

Comparison
Robert L. Goldstone
Sam Day
Ji Y. Son
Indiana University

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

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

Introduction

It might not be immediately clear why the topic of comparison warrants a whole chapter in a book on human thinking. Of course, we are often required to make decisions that involve comparing two or more alternatives and assessing their relative value. Which car should I buy? Which job is more suited to my long-term goals? Would I rather have the soup or the salad? But in the grand scheme of human cognition, it might seem that such processes could be relegated to a subheading in a chapter on decision making.

In fact, comparison is one of the most integral components of human thought. Along with the related construct of *similarity*, comparison plays a crucial role in almost everything that we do. Furthermore, comparison itself is a powerful cognitive tool—in addition to its supporting role in other mental processes, research has demonstrated that the simple act of comparing two things can produce important changes in our knowledge.

One primary function of comparison is simply to assess the similarity of two things. To understand why this is such an important part of cognition, consider the variety of processes that are hypothesized to use similarity as an input. In models of memory, recognition and reminding have been argued to rely on the similarity between a stimulus and a long-term representation (Hintzman, 1986; Shiffrin & Steyvers, 1997). Models of categorization have proposed that new examples are classified based on their similarity to other category members (Medin & Schaffer, 1978; Nosofsky, 1984), or to a prototype of a category (Reed, 1972). When making inferences about unknown properties, people often appear to rely on their knowledge about other similar entities and situations to make reasonable predictions (Osherson, et al, 1990; Shepard, 1987), and people are very likely to look to similar situations from their past when understanding and solving new problems (Holyoak & Koh, 1987; Ross, 1989). Thus, it is a rare moment in our lives when comparison and similarity do *not* seem to play a role.

However, comparison does more than simply assess existing representations—it can also affect our understanding of the things that are being compared. For example, research in decision making has shown that people's judgments and preferences may vary significantly based on the particular comparisons that are made (Huber, Payne & Puto, 1982; Simonson, 1989). More direct evidence comes from Medin and colleagues (1993), who found that participants interpreted the features of an item differently when it had been compared to different alternatives. For example, in the top row of Figure 1, when the ambiguous object B

is compared to A, participants often write that a similarity between the pair is that both shapes have three prongs. However, when B is paired with C instead, participants often write that a similarity between the pair is that they both possess four prongs, and a difference is that one of B's prongs is warped or stunted. In other words, the comparison process seems to determine the content of our representations.

[Insert Figure 1 about here.]

Importantly, these representational changes often appear to be of a very beneficial kind: comparison can allow an individual to look past simple “surface” features, and to focus instead on potentially more meaningful structural commonalities and differences. For example, Gentner & Namy (1999; Namy & Gentner, 2002) found that comparing two objects allowed young children to overcome their strong bias for perceptual similarity, and to group objects instead on common taxonomic membership. Even more impressively, research has shown that a previous comparison can change the way that people interpret *new* situations. When people compare two cases that share the same underlying principle, they are far more likely to recognize new cases where that principle is applicable (e.g., Gick & Holyoak, 1983; Gentner, Loewenstein & Thompson, 2003). This improvement does not occur if the two cases are evaluated independently, without comparison (see Gentner's chapter on analogy in this book for a more detailed account of these kinds of effects). Even comparing situations that have slightly different underlying structures can be very beneficial, because it tends to highlight those structural differences (so-called “near miss” cases; Winston, 1975).

Comparison therefore provides an invaluable tool for learning, allowing people to see how two things are alike and different, and to see important features of each case that might otherwise have been overlooked. This helps to explain why educational assignments that ask a student to “compare and contrast” are such a powerful tool (i.e., Bransford & Schwartz, 1999), and makes it that much more puzzling that these types of assignments seem to have fallen out of favor in recent years.

Models of Similarity

Given the cognitive importance of comparison, it is understandable that there have been several attempts to formalize the comparison process. The formal treatments frequently center on the question of what makes things seem similar to people. One of the prominent goals of comparison is to determine how, and in what ways, two objects, scenes, or entities are similar to one another.

The formal treatments of similarity 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, 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).

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 obtained the results shown in Figure 2A (Figure 2B shows a second analysis involving animals more generally). The MDS algorithm produced this geometric representation by positioning the birds in a two-dimensional space such that birds that are rated as being highly similar are very close to each other in the space. 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. 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).

Insert Figure 2 about here.

Another 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.

Obtaining all pairwise similarity ratings among a large set of objects is, experimentally speaking, effortful. For N objects, N^2 ratings are required as input to a standard MDS algorithm. However, geometric models of similarity have received a recent boost from

automated techniques for analyzing large corpora of text. A computational approach to word meaning that has received considerable recent attention has been to base word meanings solely on the patterns of co-occurrence between a large number of words in an extremely large text corpus (Burgess & Lund, 2000; Griffiths, Steyvers, & Tenenbaum, 2007; Landauer & Dumais, 1997). Mathematical techniques are used to create vector encodings of words that efficiently capture their co-occurrences. If two words, such as “cocoon” and “butterfly” frequently co-occur in an encyclopedia or enter into similar patterns of co-occurrence with other words, then their vector representations will be highly similar. The meaning of a word, its vector in a high dimensional space, is completely based on the contextual similarity of words to other words. Within this high dimensional space, Landauer and Dumais (1997) conceive of similarity as the cosine of the angle between two words rather than their distance. With these new techniques, it is now possible to create geometric spaces with tens of thousands of words.

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:

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.

Violations of all three assumptions have been empirically obtained (Polk et al., 2002; Tversky, 1977; Tversky & Gati, 1982; Tversky & Hutchinson, 1986). 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:

$$S(A,B) = \theta f(A \cap B) - \alpha f(A - B) - \beta f(B - A)$$

The similarity of A to B is expressed as a linear combination of the measure of the common and distinctive features. The term $(A \cap B)$ 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. θ , α , and β 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 coherent as opposed to non-coherent 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 α is not constrained to equal β 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 α is greater than β . The Contrast Model can also account for non-mirroring between similarity and difference judgments. The common features term ($A \cap B$) 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.

A number of models are similar to the Contrast model in basing similarity on features and in using some combination of the ($A \cap B$), ($A - B$), and ($B - A$) components. Sjoberg (1972) proposes that similarity is defined as $f(A \cap B)/f(A \cup B)$. Eisler and Ekman (1959) claim that similarity is proportional to $f(A \cap B)/(f(A) + f(B))$. Bush and Mosteller (1951) defines similarity as $f(A \cap B)/f(A)$. These three models can all be considered specializations of the general equation $f(A \cap B)/[f(A \cap B) + \alpha f(A - B) + \beta f(B - A)]$. 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 $[(A-B) + (B-A)]/[f(A \cap B)]$.

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, 2004). 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 $(A \cap B) \supset (A \cap C)$ and $(B \cap C) \supset (A \cap C)$. For example, if orange lies between red and yellow on the hue dimension, then this can be featurally represented by orange sharing features with both red and yellow that features that red and yellow do not share between themselves.

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, 2002; Navarro & Griffiths, 2007; 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*.

Geometric and featural accounts of similarity fall short of a truly general capacity to handle structured inputs. Figure 3 shows an example of the need for structured representations is. Using these materials, 20 undergraduates were shown triads consisting of A, B, and T, and we asked them to decide 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.

[Insert Figure 3 about here.]

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 3, 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. More sophisticated objections have been raised about these approaches by John Hummel and colleagues (Doumas & Hummel, 2005; Hummel,

2000, 2001; Hummel & Biederman, 1992; Hummel & Holyoak, 1997, 2003; Holyoak & Hummel, 2000). 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, see Chapter XX in this volume; Holyoak, 2005; Holyoak & Thagard, 1995), 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. 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; Larkey & Love, 2003). Inspired by work in analogical reasoning (Gentner, 1983; 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. 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). The first effect is shown in Figure 4. Participants were asked to judge the similarity of scenes made up of two butterflies. The average similarity for the top panel comparison is greater than the middle panel comparison, because the weighting of feature match is affected by its alignment. In the top panel, the matching body pattern occurs between butterflies that are likely to be placed into alignment on the basis of their other feature matches. However, typically the unaligned feature matches (Matches Out of Place) still increase similarity somewhat, and hence the average similarity is higher for the middle than lowest panel comparisons.

Insert Figure 4 about here.

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. The role of structural alignment in influencing similarity of more natural stimuli, including words (Bernstein, Demorest, & Eberhardt, 1994; Frisch, Broe, & Pierrehumbert, 1995; Hahn & Bailey, 2005), sentences (Bassok & Medin, 1997), consumer products (Zhang & Markman, 1998), and legal reasoning (Hahn & Chater, 1998; Simon & Holyoak, 2002).

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.

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, empirically finding that as the number of transformations required to make two strings identical increased, so did the strings' dissimilarity.

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 of transformation, from one to 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.

Conclusions

The study of similarity and comparison 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 (Rips, 1989).

Despite the growing body of evidence that similarity comparisons 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 (Allen & Brooks, 1991; Palmeri, 1997; Smith & Sloman, 1994). People seem to have difficulties ignoring similarities between old and new patterns, even when they know a straightforward and perfectly accurate categorization rule. There appears to be a mandatory consideration of similarity in many categorization judgments (Goldstone, 1994b).

Similarity and comparison play powerful roles in cognition in situations where we do not know in advance exactly what properties of a situation are critical for its properties. We rely on comparison 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. Accordingly, comparison is an excellent general purpose cognitive strategy. For example, even if we do not know why sparrows have hollow bones, by comparing sparrows to warblers, we may be led to infer that if sparrows have hollow bones, then probably warblers do as well because of their similarity to sparrows. Similarities revealed through comparison thus play a crucial role in making predictions because, tautologically, similar things usually look and behave similarly. Furthermore, once sparrows and warblers are compared, we may not only come to realize that they share the property of hollow bones, but we may even generate an explanation for this trait involving weight, energy requirements to lift a mass, and the importance of flight for the ecological niche of birds. This explanation can cause us to look at birds in a new way. In this way, comparison not only takes representations as inputs to establish similarities, but also uses similarity to establish new representations (Hofstadter, 1997; Medin et al., 1993; Mitchell, 1993). When we compare entities, our understanding of the entities changes, and this may turn out to be a far more important consequent of comparison than simply deriving an assessment of similarity.

References

- Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. Journal of Experimental Psychology: General, 120, 3-19.
- Anderson, J. R. (1982). Acquisition of a cognitive skill. Psychological Review, 89, 369-406.
- 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.
- 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.
- Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory. In E. Diettrich & A. B. Markman (Eds.) Cognitive dynamics: Conceptual change in humans and machines. (pp. 117-156). Mahwah, NJ: Lawrence Erlbaum Associates.
- Bush, R. R., & Mosteller, F. (1951). A model for stimulus generalization and discrimination. Psychological Review, 58, 413-423.
- 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.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8, 240-247
- Doumas, L. A. A., & Hummel, J. E. (2005). Approaches to modeling human mental representation: What works, what doesn't, and why. In K. J. Holyoak & R. G. Morrison (Eds.) The Cambridge Handbook of Thinking and Reasoning (pp. 73-91). Cambridge, England: Cambridge University Press.
- Eisler, H., & Ekman, G. (1959). A mechanism of subjective similarity. Acta Psychologica, 16, 1-10.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. Artificial Intelligence, 41, 1-63.
- 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.
- Garner, W. R. (1974). The processing of information and structure. New York: Wiley.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. Cognitive Psychology, 16, 341-370.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.
- Gentner, D., & 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.
- Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. Psychological Science, 2, 50-55.
- 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.
- Griffiths, T.L., Steyvers, M., & Tenenbaum, J.B.T. (2007). Topics in Semantic Representation. Psychological Review, 114(2), 211-244.
- Hahn, U. (2003). Similarity. In L. Nadel (Ed.) Encyclopedia of Cognitive Science. London: Macmillan.
- Hahn, U., & Bailey, R. M. (2005). What makes words sound similar? Cognition, 97, 227-267.
- 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.
- 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
- Hofstadter, D. (1997). Fluid concepts and creative analogies: computer models of the fundamental mechanisms of thought. New York: Basic Books.
- 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.
- 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.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. Child Development, 67, 2797-2822.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. Psychological Review, 104, 211-240.
- Larkey, L. B., & Love, B. C. (2003). CAB: Connectionist Analogy Builder. Cognitive Science, 27, 781-794.
- Lee, M.D. (2002). A simple method for generating additive clustering models with limited complexity. Machine Learning, 49, 39-58.
- Li, M., & Vitanyi, P. (1997). An introduction to Kolmogorov complexity and its applications, (2nd ed.). New York: Springer-Verlag.
- 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.
- Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. Psychological Review, 100, 254-278.
- 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
- Neisser, U. (1967). Cognitive Psychology. New York: Appleton-Century-Crofts
- Navarro, D. J., & Griffiths, T. L. (2007). A nonparametric Bayesian method for inferring features from similarity judgments. *Advances in Neural Information Processing Systems* *19*.
- Navarro, D.J., & Lee, M.D. (2004). Common and distinctive features in stimulus representation: A modified version of the contrast model. *Psychonomic Bulletin & Review*, *11*(6), 961–974.
- 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.
- 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.
- 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.
- Reed, S. K. (1972). Pattern recognition and categorization. Cognitive Psychology, 3, 382-407.
- 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.
- 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.
- 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., & Arable, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. Psychological Review, 86, 87-123.

- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM: Retrieving effectively from memory. *Psychonomic Bulletin and Review*, 4(2), 145-166.
- 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. (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.
- 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., 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. (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.
- 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 Department of Education, Institute of Education Sciences grant R305H050116, 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://cognitn.psych.indiana.edu>

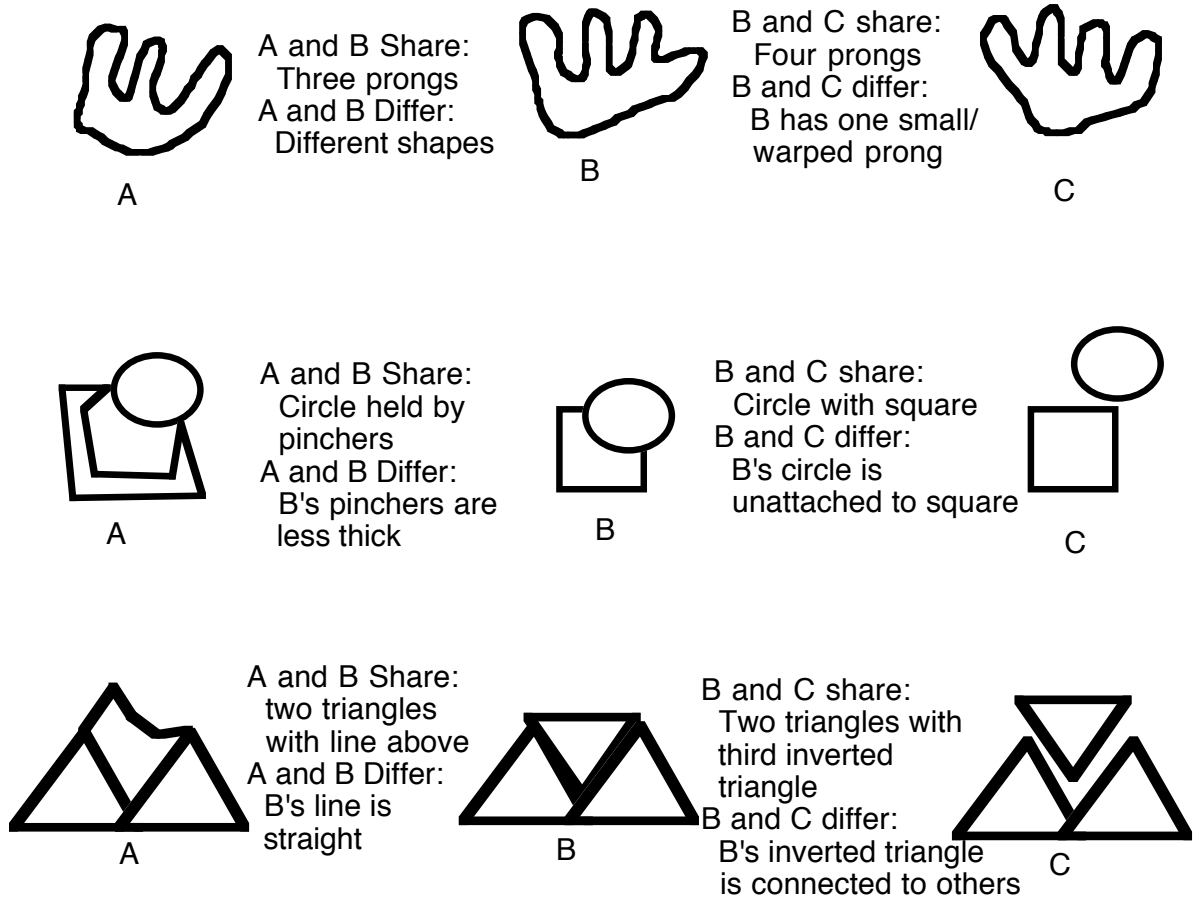


Figure 1. Examples of stimuli from Medin, Goldstone, and Gentner (1993). Subjects were asked to describe features that were shared and difference between pairs of objects. The middle objects labeled B are ambiguous, and tend to be interpreted in a manner that is consistent with the objects (A or C) with which they are paired. When determining both common and distinctive features, people apparently first interpret objects so as to make them more comparable.

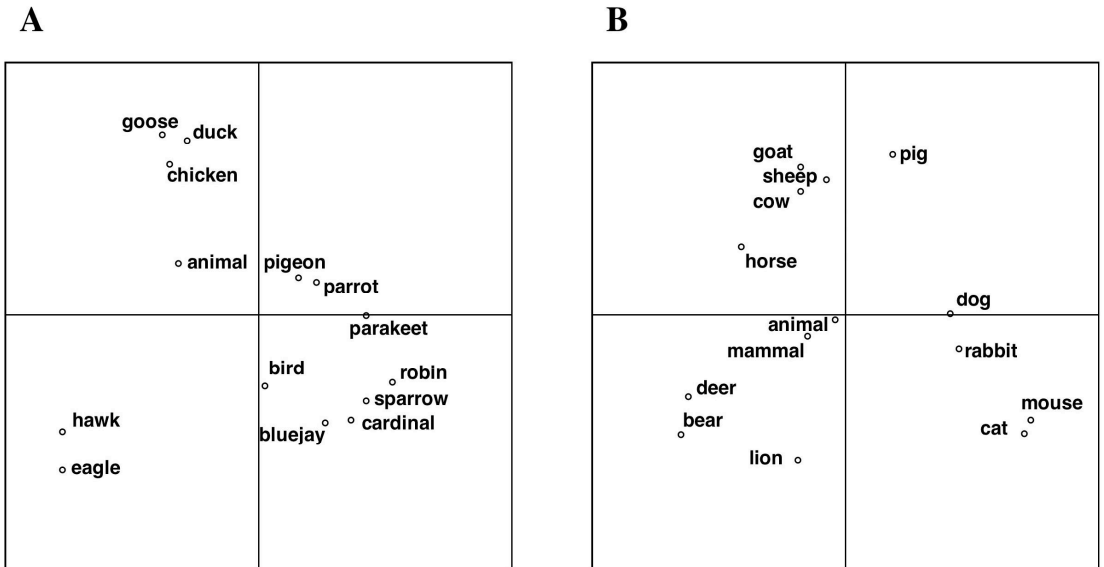


Figure 1. Two multidimensional scaling (MDS) solutions for sets of birds (A) and animals (B). The distances between the animals in the space reflect their psychological dissimilarity. Once an MDS solution has been made, psychological interpretations for the dimensions may be possible. In these solutions, the horizontal and vertical dimensions may represent size and domesticity, respectively (Reprinted from Rips, Shoben, & Smith, 1974, by permission).




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.

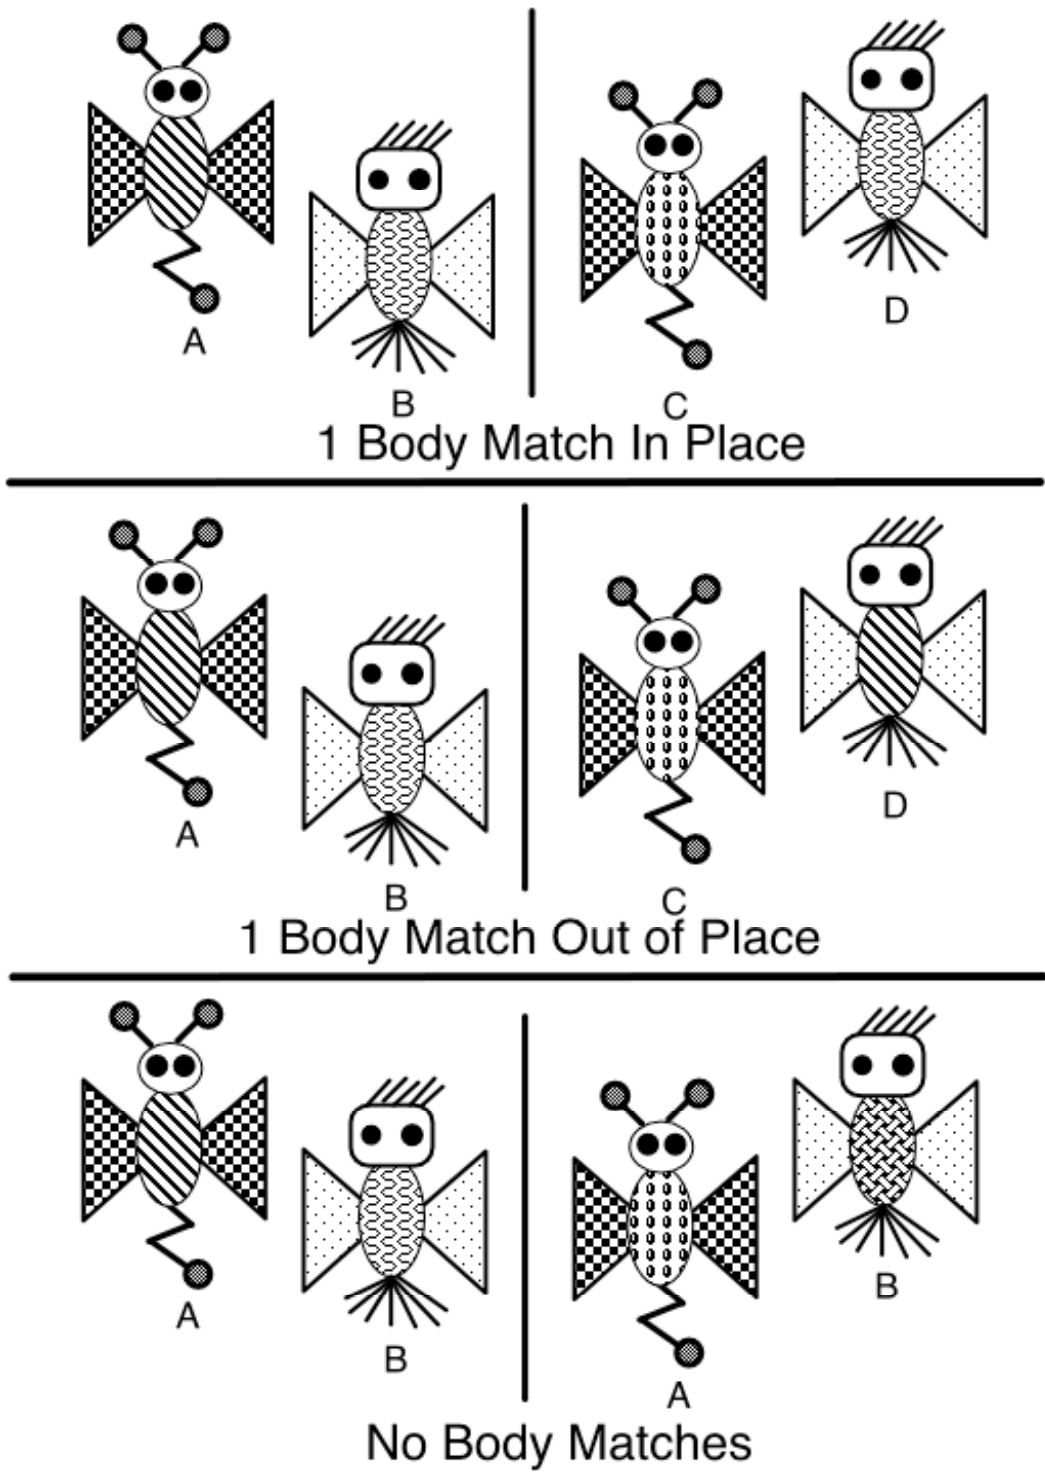


Figure 3. Sample scenes from Goldstone (1994a). In the top panel, the two butterflies that share a matching body pattern are aligned with each. In the middle panel, they are not aligned. In the lowest panel, there are no matching body patterns. Assessments of the similarity between the scenes decreases as we descend the panels.