Cognitive Science 29 (2005) 1061–1076

Processes of Similarity Judgment

Levi B. Larkey[a], Arthur B. Markman[b]

[a]*Computer and Computational Sciences Division, Los Alamos National Laboratory*
[b]*Department of Psychology, The University of Texas at Austin*

Received 9 July 2004; received in revised form 18 February 2005; accepted 11 April 2005

Abstract

Similarity underlies fundamental cognitive capabilities such as memory, categorization, decision making, problem solving, and reasoning. Although recent approaches to similarity appreciate the structure of mental representations, they differ in the processes posited to operate over these representations. We present an experiment that differentiates among extant structural accounts of similarity in their ability to account for patterns of similarity ratings. These data pose a challenge for transformation-based models and all but one mapping-based model, the Similarity as Interactive Activation and Mapping (SIAM) model of similarity.

Keywords: Cognitive modeling; Similarity; Knowledge representation; Psychology

1. Introduction

Similarity serves as factotum to cognition. Memory traces are activated according to their similarity to probes (Hintzman, 1986). Objects are categorized according to their similarity to category exemplars (Medin & Schaffer, 1978; Nosofsky, 1986, 1992) or category prototypes (Posner & Keele, 1968; Reed, 1972; Rosch & Mervis, 1975). Decisions may be based on the similarity of the situation that would result from a choice to an ideal situation (Medin, Goldstone, & Markman, 1995). Strategies used to solve previous problems are applied to new problems that are similar (Bassok, 1990; Kolodner, 1993; Novick, 1988, 1990). The strength of an inductive argument depends on the similarity of the target of the argument to the base of the argument (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990).

Two of the most influential approaches to similarity are spatial approaches and feature-set approaches. Spatial approaches, as exemplified by multidimensional scaling (Shepard, 1962a, 1962b), define similarity as inversely related to the distance between stimuli in a dimensionally organized metric space. Feature-set approaches, as exemplified by Tversky's

Requests for reprints should be sent to Levi B. Larkey, Los Alamos National Laboratory, MS B256, Los Alamos, NM 87545. E-mail: larkey@lanl.gov

(1977) Contrast Model, assume similarity increases as a function of the common features and decreases as a function of the distinctive features of compared items.

These approaches have been incorporated into many other cognitive theories (e.g., see Nosofsky, 1986, and Smith, 1990), although they do have limitations. In particular, multidimensional spaces and feature sets do not readily capture the structure of mental representations (Biederman, 1985; Fodor & Pylyshyn, 1988). For example, to represent a red square above a blue triangle, one must bind the features "red" and "square" to the top entity and the features "blue" and "triangle" to the bottom entity, as well as bind the top entity and the bottom entity to their respective roles in the "above" relation. Spatial and feature-set approaches lack general means for binding features to particular entities and entities to particular relations and, as a result, have difficulty distinguishing between, for example, a red square above a blue triangle and a red triangle above a blue square. Capturing such bindings requires structured representations (Hummel & Holyoak, 1997; Markman, 1999).

Although recent approaches address the role of structured representations in similarity comparisons, they differ in the processes posited to operate over these representations. In this article, we present an experiment that differentiates extant structural accounts of similarity.[1] The remainder of this article is organized as follows. We first describe models that have been applied to similarity comparisons involving structured representations. These models can be divided into two general approaches: the transformational approach, as instantiated by Representational Distortion (RD; Hahn, Chater, & Richardson, 2003), and the structure-mapping approach, as instantiated by the Structure-Mapping Engine (SME; Falkenhainer, Forbus, & Gentner, 1989), Similarity as Interactive Activation and Mapping (SIAM; Goldstone, 1994), and Connectionist Analogy Builder (CAB; Larkey & Love, 2003). We then describe the experiment and contrast the results with the predictions of each model.

2. Representational Distortion

According to RD, the similarity between a pair of items is inversely related to the number of basic transformations involved in distorting the representation of one item into the representation of the other item. For example, XXXOOXO is more similar to OXOOXXX than to OXOOXX because OXOOXXX involves a mirror transformation of XXXOOXO, but OXOOXX involves a mirror transformation of XXXOOXO plus a deletion of the rightmost X from OXOOXXX (Imai, 1977). Post hoc derivations of different sets of transformations for different participants have yielded high correlations between the number of transformations and similarity ratings (Wiener-Ehrlich, Bart, & Millward, 1980).

RD can be tested by examining how well the number of transformations predicts perceived similarity (Hahn et al., 2003); however, this approach requires a priori constraints on what counts as a unitary psychological transformation. These constraints remain elusive, but the central claim of RD that a given transformation has a fixed, nonpositive influence on similarity can be straightforwardly tested independent of assumptions about the set of basic transformations (Narvaez & Markman, 2004). This prediction is tested in the experiment.

3. Structure-mapping models

There is a growing body of evidence suggesting that similarity comparisons involve the same structure-mapping process involved in analogical reasoning (Gentner & Markman, 1997; Goldstone 1994; Goldstone & Medin, 1994; Markman & Gentner, 1993a, 1993b, 1996). This work extends Tversky's (1977) theory that the similarity of a pair of items is a function of the common and distinctive constituents of their mental representations by distinguishing between two types of commonalities and differences: those between constituents that correspond and those between constituents that do not correspond. In particular, matching or mismatching constituents have a greater influence on similarity if they correspond than if they do not correspond.

There are two types of differences between compared items (Markman & Gentner, 1993a). Alignable differences are differences between corresponding elements of compared items. For example, an alignable difference between a car and a motorcycle is the number of wheels they have. Nonalignable differences are differences between elements that do not correspond or differences where an element in one representation does not correspond to any element in the other representation. For example, a seat belt is a nonalignable difference between a car and a motorcycle because a motorcycle has no restraining device that corresponds to a car's seat belt. Alignable differences and nonalignable differences are psychologically distinct. Similar items tend to have more alignable differences than dissimilar items (Markman & Gentner, 1993a). Alignable differences are easier to list, serve as better memory probes, and of particular interest here, have a greater influence on similarity than nonalignable differences (Markman & Gentner, 1993a, 1996, 1997).

Likewise, there are two types of commonalities between compared items (Goldstone, 1994). A match in place (MIP) is a match between corresponding elements of compared items. A match out of place (MOP) is a match between elements that do not correspond. For example, when comparing a bird with a gray head and red wings to a bird with a gray head and a red tail, the colors of the birds' heads constitute a MIP because the heads correspond, whereas the red wings and the red tail are a MOP because the wings and tail do not correspond. MIPs have a greater influence on similarity than MOPs.

The psychological distinction between alignable differences, nonalignable differences, MIPs, and MOPs suggests that the process by which correspondences between compared items emerge is an essential component of similarity comparisons. Three prominent models of this mapping process in similarity comparisons that have been proposed are SME, SIAM, and CAB.

3.1. The Structure-Mapping Engine

Constraints governing correspondences between compared items have been studied extensively with respect to analogical reasoning. SME was designed as a simulation of Gentner's (1983) structure-mapping theory of analogy. In addition to analogy, SME can be run in literal similarity mode, which we describe here.[2]

SME takes as input two propositional representations composed of entities, predicates, and functions. Entities are objects and constants (e.g., car, John). Predicates can represent attributes or relations. An *attribute* is a unary predicate that describes some property of an entity—for example, red(x), square(x). A *relation* is a predicate with multiple arguments that can be entities or other predicates—for example, above(x, y), knows(x, y). Whereas predicates map into truth values, functions map one or more entities into another entity and can be used to represent dimensions—for example, height(x) = y, temperature(x) = y.

SME determines correspondences between the two representations given as input using a local-to-global alignment process. SME begins by finding all possible local correspondences between the two representations. Local correspondences are created between identical predicates and between the arguments of identical predicates. The process of using corresponding predicates to place their arguments into correspondence is applied recursively. At this stage, the mapping typically contains inconsistent local correspondences that place an element in one representation into correspondence with more than one element in the other representation. Next, SME produces one or a few globally consistent interpretations by coalescing combinations of consistent local correspondences. Finally, SME calculates scores reflecting the perceived similarity associated with each interpretation using a cascade-like algorithm that favors mappings preserving interconnected, higher-order relational structure. The interpretation with the highest score is selected as the preferred interpretation.

All interpretations generated by SME strictly impose the structural constraints of parallel connectivity and one-to-one mapping (Markman & Gentner, 2000). Parallel connectivity requires that the arguments of corresponding predicates themselves be placed into correspondence. For example, when comparing “John loves Lisa” and “Jane loves Tarzan,” the common “loves” relation supports mapping John to Jane and Lisa to Tarzan because John and Jane fill the “lover” roles and Lisa and Tarzan fill the “loved” roles in their respective “loves” relations. One-to-one mappings limit any element in one representation to corresponding to at most one element in the other representation. For example, John might be put into correspondence with Jane or with Tarzan, but not both Jane and Tarzan. Importantly, SME’s strict adherence to one-to-one mappings leads to the prediction that MOPs will not influence rated similarity.

3.2. Similarity as Interactive Activation and Mapping

Inspired by McClelland and Rumelhart’s (1981) interactive activation model of letter perception, SIAM is a localist connectionist model that determines similarity via a dynamic process of interactive activation among feature, object, and relational role correspondences. SIAM’s architecture consists of a network of nodes that represent all possible feature-to-feature, object-to-object, and role-to-role correspondences between compared stimuli.[3] The activation of a particular node indicates the strength of the correspondence it represents.

Excitatory and inhibitory connections between nodes in the network are set up according to the structure of the stimuli being compared. Objects are placed into correspondence according to correspondences between their features and roles. At the same time, features and roles are placed into correspondence according to correspondences between the objects they describe. There are also excitatory and inhibitory connections between feature-to-feature nodes, between object-to-object nodes, and between role-to-role nodes. These connections are inhibitory if the two nodes taken together place an element in one representation into correspondence with two elements in the other representation, and excitatory otherwise. In addition to being influenced by other nodes, feature-to-feature and role-to-role nodes are influenced by match values that represent perceptually determined similarities between their associated features or roles.

SIAM first creates correspondences between features and roles according to match values. Once features and roles begin to be placed into correspondence, activation spreads through the network, and SIAM begins to place objects into correspondence that are consistent with the feature and role correspondences. Once objects begin to be placed into correspondence, activation flows back to feature and role nodes that are consistent with the object correspondences.

SIAM defines similarity as a function of the match values for each feature-to-feature node weighted by the current activation of that node. Matching features increase similarity to a greater extent if their associated feature-to-feature node is highly activated than if it is not. Likewise, mismatching features decrease similarity to a greater extent if their associated feature-to-feature node is highly activated than if it is not. Thus, similarity is a function of common and distinctive features, where the importance of these features is determined by their degree of correspondence. Whereas SME's adherence to one-to-one mappings is strict, SIAM treats one-to-one mapping as a soft constraint (Holyoak & Thagard, 1989); one-to-one mapping is one of several parallel constraints guiding the mapping process. As a result, and because SIAM allows for different degrees of correspondence (unlike SME's all-or-nothing notion of correspondence), both MIPs and MOPs influence similarity, with MIPs having a stronger influence than MOPs.

3.3. Connectionist Analogy Builder

CAB is a performance model of comparison that makes time-course predictions as well as predictions concerning the role of working memory in the comparison process. CAB is a connectionist model that bears a resemblance to SIAM, but unlike SIAM, CAB is capable of mapping more complex relational structures involved in analogical reasoning.

CAB takes as input two directed graphs (see Fig. 1) that can be translated to and from frame systems (Minsky, 1981). Like SIAM, CAB determines similarity via a dynamic process of interactive activation among correspondences. CAB's architecture can be conceptualized as a network of nodes that represent possible correspondences between the two graphs given as input. The activation of a particular node indicates the strength of the correspondence it represents.

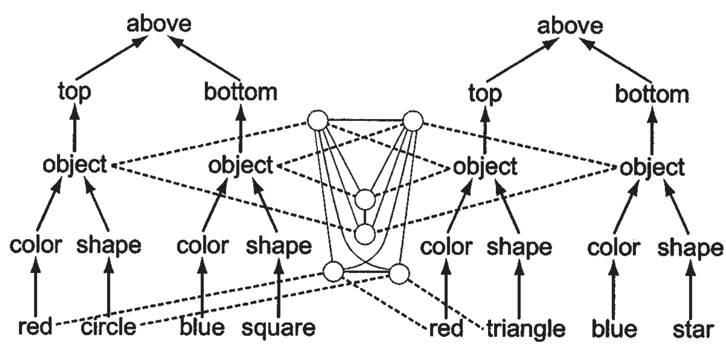

Fig 1. Comparison of a red circle above a blue square to a red triangle above a blue star. For clarity, only a few nodes are shown.

Early in processing, CAB is sensitive to any semantic commonalities between the two graphs. This semantic influence on mapping is captured by initially activating nodes representing correspondences between identical vertices. For example, both “object” vertices on the left of Fig. 1 would be tentatively placed into correspondence with both “object” vertices on the right. Whereas matching features and roles continually excite SIAM’s network, they only influence the initial activations of CAB’s network. From this starting point, activation spreads through the network, and CAB learns and unlearns correspondences.

Nodes are excited by other nodes that represent parallel correspondences. Two correspondences are parallel if the directed paths between their associated vertices are the same. For example, in Fig. 1 the correspondence between the “red” vertex on the left and the “red” vertex on the right is parallel to the correspondence between the “circle” vertex on the left and “triangle” vertex on the right because the edge directions of the path between “red” and “circle” on the left are the same as the edge directions of the path between “red” and “triangle” on the right. The strength of these excitatory connections decreases exponentially with the length of the path. This serves as a proxy for working-memory capacity because it governs how much information is considered simultaneously (see Hummel and Holyoak, 1997, for a similar view of working memory and its effects on mapping).

In addition to being excited by other nodes, nodes compete for excitation and are inhibited by other nodes. Excitation placing two vertices into correspondence is suppressed unless it is the maximal excitation received by any node associated with either of the vertices. This winner-takes-all rule imposes a strict one-to-one constraint on excitation. In addition, activations compete with each other to establish one-to-one mappings. Although not as strict as SME’s all-or-nothing notion of correspondence, CAB’s adherence to one-to-one mappings is stricter than SIAM’s.

Because the mapping process is dynamic and incremental, similarity can be calculated at any stage in processing using the current object-to-object correspondences to weight the influence of matching features. Matching features increase similarity to a greater extent if they describe objects placed into strong correspondence.

4. Experiment

The experiment was designed to answer several questions. First, does a given transformation have a fixed, nonpositive influence on similarity as predicted by RD? Second, are MOPs irrelevant to similarity as predicted by SME? Third, if MOPs do influence similarity, does SIAM or CAB capture the nature of this influence?

To test RD, SME, SIAM, and CAB, we asked participants to rate the similarity of two configurations, each consisting of a pair of objects. Objects varied in color and shape in the original study and color and texture in a replication. Objects within a pair were configured such that either one object was above the other object or one object was beside the other object.

To create one pair of objects, two different colors, two different shapes (or textures), and a spatial relation were randomly selected. To create the second pair of objects, the colors, shapes (or textures), and relation in the first configuration were independently transformed. The transformations were designed to test whether a given transformation has a fixed, nonpositive influence on similarity, and to systematically vary the number of MIPs and MOPs.

4.1. Method

4.1.1. Participants

A total of 116 people participated either for an $8 payment or to fulfill a course requirement at the University of Texas at Austin. Fifty-eight people participated in a study involving color and shape, and fifty-eight in a study involving color and texture.

4.1.2. Materials

Participants were shown 162 displays on a 43-cm wide × 34-cm high LCD screen. Each display contained two pairs of objects. In the original study, objects varied in color (green, red, blue, or yellow) and shape (triangle, circle, square, or star). In the replication, square objects varied in color (green, red, blue, or yellow) and texture (bark, carpet, pasta, or bubbles). Each individual object was approximately 3 cm wide × 3 cm high. Objects within each pair were configured such that either one object was above the other object or one object was beside the other object. The distance between objects within a pair was 1 cm. Pairs were randomly positioned on the screen subject to the constraint that the distance between the pairs was 11 cm.

4.1.3. Design

On each trial, one pair (the "base pair") was constructed by randomly selecting two different colors, two different shapes (or textures), and a spatial relation. The other pair (the "target pair") was constructed by selectively altering the features and relation of the base pair.

On half the trials, the relation between the objects was not altered; on the remaining half, the relation was different (e.g., if the objects in the base pair were one above the other, the objects in the target pair were one beside the other). When the relation was different, the top object in one pair was randomly selected as either the left or right object in the other pair.

Each feature dimension of the base pair was altered in one of nine ways, as illustrated in Fig. 2. A pair's respective values on a particular dimension can be abstractly represented by letters. For example, a pair with shapes represented by AB has shape A for one object and shape B for the other object. If AB denotes the base pair's respective values on a particular dimension, then the methods used to transform each dimension were AB → AB (no change), AB → BA (switch values), AB → AA (copy one value), AB → BB (copy the other value), AB → AC (replace one value), AB → CB (replace the other value), AB → CA (replace one value and switch values), AB → BC (replace the other value and switch values), AB → CD (replace both values). Each feature dimension and the relation were transformed independently, creating a total of 162 unique trials (9 methods of changing one dimension × 9 methods of changing the other dimension × 2 methods for changing the relation).

To control for the possibility that the absolute positions of objects or spatial relations between pairs might act as cues for mapping the objects, pairs were randomly positioned on the screen subject to the constraint that the distance between the pairs was always the same. Thus the only consistent spatial cue for mapping was the location of an object relative to the other object in its pair.

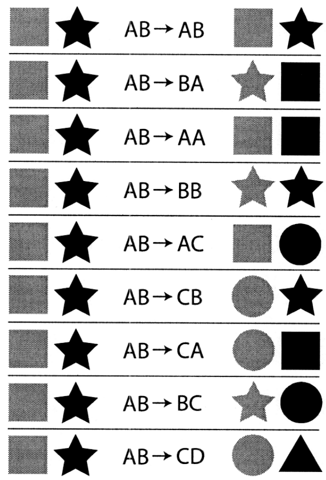

Fig. 2. Methods for altering feature dimensions. The target pairs (right column) were constructed by altering the shapes of the base pair (left column) according to each transformation (middle column).

4.1.4. Procedure

Each trial began with the simultaneous display of the base pair and the target pair. The participants' task was to rate the pairs' similarity on a scale from 1 (*low similarity*) to 6 (*high similarity*) that was displayed at the bottom of the screen. It was emphasized that participants should compare the pairs and rate the similarity between the pairs. After participants submitted a rating by clicking on the appropriate button, the screen was erased, and participants proceeded to the next trial. Each participant was presented with all 162 trials in randomized order.

4.1.5. Results

The primary data we used to test the models are the participants' similarity ratings for each method of changing one dimension when the method of changing the other dimension is AB $\rightarrow$ AB (i.e., the values are not changed) or AB $\rightarrow$ BA (i.e., the values are switched). Note that the methods AB $\rightarrow$ AA and AB $\rightarrow$ BB, AB $\rightarrow$ AC and AB $\rightarrow$ CB, and AB $\rightarrow$ CA and AB $\rightarrow$ BC are functionally equivalent when the method of changing the other dimension is symmetric (i.e., AB $\rightarrow$ AB or AB $\rightarrow$ BA). Thus ratings for these methods are collapsed and labeled by the former method (e.g., AB $\rightarrow$ AA denotes both AB $\rightarrow$ AA and AB $\rightarrow$ BB).

The data for the study involving color and shape are shown in Fig. 3. All pairwise differences between the methods of changing one dimension when the method of changing the other dimension is AB $\rightarrow$ AB are significant ($p < .001$, using Tukey honestly significant difference [HSD] test), except AB $\rightarrow$ AA and AB $\rightarrow$ AC, and AB $\rightarrow$ CA and AB $\rightarrow$ CD. All pairwise differences when the method of changing the other dimension is AB $\rightarrow$ BA are significant ($p < .001$, using Tukey HSD), except AB $\rightarrow$ AA and AB $\rightarrow$ CA, and AB $\rightarrow$ AC and AB $\rightarrow$ CD.

The data for the study involving color and texture are shown in Fig. 4. All pairwise differences between the methods of changing one dimension when the method of changing the other dimension is AB $\rightarrow$ AB are significant ($p < .001$, using Tukey HSD), except AB $\rightarrow$ BA, AB $\rightarrow$ AA, and AB $\rightarrow$ AC; and AB $\rightarrow$ CA and AB $\rightarrow$ CD. All pairwise differences when the method

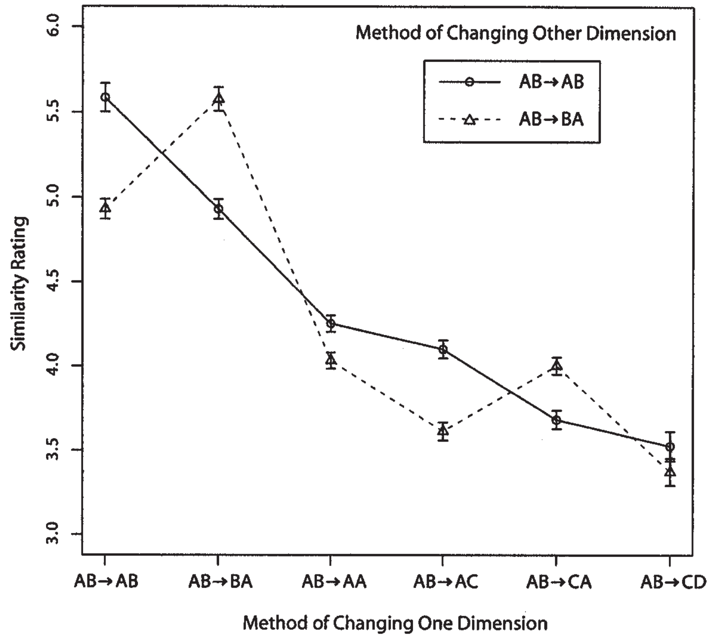

Fig. 3. Mean similarity ratings for each method of changing one dimension when the method of changing the other dimension is AB → AB or AB → BA. Feature dimensions are color and shape. Error bars denote standard errors.

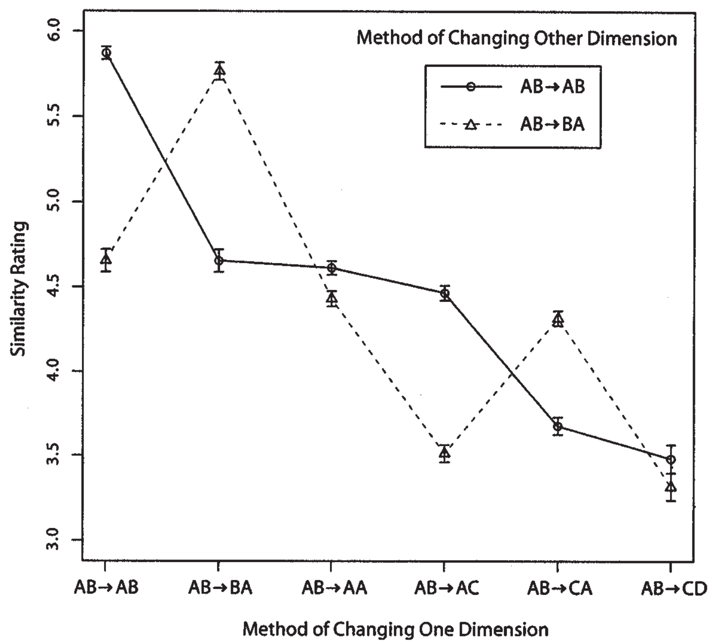

Fig. 4. Mean similarity ratings for each method of changing one dimension when the method of changing the other dimension is AB → AB or AB → BA. Feature dimensions are color and texture. Error bars denote standard errors.

of changing the other dimension is AB → BA are significant ($p < .001$, using Tukey HSD), except AB → AB and AB → AA, AB → AA and AB → CA, and AB → AC and AB → CD.

Ordinal relationships between ratings did not depend on the particular dimensions altered or whether the relation between the objects was altered. The latter is consistent with Davenport and Keane's (1999) finding that relations are downplayed if comparisons involving matching relations are randomly distributed among comparisons involving different relations. Two analyses were conducted with respect to individual differences. First, graphical examination of individual participants' responses revealed no discrete clusters of strategies. Second, for each trial type the distribution of responses had a single peak near the mean.

4.2. Comparison to model predictions

Figs. 5 and 6 show the ordinal relationships between methods of changing one dimension when the method of changing the other dimension is AB → AB or AB → BA. The first two columns show the ordinal relationships for the two studies. Horizontal lines separate methods that are significantly different. The next four columns show the ordinal predictions of the mod-

color & shape	color & texture	RD	SME	CAB	SIAM
AB→AB	AB→AB	AB→AB	AB→AB	AB→AB	AB→AB
AB→BA	AB→BA	AB→BA	AB→AA	AB→AA	AB→BA
AB→AA	AB→AA	AB→AA	AB→AC	AB→BA	AB→AA
AB→AC	AB→AC	AB→AC	AB→BA	AB→AC	AB→AC
AB→CA	AB→CA	AB→CA	AB→CA	AB→CA	AB→CA
AB→CD	AB→CD	AB→CD	AB→CD	AB→CD	AB→CD

Fig. 5. Ordinal relationships between methods of changing one dimension when the method of changing the other dimension is AB→AB.

color & shape	color & texture	RD	SME	CAB	SIAM
AB→BA	AB→BA	AB→AB	AB→BA	AB→BA	AB→BA
AB→AB	AB→AB	AB→BA	AB→AA	AB→AA	AB→AB
AB→AA	AB→AA	AB→AA	AB→CA	AB→AB	AB→AA
AB→CA	AB→CA	AB→AC	AB→AB	AB→CA	AB→CA
AB→AC	AB→AC	AB→CA	AB→AC	AB→AC	AB→AC
AB→CD	AB→CD	AB→CD	AB→CD	AB→CD	AB→CD

Fig. 6. Ordinal relationships between methods of changing one dimension when the method of changing the other dimension is AB→BA. In the second column, the placement of AB→AA on the line separating AB→AB and AB→CA denotes that AB→AA is not significantly different than AB→AB or AB→CA, but AB→AB is significantly different than AB→CA.

els. For each model, a similarity score was generated for each method, and the methods were ordered according to their scores. Horizontal lines separate methods that result in different predicted similarities. Methods are listed in descending order of similarity.

4.2.1. Transformational accounts

RD has difficulty accounting for the data. Post hoc weighting of transformations allows RD to fit either the pattern when the method of changing the other dimension is AB → AB, or the pattern when the method of changing the other dimension is AB → BA, but not both. No weighting of transformations can simultaneously fit both patterns of data. There are two reasons for this failure. First, a given transformation does not have a fixed influence on similarity in all contexts. For example, AB → CA decreases similarity more than AB → AC when the method of changing the other dimension is AB → AB, but AB → CA decreases similarity less than AB → AC when the method of changing the other dimension is AB → BA. Second, a given transformation can have a negative influence on similarity in some cases and a positive influence on similarity in others. For example, AB → BA decreases similarity when the method of changing the other dimension is AB → AB, but increases similarity when the method of changing the other dimension is AB → BA.

One way to mitigate RD's difficulties is to consider additional transformations that operate on whole objects. For example, if changing both dimensions via AB → BA is recast as changing the whole objects via AB → BA, and if this new transformation is given a lower weight than changing a single dimension via AB → BA, then RD is consistent with the data. However, RD provides no explanation for why changing a single dimension via AB → BA decreases similarity more than changing the whole objects via AB → BA. In contrast, this sort of prediction is the bread and butter of structure-mapping accounts. According to structure-mapping accounts, changing a single dimension via AB → BA decreases similarity more than changing the whole objects via AB → BA because the former results in structurally inconsistent feature correspondences, whereas the latter results in structurally consistent feature correspondences.

4.2.2. Structure-mapping approaches

Structure-mapping accounts differ from transformational accounts in two fundamental ways. First, structure-mapping accounts specify the process by which correspondences are determined, whereas transformational accounts do not. Second, structure-mapping accounts define similarity in terms of matching or mismatching representational constituents that correspond, whereas transformational accounts require additional processing to determine the transformations that distort one representation into the other. It might be possible to fit the data using a weighted transformational account that uses structure mapping to determine correspondences, but because a structure-mapping account can alone account for the data, the additional suppositions of transformational accounts are superfluous.

Structure-mapping accounts are alike in that correspondences between representational constituents modulate judgments of similarity, but they differ in how correspondences emerge from the comparison process. Unlike SIAM and CAB, SME ignores matching features that are inconsistent with the dominant mapping (i.e., MOPs). According to SME, similarity is a monotonically increasing function of the number of MIPs between compared items. Thus a critical test of SME is whether the number of MIPs predicts the ordinal relationship between methods of changing one dimension when the method of changing the other dimension is AB → AB. The ordinal relationship predicted by SME is shown in Fig. 5. SME fits the participants' similarity ratings with the exception of method AB → BA. For example, when comparing a red square above a blue circle to a red circle above a blue square, SME places the red square into correspondence with the red circle and the blue circle into correspondence with the blue square. According to SME, the two matching colors that correspond (MIPs) increase the similarity of the pairs, but the two matching shapes that do not correspond (MOPs) do not increase similarity. The discrepancy between SME's predictions and the data is due to participants' sensitivity to MOPs as well as MIPs.

Both SIAM and CAB are sensitive to MOPs as well as MIPs. Although MIPs and MOPs increase similarity, inconsistent MIPs and MOPs indirectly *decrease* similarity by competing for activation. For example, when comparing a red square above a blue circle to a red circle above a blue square, the matching shapes compete with the matching colors for activation because the shape correspondences are inconsistent with the color correspondences. The shapes vote for placing the red square into correspondence with the blue square and the blue circle into correspondence with the red circle, whereas the colors vote for placing the red square into correspondence with the red circle and the blue circle into correspondence with the blue square. These correspondences are inconsistent because taken together they constitute a many-to-one mapping between the pairs.

Competition between inconsistent MIPs and MOPs decreases similarity more in CAB than in SIAM. First, whereas both SIAM and CAB implement a soft one-to-one constraint on activation using inhibitory connections between nodes that taken together violate one-to-one correspondence, CAB in addition implements a strict one-to-one constraint on excitation using a winner-takes-all rule for incoming excitation. Because of this additional strict one-to-one constraint on excitation, competition between inconsistent MIPs and MOPs is stronger in CAB than in SIAM. Second, whereas semantic commonalities only influence the initial activations of CAB's network, SIAM continually incorporates the influence of match values. As a result, match values compensate for competition between inconsistent MIPs and MOPs by continually exciting MIPs and MOPs.

CAB fits the ordinal relationship between methods of changing one dimension when the method of changing the other dimension is AB → AB, with the exception of the relationship between AB → BA and AB → AA. CAB has difficulty with the relationship between AB → BA and AB → AA because it overestimates competition between the AB → AB method of changing one dimension and the AB → BA method of changing the other dimension.

SIAM predicts the data using its default parameters (Goldstone, 1994). SIAM's ability to predict the data arises because (unlike RD) it is sensitive to subjective correspondences between compared items, (unlike SME) it is sensitive to MOPs as well as MIPs, and (unlike CAB) it does not overestimate competition between inconsistent MIPs and MOPs. Impressively, SIAM predicts both patterns of data without post hoc parameter fitting.

5. Conclusions

These results are consistent with the hypothesis that correspondences between representational constituents modulate similarity judgments. The data are inconsistent with transformational accounts of similarity because a given transformation does not have a fixed, nonpositive influence on perceived similarity. Within structure-mapping accounts of similarity, only SIAM fully captures participants' similarity judgments. Inconsistent with SME's predictions, MOPs affect rated similarity. Inconsistent with CAB's predictions, competing MIPs and MOPs do not drastically decrease similarity.

SIAM has both conceptual strengths and weaknesses as a model of similarity. An important strength is its ability to model the dynamic time course of similarity comparisons (see Goldstone & Medin, 1994). Perhaps SIAM's greatest weakness is that, unlike SME and CAB, it does not process complex relational structures like those involved in analogical reasoning.

Another serious limitation of SIAM is that its representations are hand-coded. There have been several notable attempts to address the issue of representation construction. Copycat (Hofstadter, 1984; Mitchell, 1993) solves proportional letter-string analogies such as "abc is to abd as mrrjjj is to *blank*." Copycat's interpretations of analogies are built up from representational primitives, including letters, groups of letters, and relationships between letters. At the heart of Copycat is the idea that representation construction and mapping are inseparable (Hofstadter, 1995). However, Forbus, Gentner, Markman, and Ferguson (1998) argued that Copycat's restriction to a microdomain is a critical limitation:

> Copycat is unable to make correspondences between classes of statements that are not explicitly foreseen by its designers. Copycat cannot learn, because it cannot modify or extend these hand-coded representations that are essential to its operation. More fundamentally, it cannot capture what is perhaps the most important creative aspect of analogy: the ability to align and map systems of knowledge from different domains. (p. 245)

Several models create postulated representations from obtained similarity data (Navarro & Lee, 2003; Shepard, 1962a, 1962b, 1972; Shepard & Arabie, 1979; Tenenbaum, 1996). Lastly, Yan, Forbus, and Gentner (2003) proposed a theory of re-representation in analogical mapping that divides the problem into detecting opportunities for re-representation, generating re-representation suggestions based on libraries of general methods, and controlling the re-representation process.

Although SIAM has been successfully applied to many aspects of similarity judgments, the difficulties described earlier may preclude SIAM from serving as a similarity measure within models of memory, categorization, decision making, problem solving, and reasoning. To understand how these cognitive capabilities operate in complex environments involving relations between multiple entities, future research will need to address these difficulties. Perhaps some of SIAM's strengths can be incorporated into models that process more complex relational representations and construct their own representations.

Notes

1. Differentiating between structural accounts and spatial or feature-set accounts is beyond the scope of this article. We refer the reader to Markman and Gentner (1993a, 1993b, 1996), Gentner and Markman (1997), Goldstone (1994), Goldstone and Medin (1994), and Hahn et al. (2003).

2. Other models of analogical mapping like LISA (Hummel & Holyoak, 1997) and IAM (Keane & Brayshaw, 1988) do not address similarity comparisons, though it might be possible to add similarity modes to these models.
3. Goldstone (1998) describes a version of SIAM that does not initially create nodes for all possible correspondences, but rather creates nodes as needed during the mapping process.

Acknowledgments

This research was supported by an NDSEG Fellowship and a IC Postdoctoral Research Fellowship to L. B. Larkey and NIDA Grant NIH #1 R21 DA015211–01A1 to A. B. Markman.

We thank Rob Goldstone, Michael Lee, and Bob French for thoughtful and constructive comments on earlier versions of this work. We also thank Rob Goldstone for sending his source code for SIAM.

References

Bassok, M. (1990). Transfer of domain-specific problem-solving procedures. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 522–533.

Biederman, I. (1985). Human image understanding: Recent research and a theory. *Computer Vision, Graphics, and Image Processing, 32,* 29–73.

Davenport, J., & Keane, M. T. (1999). Similarity & structural alignment: You can have one without the other. In H. Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society* (pp. 132–137). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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

Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28,* 3–71.

Forbus, K. D., Gentner, D., Markman, A. B., & Ferguson, R. W. (1998). Analogy just looks like high-level perception: Why a domain-general approach to analogical mapping is right. *Journal of Experimental and Theoretical Artificial Intelligence, 10,* 231–257.

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

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52,* 45–56.

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

Goldstone, R. L. (1998). Hanging together: A connectionist model of similarity. In J. Grainger & A. M. Jacobs (Eds.), *Localist connectionist approaches to human cognition* (pp. 283–325). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Goldstone, R. L., & Medin, D. L. (1994). The time course of comparison. *Journal of Experimental Psychology, Learning, Memory, and Cognition, 20,* 29–50.

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

Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review, 93,* 411–428.

Hofstadter, D. R. (1984). The Copycat project: An experiment in nondeterminism and creative analogies (AI Memo 755). Cambridge, MA: MIT Artificial Intelligence Laboratory.

Hofstadter, D. R. (1995). *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.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427–466.

Imai, S. (1977). Pattern similarity and cognitive transformations. *Acta Psychologica, 41,* 433–447.

Keane, M. T., & Brayshaw, M. (1988). The incremental analogical machine: A computational model of analogy. In D. Sleeman (Ed.), *European Working Session on Third Learning* (pp. 53–62). London: Pitman.

Kolodner, J. L. (1993). *Case-based reasoning.* San Mateo, CA: Kaufmann.

Larkey, L. B., & Love, B. C. (2003). CAB: Connectionist analogy builder. *Cognitive Science, 27,* 781–794.

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Markman, A. B., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language, 32,* 517–535.

Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431–467.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition, 24,* 235–249.

Markman, A. B., & Gentner, D. (1997). The effects of alignability on memory. *Psychological Science, 8,* 363–367.

Markman, A. B., & Gentner, D. (2000). Structure mapping in the comparison process. *American Journal of Psychology, 113,* 501–538.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375–407.

Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processes and decision processes. *Psychonomic Bulletin and Review, 2,* 1–19.

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

Minsky, M. (1981). A framework for representing knowledge. In J. Haugland (Ed.), *Mind design* (pp. 95–128). Cambridge, MA: MIT Press.

Mitchell, M. (1993). *Analogy-making as perception: A computer model.* Cambridge, MA: MIT Press.

Narvaez, L., & Markman, A. B. (2005). *Studying similarity requires making representational assumptions.* Manuscript in preparation.

Navarro, D. J., & Lee, M. D. (2003). Combining dimensions and features in similarity-based representations. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems* (Vol. 15, pp. 67–74). Cambridge, MA: MIT Press.

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

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–394). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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

Novick, L. R. (1990). Representational transfer in problem solving. *Psychological Science, 1,* 128–132.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97,* 185–200.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract idea. *Journal of Experimental Psychology, 77,* 353–363.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3,* 382–407.

Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology, 7,* 573–605.

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 J. E. E. David & P. B. Denes (Eds.), *Human communication: A unified view.* New York: McGraw-Hill.

Shepard, R. N., & Arabie, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. *Psychological Review, 86,* 87–123.

Smith, E. E. (1990). Categorization. In D. N. Osherson & E. E. Smith (Eds.), *Thinking: An invitation to cognitive science* (pp. 33–54). Cambridge, MA: MIT Press.

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* (Vol. 8, pp. 4–9). Cambridge, MA: MIT Press.

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

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

Yan, J., Forbus, K., & Gentner, D. (2003). A theory of rerepresentation in analogical matching. In R. Alterman & D. Kirsch (Eds.), *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.