetries induced by manipulating exposure frequency, Cognition, *82*, B75-B88.
- Pylyshyn, Z.W. (1985). Computation and Cognition, Cambridge, Mass.: MIT press.
- Quine W. V. (1969) *Ontological Relativity and Other Essays*. New York, Columbia University Press. 165 p.
- Quine, W. V. (1977). Natural kinds. In S. P. Schwartz, ed., Naming, necessity, and natural kinds. Ithaca, NY: Cornell University Press.

- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93-134.
- Richardson, M. W. (1938). Multidimensional psychophysics. *Psychological Bulletin*, *35*, 659-660.
- Ripley B. D. (1996). Pattern recognition and neural networks. Cambridge: Cambridge University Press.
- Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadu & A. Ortony (Eds.), Similarity, analogy, and thought. (pp. 21-59). Cambridge: Cambridge University Press.
- Rips, L. J., & Collins, A. (1993). Categories and resemblance. *Journal of Experimental Psychology: General*, *122*, 468-486.
- Ritov, I., Gati, I., & Tversky, A. (1990). Differential weighting of common and distinctive components. *Journal of Experimental Psychology: General*, *119*, 30.
- Ross, B. H. (1987). This is like that: the use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 629-639.
- Ross, B. H., (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 456-468.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science*, *26*, 521-562.
- Schank, R. C. (1982). Dynamic memory: A theory of reminding and learning in computers and people. Cambridge: Cambridge University Press
- Schvaneveldt, R. (1985). Measuring the structure of expertise. *International Journal of Man-Machine Studies*, *23*, 699-728.
- Sergent, J., & Takane, Y. (1987). Structures in two-choice reaction-time data. *Journal of Experimental Psychology: Human Perception and Performance*, *13*, 300-315
- Shepard, R. N. (1962a) The analysis of proximities: Multidimensional scaling with an unknown distance function. Part I. *Psychometrika*, *27*, 125-140.
- Shepard, R. N. (1962b) The analysis of proximities: Multidimensional scaling with an unknown distance function. Part II. *Psychometrika*, *27*, 219-246.

- Shepard, R. N. (1972). Psychological representation of speech sounds. In E. E. David, Jr., & P. B. Denes (Eds.), Human communication: A unified view. New York: McGraw-Hill.
- Shepard, R. N. (1982). Geometrical approximations to the structure of **musical** pitch. Psychological Review, *89*, 305-333.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. Science, *237*, 1317-1323.
- Shepard, R. N., & Arabie, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. Psychological Review, *86*, 87-123.
- Simon, D., & Holyoak, K. J. (2002). Structural dynamics of cognition: From consistency theories to constraint satisfaction. Personality & Social Psychology Review, *6*, 283-294.
- Sjoberg, L. (1972). A cognitive theory of similarity. Goteborg Psychological Reports, *2*(10).
- Sloman, S. A. (1993). Feature-based induction. Cognitive Psychology, *25*, 231-280.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. Psychological Bulletin, *119*, 3-22.
- 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.
- Smith, E. E., & Sloman, S. A. (1994). Similarity-versus rule-based categorization. Memory & Cognition, *22*, 377-386.
- Smith, L. B. (1989). From global similarity to kinds of similarity: The construction of dimensions in development. In S. Vosniadou and A. Ortony (Eds.), Similarity and analogical reasoning (pp. 146 -178). Cambridge: Cambridge University Press.
- Smith, L.B., & Kemler, D.G. (1978). Levels of experienced dimensionality in children and adults. Cognitive Psychology, *10*, 502-532.
- Smith, J.D., & Kemler, D.G. (1984). Overall similarity in adults' classification: The child in all of us. Journal of Experimental Psychology: General, *113*, 137-159.
- Suzuki, H., Ohnishi, H., & Shigemasa, K. (1992). Goal-directed processes in similarity judgment. Proceedings of the Fourteenth Annual Conference of the Cognitive

- Science Society. (pp. 343-348). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Tarr, M.J., & Gauthier, I. (1998). Do viewpoint-dependent mechanisms generalize across members of a class? Cognition. Special Issue: Image-based object recognition in man, monkey, and machine, 67, 73-110.
- Tenenbaum, J.B. (1996). Learning the structure of similarity. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), Advances in Neural Information Processing Systems 8. Cambridge, MA: MIT Press, 4-9.
- Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. Advances in Neural Information Processing Systems 11, M. S. Kearns, S. A. Solla, & D. A. Cohn (eds.). Cambridge, MA: MIT Press.
- Tenenbaum, J. B., De Silva, V, & Lanford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. Science, 290, 22-23.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity and Bayesian inference. Behavioral & Brain Sciences, 24, 629-640.
- Torgerson, W. S. (1958). Theory and methods of scaling. New York: Wiley.
- Torgerson, W. S. (1965). Multidimensional scaling of similarity. Psychometrika, 30, 379-393.
- Treisman, A. M. (1986). Features and objects in visual processing. Scientific American, 255, 106-115.
- Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.
- Tversky, A., & Gati, I. (1982). Similarity, separability, and the triangle inequality. Psychological Review, 89, 123-154
- Tversky, A., & Hutchinson, J.W. (1986). Nearest neighbor analysis of psychological spaces. Psychological Review, 93, 3-22
- Ullman, S. (1996). High-level vision: object recognition and visual cognition. London: MIT Press.
- Ward, T.B. (1983). Response tempo and separable-integral responding: Evidence for an integral-to-separable processing sequence in visual perception. Journal of Experimental Psychology: Human Perception and Performance, 9, 103-112.
- Wedell, D. (1994). Context effects on similarity judgments of multidimensional stimuli: Inferring the structure of the emotion space. Journal of Experimental Social Psychology, 30, 1-38.

Wiener-Ehrlich, W.K., Bart, W.M., & Millward, R. (1980). An analysis of generative representation systems. Journal of Mathematical Psychology, 21(3), 219-246.

Zhang, S, & Markman, A.B. (1998). Overcoming the early entrant advantage: The role of alignable and nonalignable differences. Journal of Marketing Research, 35, 413-426.

Author Notes

This research was funded by NIH grant MH56871 and NSF grant 0125287. Correspondence concerning this chapter should be addressed to rgoldsto@indiana.edu or Robert Goldstone, Psychology Department, Indiana University, Bloomington, Indiana 47405. Further information about the laboratory can be found at <http://cognitrn.psych.indiana.edu>

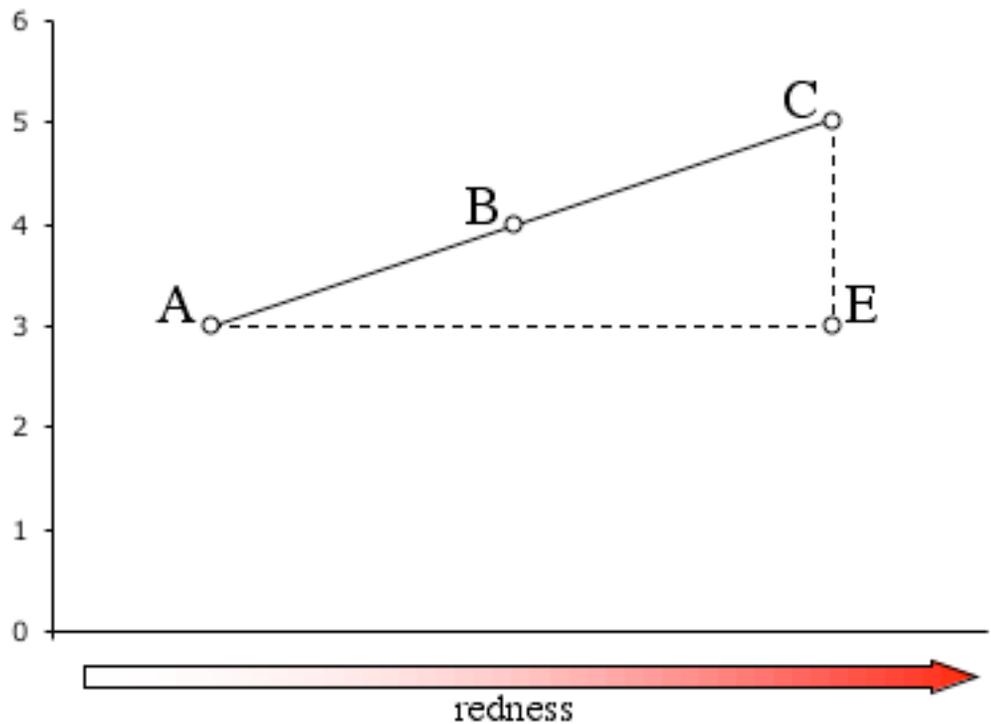


Figure 1. The triangle inequality assumption is requires the path from A to C going through B to be shorter than the path going through E.

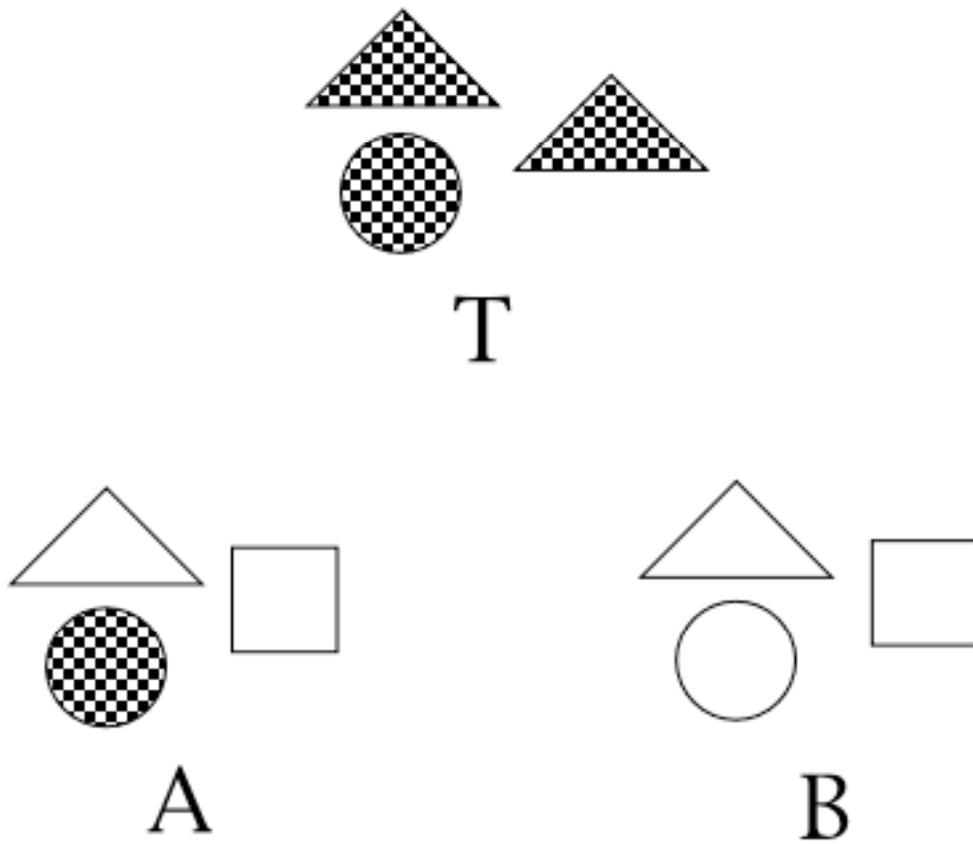


Figure 2. The set of objects in B is selected as both more similar to, and more different from, the set of objects in T.

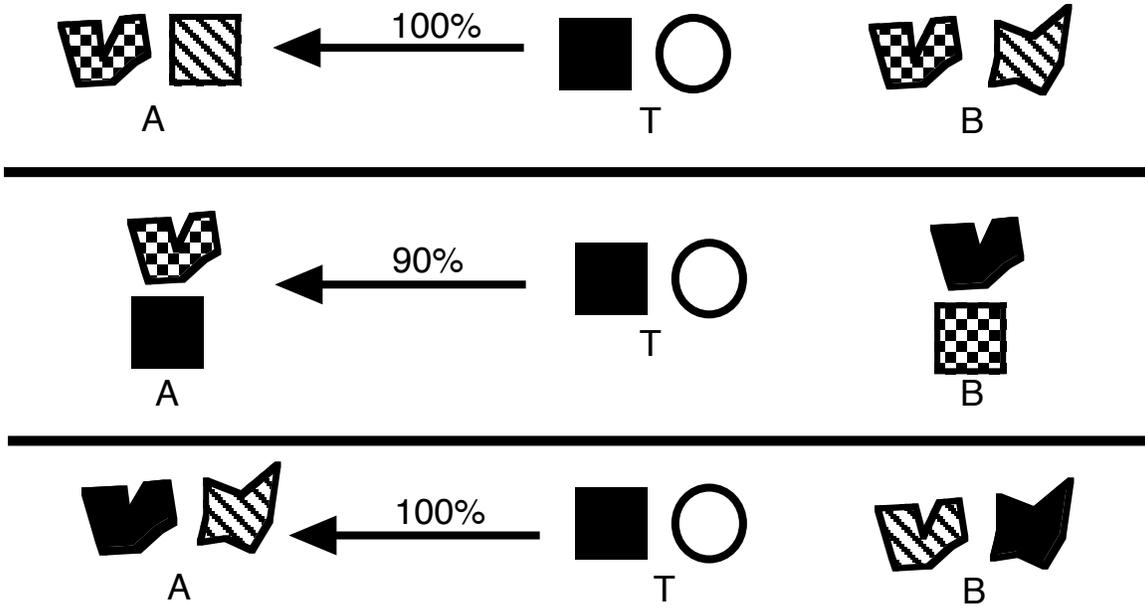


Figure 2. The sets of objects T are typically judged to be more similar to the objects in the A sets than the B sets. These judgments show that people pay attention to more than just simple properties like “black” or “square” when comparing scenes.

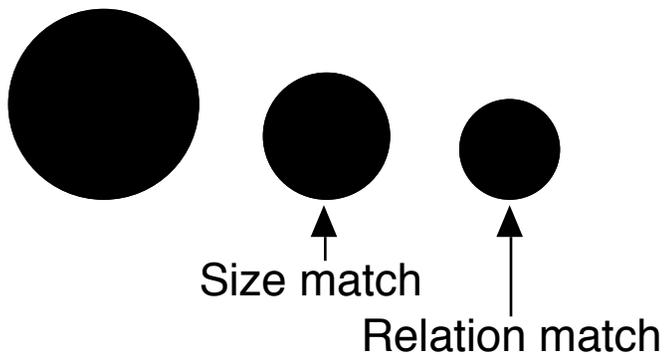
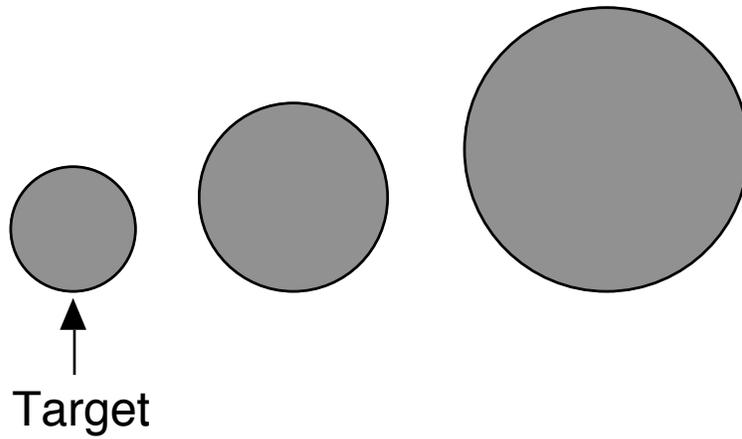


Figure 3. The Target from the gray circles could match either the middle black object because they are the same size, or the right-most object because both objects are smallest objects in their sets.