Similarity

Robert L. Goldstone

Samuel B. Day

Indiana University

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

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

Introduction

An enormous amount of ink has been spilled in the psychology literature on the topic of similarity. There are two reasons that this seemingly intuitive and prosaic concept has been the subject of such intense scrutiny. First, there is virtually no area of cognitive processing in which similarity does not seem to play a role. William James observed that "This sense of Sameness is the very keel and backbone of our thinking" (James 1890/1950: 459). Ivan Pavlov first noted that dogs would generalize their learned salivation response to new sounds as a function of their similarity to the original tone, and this pattern of generalization appears to be ubiquitous across species and stimuli. People group things together based on their similarity, both during visual processing and categorization. Research suggests that memories are retrieved when they involve similar features or similar processing to a current situation. Problem solutions are likely to be retrieved from similar prior problems, inductive inference is largely based on the similarity between the known and unknown cases, and the list goes on and on.

Psychology clearly has an enormous stake in similarity. In fact, the depth of this stake makes it that much more unsettling that similarity can be such a slippery and temperamental construct. Perceived similarity can vary considerably with context. In fact, the act of comparison itself can change people's representations, leading them to construe things as more comparable by re-interpreting the given features or even creating new features. Similarity ratings are often asymmetric; for instance, people view North Korea to be more similar to China than the reverse. Likewise, "similarity" and "difference" are not always the inverse of one another. When given a choice between XX and OX, people choose XX as both more similar to, and more different from, OO. Similarity can change systematically with temporal distance and physical distance, and there is a growing body of evidence for consistent individual differences in which kinds of features drive a person's similarity judgments. Judgments of similarity can increase with simple association. For example, coffee is judged to be similar to cream because it is contextually associated with cream. In the absence of objective ways to measure psychological similarity, researchers are left to rely on participants' subjective judgments of their own processing, which are potentially unreliable, or on data from tasks in which similarity is proposed to play a role, leading to conclusions that are potentially circular. Thus, it is possible that similarity is both the most essential and the most problematic construct in cognitive science. In this chapter, we discuss some of the important findings from the extensive literature on similarity, and the ways in which psychologists have attempted to address a topic that is so critical to our understanding of the mind and yet so elusive.

Theories of Similarity

As one might expect given the importance of similarity for thinking, it is understandable that there have been several attempts to formalize the process of determining similarity. These formal accounts stipulate how similarity is to be empirically measured, and provide theoretic accounts of how similarity should be conceptualized. The resulting models have had a profound practical impact in knowledge representation, automatic pattern recognition by machines, search engines, data mining, and marketing (e.g. online stores that recommend new products to you based on your similarity to previous buyers). We survey four of the most prominent models of similarity:: geometric, feature-based, alignment-based, and transformational.

Geometric models and multidimensional scaling

Geometric models work under the premise that what it means for two things to be similar is for them to be close to one another in a psychological space. These approaches are exemplified by the statistical modeling technique of multidimensional scaling (MDS). MDS models represent similarity relations between entities in terms of a geometric model that consists of a set of points embedded in a space. The input to MDS routines may be similarity judgments, confusions between entities, patterns of co-occurrence in large samples of text, or any other measure of pairwise proximity. Most straightforwardly, participants may be asked to judge how similar every object in a set is to every other object. The output of an MDS routine is a geometric model of the set of objects, with each object 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:

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

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). Empirically, the Euclidean distance measure typically fits human data better when the stimuli being compared consist of perceptual dimensions that are psychologically fused together. For example, brightness is a subjective dimension related to the amount of luminance energy coming off of an object. Saturation is subjective dimension related to the amount of monochromatic light mixed into a color. Brightness and saturation are psychologically fused in the sense that it is difficult to pay attention to brightness differences between objects without also being influenced by saturation differences. For objects differing on saturation and brightness, their subjective similarity is best measured by a distance calculation that fully integrates saturation and brightness differences together, namely r=2. Conversely, if objects differ on brightness and size, then their similarity is best measured by computing their distance on brightness, and then adding this to their distance on size, namely r=1.

A classic example of MDS comes from Smith, Shoben, and Rips' study of animal concepts. They asked participants to provide similarity ratings on many pairs of birds or other animals. Submitting these pair-wise similarity ratings to MDS analysis, they obtained the results shown in Figure 2. 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.

Insert Figure 1 about here.

One practical limitation of MDS is that obtaining all pairwise similarity ratings among a large set of objects requires a substantial commitment of time and effort. If similarity ratings are used as the input to MDS, then standardly N^2 ratings are required for N objects. This number is halved if one assumes that the similarity of A to B is the same as the similarity of B to A. Even with this halving, the number of ratings still becomes prohibitively large as N increases. Fortunately, automated techniques for analyzing large corpora of text can provide the input to MDS models instead of relying on "manually" provided ratings. Using this method, Latent Semantic Analysis is a computational approach to word meaning that has received considerable recent attention. It bases word meanings solely on the patterns of co-occurrence between a large number of words in an extremely large text corpus such as an encyclopedia or thousands of email messages. It employs the mathematical analysis tool of Singular Value Decomposition (SVD) to create vector encodings of words that efficiently capture their co-occurrences. SVD is used to create encodings of words that represent each word by an ordered set of numbers, that is, a vector. The similarities between two words' vectors 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 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 the 1970s, it was observed that subjective assessments of similarity do 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. where D(A,B) is interpreted as the dissimilarity between items A and B. Minimality captures the simple ideas that no object should be more dissimilar to itself than it is to another object, and the dissimilarity of an object to itself should be 0. Symmetry captures the notion that the dissimilarity of Object A to B should be the same as the dissimilarity of Object B to A. Whatever Lady Gaga's similarity is to Madonna should be the same as Madonna's similarity to Lady Gaga. The notion behind the triangle inequality is that the length of the direct path from A to C should be no longer than the path from A to B plus the path from B to C. For similarities, this means that if a red square (A) is fairly distance to (dissimilar from) a blue circle (C), then the red square's distance to a red circle (B), and the red circle's distance to the blue circle cannot both be very short, otherwise the two-legged detour route from A to C going through B will be shorter than the direct route from A to C.

In fact, violations of all three assumptions have been empirically obtained. In response to these violations of the geometric model of similarity, some researchers have proposed fixes that allow, for example, the local density of objects in a region to warp the calculation of distance/dissimilarity. More radically, in 1977 Amos Tversky proposed to model similarity in terms of matching and mismatching features rather than distance on psychological dimensions. In his model, entities are represented as a collection of features and similarity is computed by:

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

Where S(A,B) is the similarity of Object A to Object B, and is expressed as a linear combination of the measure of the common and distinctive features. The term $(A \cap B)$ represents the features that Objects 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, reflecting how important each component is for determining similarity. For example, in Figure 2 we imagine comparing robots (A) to zombies (B). This would be accomplished by first determining all of the features of each of these two objects. Then, their similarity is calculated to be a positive function of their shared features, and negative functions of the features possessed by robots but not zombies (A – B) and features possessed by zombies but not robots (B-A).

Alignment-based Models

MDS and featural models make different assumptions about similarity, but they also share a number of similarities. An important commonality between geometric and featural representations is that both use relatively unstructured representations. Entities are structured as sets of features or dimensions with no relations between these attributes. However, entities such as natural objects with parts, real-world scenes, words, sentences, stories, scientific theories, and faces are not simply a "grab bag" of attributes. The relations among the parts of entities matter. A dog biting a man is not the same thing as a man biting a dog, even though they both feature a dog, a man, and biting. How these elements are related matters. Partly in response to the difficulties that the geometric and featural 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 man wearing a black tie and a woman with black shoes both share the feature *black*, but this matching feature may not increase their similarity much because the man's tie does not correspond to the woman's shoes. Drawing inspiration from a structure-matching model of analogical reasoning by Dedre Gentner, in alignment-based models, matching features influence similarity more if they belong to parts that are placed in correspondence, and conversely, parts tend to be placed in correspondence if they have many features in common and are consistent with other emerging correspondences.

Transformational Models

A final approach to similarity maintains that the similarity of two objects is directly related to the number of transformations required to turn one object into the other. A critical step for these models is to specify what transformational operations are possible.

Researchers in artificial Intelligence have claimed that an object is recognized by being aligned with memorized pictorial descriptions. The unknown object will be placed in the object category that contains the candidate model with which it best aligns. The alignment operations rotate, scale, translate, and topographically warp object descriptions.

According to Hahn, Chater, and Richardson, the similarity between two entities is based on how complex is the sequence of transformations that changes one entity to the other. The simpler the transformation sequence, the more similar the entities are assumed to be. For example, the transformational complexity connecting 1 2 3 4 5 6 7 8 and 2 3 4 5 6 7 8 9 is small, because the simple instruction “add 1 to each digit” suffices to transform one into the other. Experiments demonstrate that once reasonable vocabularies of transformation are postulated, transformational complexity does indeed predict subjective similarity ratings. Furthermore, when a new transformation is learned that turns Object A into Object B, A thereby becomes increasingly similar to B.

Conclusions

It might be argued that all four of the above approaches err on the side of treating similarity as a unitary phenomenon. It could well turn out that the calculation of similarity is fundamentally different for different kinds of entities. Taken to an extreme, this notion raises the possibility that similarity is not a coherent notion at all. Like the terms *bug* or *family values*, *similarity* may not pick out a consolidated or principled set of things. Consistent with this possibility, we anticipate that much of the real theoretical work in the future will be achieved by determining what counts as the features and relations that underlie similarity assessments for different kinds of entities and in different situations. Nonetheless, one justification for pursuing general theories of similarity is that if they do exist, a large payoff results. Even if similarity is not a monolithic *entity*, there probably will be common cognitive processes involved in different kinds of comparisons.

Some philosophers have attacked the very notion of similarity as being empty or circular. They have pointed out that the claim that Entities A and B are similar is vague and ill-defined unless one specifies the aspects under consideration when making the claim. However, part of the power of the notion of similarity is that it integrates over many aspects of entities. All four of the models of similarity can be interpreted as proposals for how the process of integrating across aspects proceeds. By integrating over many aspects, similarity is a powerful tool for cognition because it does not require the cognizer to fully understand exactly what makes entities behave as they do. It only requires that the world is a sufficiently orderly 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 tend to behave similarly in most respects.

Signed: Robert L. Goldstone and Samuel B. Day

See Also: Categorization: Psychological perspectives, Categorization: Computational perspectives, Concepts: Philosophical Issues, Concepts: Development of, Analogical mapping and reasoning

Further Readings

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

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 3-28.

Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. Cognition, 87, 1-32.

Holyoak, K. J., &Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.

James, W. (1890/1950). *The Principles of Psychology*. Dover: NewYork. (Original work published 1890.)

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.

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

Shepard, R. N. (1962a) The analysis of proximities: Multidimensional scaling with an unknown distance function. Part I. Psychometrika, 27, 125-140.

Smith, E. E., Shoben, E. J., and Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81* 214–241

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity and Bayesian inference. Behavioral & Brain Sciences, 24, 629-640.

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

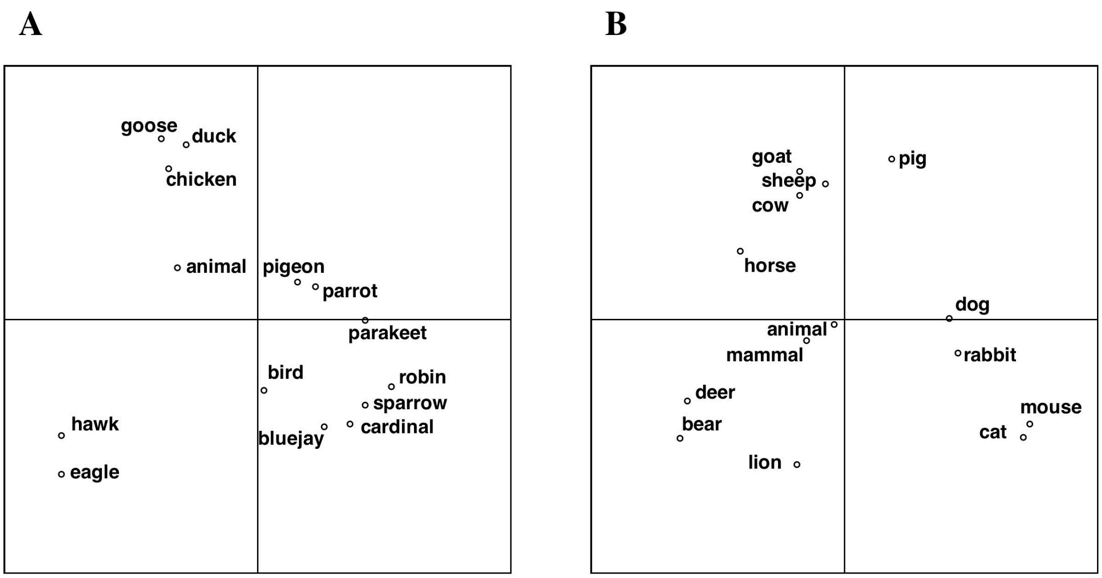

Figure 1. Two multidimensional scaling (MDS) solutions for sets of birds (A) and animals (B). The distances between the animals in the space reflect their psychological dissimilarity. Once an MDS solution has been made, psychological interpretations for the dimensions may be possible. In these solutions,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

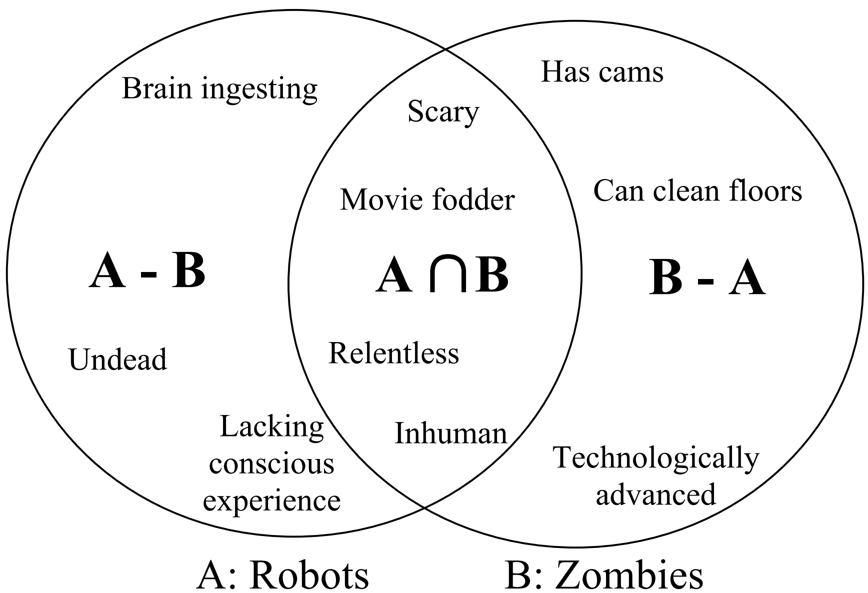

Figure 2. A Venn diagram showing some of the features of robots and zombies. Featural models take the similarity of robots to zombies to be a positive function of the features shared by both, and negative functions of the features possessed by one but not the other.