Markman, A. B., & Gentner, D. (2000). Structure mapping in the comparison process. American Journal of Psychology, 113(4), 501-538.

Structure mapping in the comparison process

ARTHUR B. MARKMAN
University of Texas at Austin

DEDRE GENTNER
Northwestern University

Carrying out similarity and analogy comparisons can be modeled as the alignment and mapping of structured representations. In this article we focus on three aspects of comparison that are central in structure-mapping theory. All three are controversial. First, comparison involves structured representations. Second, the comparison process is driven by a preference for connected relational structure. Third, the mapping between domains is rooted in semantic similarity between the relations that characterize the domains. For each of these points, we review supporting evidence and discuss some challenges raised by other researchers. We end with a discussion of the role of structure mapping in other cognitive processes.

Similarity is a central influence in human cognition. William James (1985) pointed out that all humans can extract common elements—including abstract relational elements—from comparisons, whereas other animals are unlikely to have this ability. More recently, research has searched for processes common across different kinds of comparisons, including analogy and metaphor as well as more prosaic similarity (Gentner, 1983, 1989; Gentner & Clement, 1988; Gentner & Markman, 1997; Markman & Gentner, 1993b; Medin, Goldstone, & Gentner, 1993).

In this article we take a step back from the individual theoretical and empirical statements and draw attention to three central conclusions about the nature of representation and processing that arise from research on comparison. First, cognitive representations are structured and the comparison process operates to align two structures. Second, the psychological sense of analogical relatedness depends on semantic commonalities between the relations in the two domains being compared. Third, the comparison process is driven by a search for correspondences that preserve connections between representational elements. The first and third points emphasize the importance of abstract structure. But matching abstract structure is not enough. As noted in

AMERICAN JOURNAL OF PSYCHOLOGY
Winter 2000, Vol. 113, No. 4, pp. 501–538

the second point, analogy also depends on matching conceptual relations. To examine these claims, we first describe the structure-mapping process.

Structure mapping in analogy and similarity

We introduce the comparison process with an example. Figure 1 presents a pair of scenes. At a high level of abstraction, these pictures are similar because both depict places where people live (indeed, places that the occupants may think of as small and cramped). More concretely, both contain beds, desks, and books. The pictures are by no means identical, however. One depicts a dorm room, the other a jail cell. There is a bunk bed in the jail cell and a single bed in the dorm. The books are on the bookshelf in the dorm room but on the desk in the jail cell. Furthermore, the dorm room has a mug of beer in it, but the jail cell does not. The jail cell has a toilet, but the dorm room does not.

People can compare scenes like this quickly, even if they have never seen the particular scenes before. Furthermore, the information being compared need not be presented explicitly in the stimulus to allow a comparison to happen. For example, given the statement

New York is like Chicago. (1)

people understand that this comparison involves similarities in the underlying concepts described by the words. For comparisons such as this, people are facile at generating both commonalities (i.e., "Both are large U.S. cities" and "Both contain skyscrapers") and differences (i.e., "New York is larger than Chicago" or "Chicago has a museum of science and industry, and New York does not"). Thus, people are good at quickly recognizing that there are similarities between a pair and at extracting the commonalities and differences of the pair.

How does comparison take place? In the comparison of the scenes in Figure 1 and the comparison in Sentence 1, there must be some representation of the elements being compared. As in other models of similarity (e.g., Tversky, 1977), the view we espouse assumes that the representations compared are generated by another process. Structure-mapping theory assumes that representations are structured, and thus they consist of entities (e.g., buildings, museums); attributes, which are representational elements that provide descriptive information (e.g., tall[building-1]); functions (e.g., color[building-1] = gray), which map onto values other than truth values and are used to represent psychological dimensions; and relations, which are representational elements that relate two or more entities, attributes, or other relations (e.g., taller[building-1, building-2]). Relations that take other relations as arguments are called higher-order relations (e.g., cause[knock over

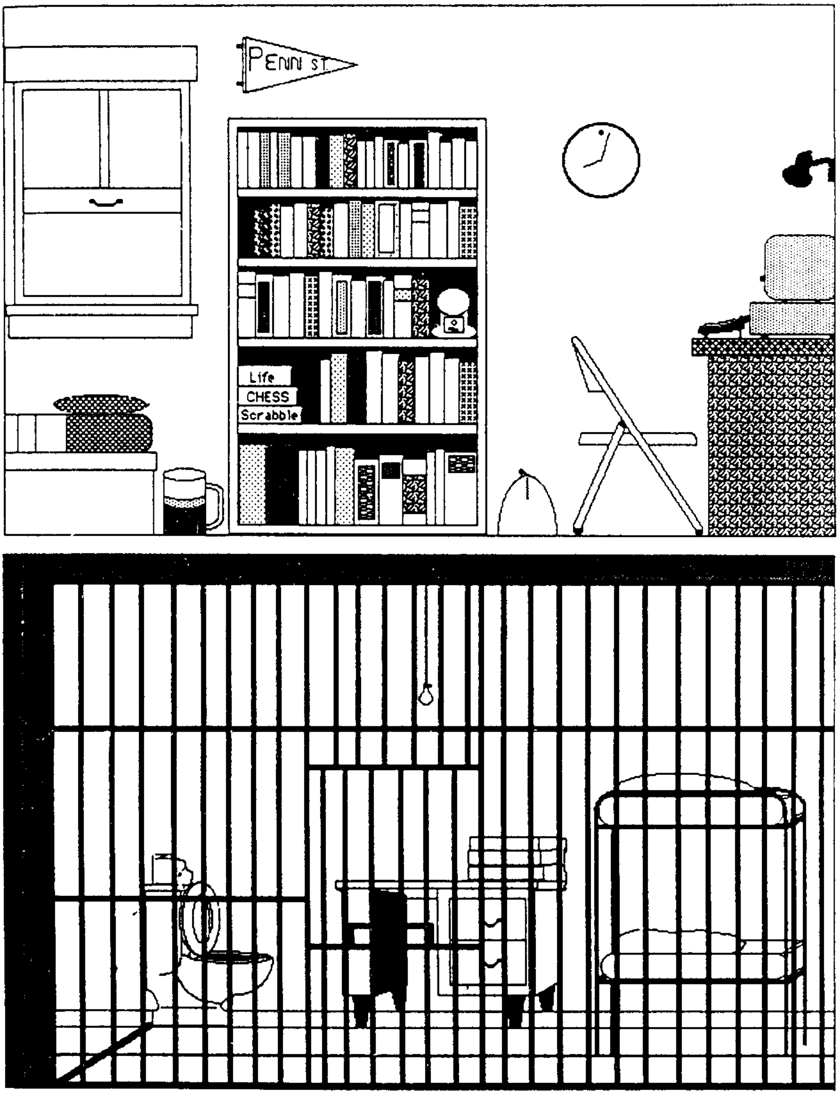

Figure 1. Pair of complex scenes that can be compared

(cow,lamp), destroy(fire,Chicago)]). Higher-order relations are particularly important in structure mapping because it is assumed that they encode important relations in a domain such as causal relationships and implications. The particular representation generated for a given situation is assumed to be the result of the person's construal of that situation. Different people might generate different representations of the same situation (as might the same person at different times).

We are now ready to set out eight signature phenomena that characterize analogical comparison. These benchmark phenomena are listed in Table 1, revised from an earlier list by Gentner and Markman (1995).[1]

The comparison of two domains involves an alignment of this representational structure. By convention, in a statement of the form "*X* is like *Y*," the *Y* domain—the given, or familiar domain—is called the base (or source) and the *X* domain—the new domain, to which the new information is to be applied—is called the target. There are three key constraints on the comparison process, given as the first three items in Table 1. The first is relational similarity. The basis of any similarity match is some kind of semantic similarity between the situations compared. In analogy, the similar elements must include some relational matches. Thus, the situations described by

The Celtics defeated the Lakers. (4)
and
Xerxes sacked Rome. (3)

can be seen as similar, even if the Celtics themselves have little in common with Xerxes.[2] In structure mapping, we take the stand that semantic

Table 1. Benchmark phenomena of analogy

Relational similarity	Analogies involve relational commonalities; object commonalities are optional.
Structural consistency	Analogical mapping involves one-to-one correspondence and parallel connectivity.
Systematicity	In interpreting analogy, connected systems of relations are preferred over sets of isolated relations.
Candidate inferences	Analogical inferences are generated via structural completion.
Alignable differences	Differences that are connected to the commonalities of a pair (and note unconnected differences) are rendered more salient by a comparison.
Interactive interpretation	Analogy interpretation depends on both terms. The same term yields different interpretations in different comparisons.
Multiple interpretations	Analogy allows multiple interpretations of a single comparison.
Cross-mapping	People typically perceive both interpretations of a cross-mapping and prefer the relational interpretation.

similarity can be expressed in terms of identical representational elements (e.g., cause(x,y) can match cause(x',y'), but not implies(x',y')). Here we assume that *sacked* and *defeated* have semantic subcomponents in common. This reliance on relational similarity does not mean that only identical elements can be placed in correspondence. Nonidentical elements can be matched if they are arguments to representational elements placed in correspondence by some other means. Here, the semantic match between *sacked* and *defeated* sets up the correspondences "Celtics–Xerxes" and "Lakers–Rome."

The second constraint is structural consistency. Structural consistency can be further broken down into the constraints of one-to-one mapping and parallel connectivity. Parallel connectivity states that if a pair of attributes or relations is placed in correspondence, then the arguments of those attributes or relations must also be placed in correspondence. One-to-one mapping requires that each element in one representation be matched to at most one element in the other representation. Parallel connectivity allows nonidentical representational elements to be placed in correspondence, as noted earlier. In this way, structure mapping permits analogies in which the objects are dissimilar or in which different psychological dimensions are placed in correspondence. We use the term *tiered identicality* to describe this pattern by which identical relational systems can cause nonidentical elements to be placed in correspondence.

Finally, systematicity (Benchmark 3) states that there is an implicit preference in analogy for matching connected systems of relations (Gentner, 1983, 1989). A matching set of relations interconnected by higher-order constraining relations makes a better analogical match than does an equal number of matching relations that are unconnected to each other. The systematicity principle captures a tacit preference for coherence and causal predictive power rather than for sets of coincidental matches.

Given an analogy that satisfies these three constraints, it is often possible to make further candidate inferences (Benchmark 4). These are statements true in the base that can be projected into the target and predicated (*mutatis mutandis*) of the target. In any such inference process, there must be some way to limit which information from the base is mapped to the target. Systematicity is the key factor here. Given a shared system between base and target, the candidate inferences are just the base predicates connected to the matching system and not yet present in the target. Finally, candidate inferences are only hypotheses; they must be checked for consistency with what is known about the target.

So far we have focused on the process that compares pairs of representations. When this process is completed, the resulting match permits the commonalities and differences of the pair to be found. The commonalities are just the identical elements that are placed in correspondence. In addition, there may be nonidentical elements that are placed in correspondence because they play the same role in a matching relational structure. For example, in Figure 1 there is a bed in the dorm room and a bunk bed in the jail cell. These elements correspond because both are beds, but they are different types of beds. There is evidence that these alignable differences—differences connected to the commonalities of a pair—are especially salient in analogy and similarity comparisons (Gentner & Markman, 1994; Markman & Gentner, 1993a, 1996). Alignable differences (Benchmark 5) can be contrasted with nonalignable differences, which are elements in one representation that have no correspondence in the other. For example, in Figure 1 there is a mug of beer in the dorm room and no mug of beer in the jail cell, making this a nonalignable difference.

The next two benchmarks express different aspects of the flexibility of analogical mapping. Indeed, Barnden (1994) suggests that analogy may be one way to reconcile the power of symbol systems with the flexibility of connectionist networks (see also Hummel & Holyoak, 1997). Interactive interpretation (Benchmark 6) captures the fact that different aspects of the same object or concept may be highlighted when it takes part in different comparisons. For example, William James (1985) noted that comparing the moon to a ball highlights the property that the moon is round, whereas comparing the moon to a lamp highlights the property that the moon is bright. Multiple interpretation (Benchmark 7) expresses the fact that the same comparison typically can be interpreted in different ways. For example, Gentner (1988) asked children to interpret double metaphors such as "A cloud is a sponge." Double metaphors permit both an object-based interpretation (e.g., both are fluffy) and a relational interpretation (e.g., both hold water). Gentner found that younger children tended to make (and to prefer) the object interpretation, whereas older children and adults could make both interpretations and typically preferred the relational interpretation.

A particular manifestation of this flexibility arises in cross-mappings (Benchmark 8). In a cross-mapping, objects in a pair that look similar play different roles in a matching relational structure (Gentner & Toupin, 1986; Markman & Gentner, 1993b). For example, one picture might show a car towing a boat, and a second might show a truck towing a car, where the two cars look similar. Most adults asked to compare these scenes can see both that there are similar cars and that the cars play different roles in the two scenes. People asked to find corresponding objects in a pair of scenes with cross-mappings generally prefer to map the objects that look similar to each other. However, if asked to compare the scenes (by making similarity or difference judgments), then they generally prefer to map the objects based on their relational similarity (Markman, 1996; Markman & Gentner, 1993b).

A computational model of comparison

How is the best alignment found? This problem is not trivial. The problem of placing two structured representations in correspondence (even excluding the candidate inference process) is one of matching two directed acyclic graphs. Graph matching is known to be in the class of NP-hard problems, meaning that the running time needed for any serial algorithm that is guaranteed to find the best match increases as an exponential function of the size of the domains being compared. Thus, any psychologically plausible process for finding analogical correspondences must either restrict itself to trivial problems (an unacceptable course) or simplify the solution process (at the risk of finding suboptimal matches).

One way to ease the computational burden is to assume that the top-level conclusion or goal of the analogy is known in advance (Greiner, 1988; Holyoak, 1985). This solution is unsatisfactory because people can process comparisons such as "Cigarettes are like ticking time bombs" without advance knowledge of their meaning (although of course such knowledge could be helpful for selecting between competing interpretations; see Gentner & Clement, 1988). A plausible computational simulation of comparison must be able to operate without advance knowledge of the final interpretation.

Another way to simplify the graph-matching process is to use semantic similarity. This is the tack taken in our simulation, the structure-mapping engine (SME; Falkenhainer, Forbus, & Gentner, 1986, 1989; Forbus, Ferguson, & Gentner, 1994; Forbus, Gentner, & Law, 1995; Forbus & Oblinger, 1990). SME arrives at a comparison of a pair of representations using a local-to-global alignment process.[3] Figure 2 shows the three stages of mapping in SME. In the first stage, SME begins blind and local by matching all identical predicates in the two representations. Semantic similarity between predicates is captured through a decomposition into partial identities. This initial mapping typically is inconsistent, containing many-to-one matches. In the second phase, these local matches are coalesced into structurally consistent clusters (called kernels). Finally, in the third stage, these kernels are merged into one or a few maximal[4] structurally consistent interpretations (i.e., mappings displaying one-to-one correspondence and parallel connectivity) fulfilling Benchmark 2. SME then produces a structural evaluation of the interpretations using a cascadelike algorithm in which evidence is

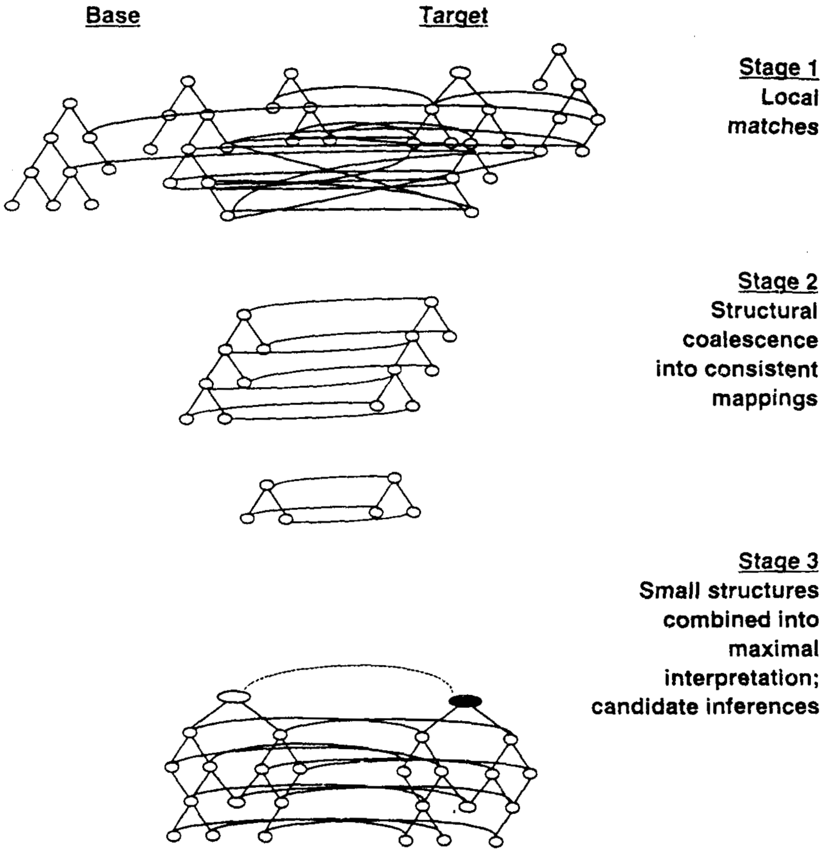

Figure 2. Stages of processing in the structure-mapping engine

passed down from predicates to their arguments, which favors systematic matches over unconnected systems of predicates (Benchmark 3; Forbus & Gentner, 1989). Finally, predicates connected to the common structure in the base but not initially present in the target are proposed as candidate inferences in the target (Benchmark 4). In this way, structural completion can lead to spontaneous unplanned inference.

SME produces more than one interpretation of a given comparison, consistent with Benchmark 7; typically, the two or three best interpretations are produced. For example, when presented with a pair of scenes that have a cross-mapping, SME generates two interpretations. One preserves the object-based similarity, and the second preserves the relational match. Because the relational match is more systematic than the object-based match, the structural evaluator assigns a higher score to the relational interpretation (Markman & Gentner, 1993b). This pattern is consistent with the observation that people typically prefer the relational correspondence when given a cross-mapping (Benchmark 8).[5]

An important extension to SME is the ability to perform incremental mappings, that is, to extend an existing analogical mapping by adding further connected material from the base domain (Forbus et al., 1994). This goes further than candidate inferences, which draw on the working memory representation of the base to fill out the common system in the target; in incremental mapping, further material is drawn from long-term memory or an ongoing stimulus presentation. The incremental mapping algorithm adds the new base material to the existing mappings, provided it is consistent with the correspondences already set up. Forbus, Ferguson, and Gentner suggested that incremental mapping is a psychologically plausible mechanism for processing extended analogies. This mechanism accords with the finding that extending a connected mapping is easier than creating a new one (Gentner & Clement, 1988; Gentner & Imai, 1992). This model was developed in part on the basis of the success of earlier incremental approaches (Burstein, 1986; Keane et al., 1994; Keane & Brayshaw, 1988).

A number of other computational models of the similarity–analogy process have been developed, including Holyoak and Thagard's (1989) ACME, Keane's (Keane & Brayshaw, 1988; Keane et al., 1994) IAM, Goldstone's (1994b; Goldstone & Medin, 1994a) similarity, interactive-activation, and mapping (SIAM), and Hummel and Holyoak's (1997) LISA. All of these models share certain basic assumptions. All assume that mental representations are structured. All further assume that structural consistency plays a role in mapping. Finally, all of the models except SIAM, which was intended primarily as a model of similarity judgments, have ways of generating new inferences. But the models differ in important respects as to how these factors play out in the comparison process. For example, SME takes structural consistency to be a strong constraint on analogy, whereas models such as ACME and SIAM treat structural consistency as only one of the several interacting pressures. In addition, SME assumes that all mappings place a strong emphasis on finding matching conceptual relations in the base and target. In contrast, in ACME, IAM, and LISA no semantic similarities are necessary for the mapping to take place. Finally, ACME and LISA allow people's pragmatic goals to influence the mapping process, whereas SME assumes that goals have their influence before mapping (by selecting relevant information from base and target) and after mapping (by influencing the evaluation and inference processes) but not during mapping. These models incorporate many of the assumptions underlying SME (Falkenhainer et al., 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997; Keane et al., 1994); other models of analogy (e.g., Halford et al., 1995; Hofstadter & Mitchell, 1994) operate on different principles. We return to this issue later.

Structure, similarity, and connectivity

With this theoretical framework in place, we now turn our attention to the empirical and theoretical implications of three central aspects of comparison: representational structure, the role of semantic similarity in mapping, and the importance of relational connectivity in comparison. We draw attention to these areas partly to pull together theory and data that have been presented in a variety of locations and partly to summarize and respond to recent theoretical disputes that have focused on these three issues.

Representational structure. As noted earlier, structure-mapping theory assumes that mental processes act on structured symbolic representations. It is assumed that people store enduring representations that preserve information about which elements take other as arguments. Given the prominence of nonstructural models in current cognitive theory (e.g., connectionist models that use vector representations), it is reasonable to ask whether a model that did not assume structured representations could account for people's behavior in comparison tasks. Although previous studies have provided evidence that common relational structure is important in analogy (Clement & Gentner, 1991; Gentner, 1988; Gentner & Clement, 1988; Gentner, Rattermann & Forbus, 1993) and similarity (Goldstone, 1994b), a direct concentrated test of the role of structure is worth pursuing.

To test whether structural relations are important in similarity processing, we carried out a similarity choice study using simple perceptual materials. Perceptual materials have the advantage that their presumed representations are less open to debate than those of causal scenarios. The idea is simply to give participants a series of choices and in each case to ask whether their choice can be captured by assuming a flat feature vector or feature set representation. Participants were shown eight forced-choice triads in a random order and were asked to say which of two figures was most similar to a standard. The triads are shown in Figure 3; in all cases, the comparison figure shown first is the one preferred by a majority of participants (at least 8 of 10 in this study, $p < .05$ by sign test).

The first triad verifies that people find configurations with similar objects in them to be more similar than configurations with dissimilar objects. Clearly this finding is compatible with a flat feature account. The next two triads (2a and 2b) show that people prefer configurations with the same relation to those without the same relation, even when the objects that play like relational roles are different. In Triad 2a, the

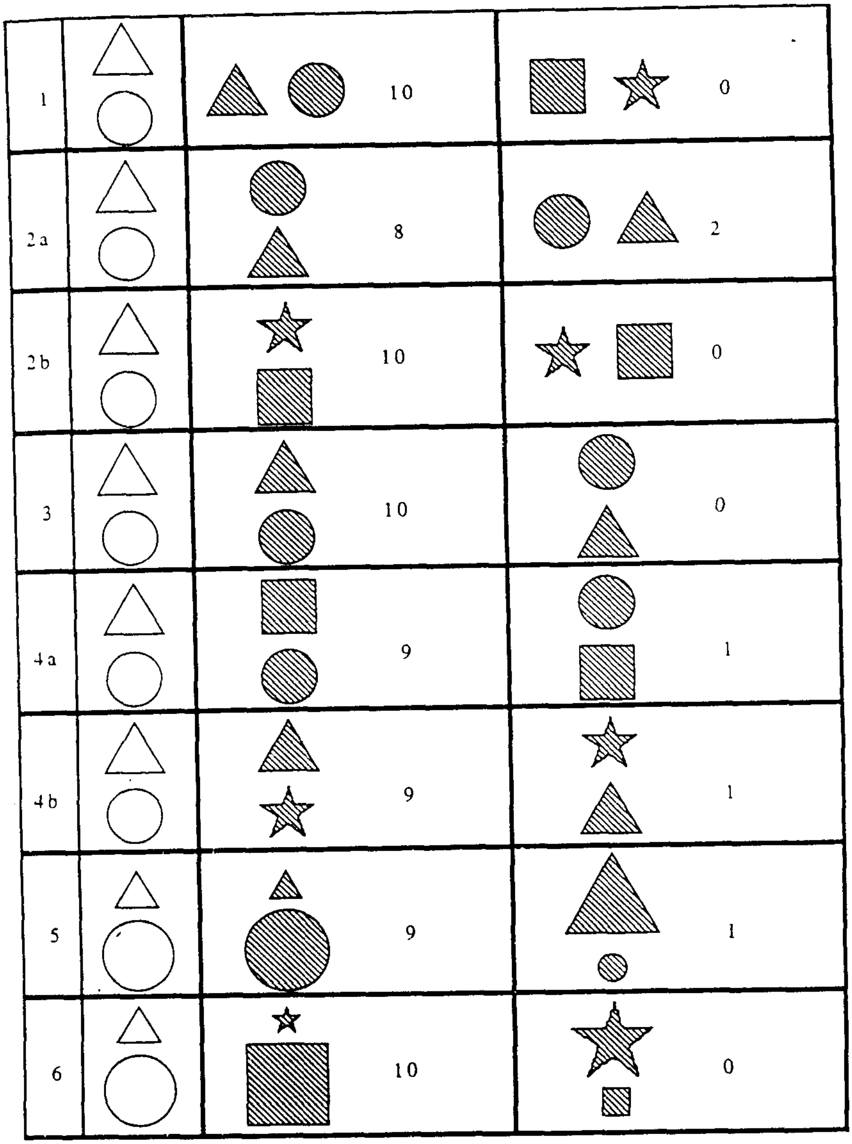

Figure 3. Stimuli in a two-alternative forced-choice task demonstrating that the mental representations used as the basis of comparisons are structured

winning figure has the same above relation as the standard, but the objects from the standard are reversed. Despite the cross-mapping, this figure is considered more similar to the standard than a figure that lacks the above relation (even though it shares objects with the standard). In Triad 2b, the shapes in the comparison figures are different from those in the standard. Again, these results are compatible with either featural representations or structured representations. A featural account would simply have to assume that the representations include the above relation as well as the triangle and the circle. So far these results do not require representing specific relational bindings.

Triad 3 begins to make the situation more complex. This triad demonstrates that having similar objects in the same relational roles makes a pair more similar than having similar objects playing different relational roles (that is, cross-mappings count against similarity). Although this evidence could suggest that relational bindings matter, a featural view can still account for this pattern by postulating the global configural feature above-triangle-circle. But Triads 4a and 4b push the argument a step further and show that this global configural feature is not enough. In these triads, having only one similar object playing the same relational role in the target and comparison figure makes the configurations more similar than having no similar objects playing the same relational roles. Thus, to maintain the feature view, further configural features must be added (e.g., circle-on-the-bottom and triangle-on-the-top) to account for people's choices in these triads.

Triad 5 pushes the issue further by demonstrating that people prefer consistency across a number of different relations in a scene. In the standard, the triangle is above the circle, and the triangle is also smaller than the circle. A comparison figure that preserves both of these relational commonalities is preferred to one that preserves only one of them. To capture this in a featural representation would require yet more configural features (e.g., smaller-triangle-circle). Finally, Triad 6 demonstrates that this preference for relational consistency holds even when all the objects are different. To capture this finding, the featural account must add yet more configural features to those listed above (e.g., smaller-top-object-bottom-object). At this point, the featural account becomes prohibitively cumbersome. In contrast, structure-mapping theory would capture the match in Triad 6 as an alignment between two structured representations:

Standard: AND [SMALLER(triangle, circle), ABOVE(triangle, circle)]
Match: AND [SMALLER(star, square), ABOVE(star, square)]

Even for simple perceptual materials, people's similarity judgments appear to be sensitive to relations and to relational structure (Gleitman, Gleitman, Miller, & Ostrin, 1996; Goldstone, 1994b; Goldstone & Medin, 1994a; Goldstone, Medin, & Gentner, 1991; Kotovsky & Gentner, 1996; Markman & Gentner, 1993b; Medin, Goldstone, & Gentner, 1993). Even for these simple materials, a featural approach would involve unwieldy representations. Relational bindings seem all the more necessary for describing the pictures in Figure 1 or for other complex items. Thus, simple models that do not permit bindings between representational elements are unlikely to provide a good account of psychological similarity in conceptual tasks (Gentner & Markman, 1993, 1995; Medin et al., 1993).

The importance of being connected. A second aspect of comparison that is critical from the perspective of structural alignment is connectivity. Not only do representations contain explicit bindings between representational elements, but these bindings influence what information people perceive to be important in a comparison. In this section, we first demonstrate that connectivity determines which commonalities in a comparison are important. Then we show that connectivity determines which differences are important. Finally, we show how systematicity governs inferences via structured pattern completion.

Connectivity and commonalities. Systematicity (Benchmark 3)—the presence of higher-order connections between lower-order relations—is important for determining which predicates enter into the interpretation of a comparison and how that interpretation is evaluated (Forbus & Gentner, 1989; Gentner, Rattermann, & Forbus, 1993). In general, a large number of different relational matches between two items could be considered. The focus on connected systems helps select which commonalities to pay attention to (Gentner & Clement, 1988).

Clement and Gentner (1991) provided direct evidence of the effects of systematicity—specifically, of causal connectivity—on analogical mapping. In their study, participants were shown pairs of rich passages describing an obvious analogy. In one pair, the first story described a robot taking in data, and its companion story described a space organism ingesting minerals. The passages contained two key facts, both of which were easily matched between the base and target but only one of which was connected to a matching causal antecedent. Participants were asked which of these two matches—the causally connected match (the shared system match) or the isolated match—was most important to the analogy. On 79% of their trials subjects chose the causally connected match. That is, they chose the match that was connected to a shared system of relations. For example, one of the key matching facts (italicized in the following excerpts) involved ceasing to use their intake systems to ingest food or information. When subjects received a matching fact (such as *stop using underbellies/stop using probes*) with a matching causal antecedent (as in the first two pairs that follow) they considered it highly important to the analogy:

(a) When the Tams exhaust the minerals in their rock, this causes them to *stop using their underbellies (to ingest minerals).*

(b) When the robots run out of new information on the planet, this causes them to *stop using their probes (to gather information).*

(c) When the robots malfunction, this causes them to *stop using their probes (to gather information).*

However, if the matching facts did not have matching causal antecedents, as in pairs (a) and (c), then participants did not consider the match important to the analogy. Thus, participants rated the matches as more important when they were connected to matching antecedents. In contrast, in a control condition in which participants were given only the target domain and asked to choose the most important fact, they showed no tendency to select these same key facts. Thus, the importance of a given fact to an analogy is determined not only by the degree of local match but also by connections to other matching pairs. These results support the claim that processing a comparison involves aligning interconnected systems of knowledge.

Connectivity and differences. Structure mapping also determines the which differences in a comparison are salient. In particular, differences that are connected to the commonalities (called alignable differences) are more salient than differences not connected to a matching structure (called nonalignable differences). We have tested four predictions that arise from the structural alignment view of alignable and nonalignable differences (Benchmark 5). The first two predictions turn on the relationship between alignable differences and commonalities. First, because alignable differences are connected to commonalities, pairs with many commonalities should yield many situations in which nonidentical representational elements are placed in correspondence. Thus, the number of alignable differences listed for a pair should be positively correlated with the number of commonalities that can be listed for it.[6] Second, alignable differences should be conceptually related to commonalities. The other two phenomena derive from the claim that alignable differences, by virtue of their connection to the matching system, are favored psychologically over nonalignable differences. Third, alignable differences should be listed more naturally than nonalignable differences. Fourth, alignable differences should matter more in judging similarity than should nonalignable differences.

Evidence for the first two claims has been obtained from studies in which participants list the commonalities and differences of pairs of words (Markman & Gentner, 1993a; Markman & Wisniewski, 1997) or pictures (Markman & Gentner, 1996). We find that highly similar word pairs, such as *hotel/motel*, have both more commonalities and more alignable differences listed than very dissimilar pairs such as *eggplant/giraffe*.[7] For example, Markman and Gentner (1996) showed participants 16 picture pairs, ranging from highly dissimilar to highly similar, like those in Figure 1 or Figure 4 (where the pair consisted either of Figure 4a and 4b or Figure 4a and 4c). Participants were asked to list either the commonalities or the differences (but not both) of the pairs. As expected, the number of alignable differences listed for these items was positive-

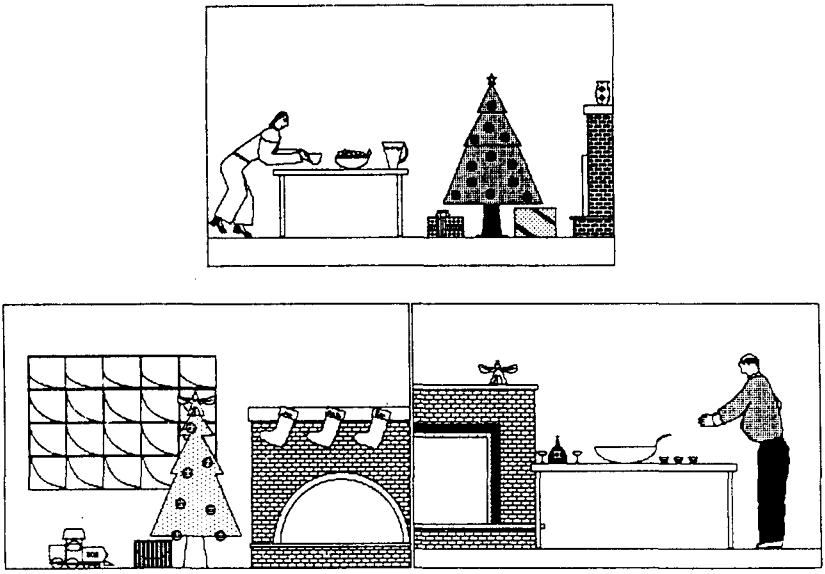

Figure 4. A triad of complex scenes used by Markman and Gentner (1996)

ly correlated with the number of commonalities listed. The studies with pairs of novel pictures are important because they demonstrate that this phenomenon is not a result of retrieving prestored differences from comparisons of existing concepts. Gunn and Gentner (in preparation) provided further evidence for this point by showing that people are more fluent at listing differences after carrying out a structural alignment (for both high- and low-similarity pairs).

The second prediction is that alignable differences are conceptually connected to the commonalities. A detailed analysis of the commonalities and differences revealed 72 consensus alignable differences (i.e., those listed by at least 5 of 10 participants). All of them were related to the commonalities given by the same participants. For example, all 10 participants who listed differences for pair 4a and 4b said that the Christmas trees (which are placed in correspondence via a commonality) have different ornaments. There were also examples of alignable differences that arose from relational commonalities. For example, 7 of 10 participants who listed differences for pair 4a and 4b said that there was a star on top of the tree in one picture and an angel on top of the tree in the other. In contrast, 7 of 10 participants who listed differences for pair 4a and 4c said that there was a vase on top of the fireplace in one scene and an angel on top of the fireplace in the other. The angels in Figures 4b and 4c are locally alike, but they do not correspond; they are put in correspondence (as alignable differences) with different objects, depending on the structural alignment.

The third phenomenon, that alignable differences are listed more fluently than are nonalignable differences, is manifest in two ways. First, most studies of commonality and difference listings in existing categories reveal more listed alignable differences for concepts than listed nonalignable differences. Second, structure-mapping theory predicts that participants should find it easier to list differences for pairs of similar items than for dissimilar items (because high-similarity pairs have many commonalities and hence many alignable differences). This rather counterintuitive hypothesis was tested by Gentner and Markman (1994), who presented participants with a page containing 40 word pairs, half similar and half dissimilar, and gave them 5 min to list one difference for as many different pairs as they could. Participants were told that they would not have time to do all 40 word pairs, and so they should do the easiest pairs first. Participants listed almost twice as many differences for similar pairs as for dissimilar pairs. This difference was concentrated in the alignable differences. Subjects listed few nonalignable differences and did so at about the same rate for similar and dissimilar pairs.

Finally, the fourth phenomenon, that variations in alignable differences affect perceived similarity more than do variations in nonalignable differences, is demonstrated in a study by Markman and Gentner (1996). Participants were asked to make similarity ratings on eight sets of four pairs like the one in Figure 5 in which the comparison figure at the top was paired separately with each of the other four pictures. In the first pair, the pictures were the same except that the car being fixed was replaced by a truck (an alignable difference). In the second pair, the car was replaced by a robot being fixed (also an alignable difference). Assuming that participants would see the truck as more similar to the car than the robot, the prediction is that the rated similarity with the standard should be greater for the truck picture than for the robot picture. This result was indeed obtained ($M = 6.54$ for the truck, $M = 4.75$ for the robot). But when same two items (the truck and the robot) were added to the scene as nonalignable differences, as in the bottom two pictures in Figure 5, we predicted little if any difference in the rated similarity (because the nonalignable differences were expected not to matter much in the similarity comparison). Again, this result was obtained: the rated similarity with the standard was about the same for both the bottom pictures ($M = 6.17$ for the truck, $M = 5.58$ for the robot). This finding suggests that, as predicted, participants attend more to alignable differences when rating similarity than to nonalignable differences.

These findings suggest that alignable differences are focal in similarity comparisons. As an additional test of this hypothesis, Markman and

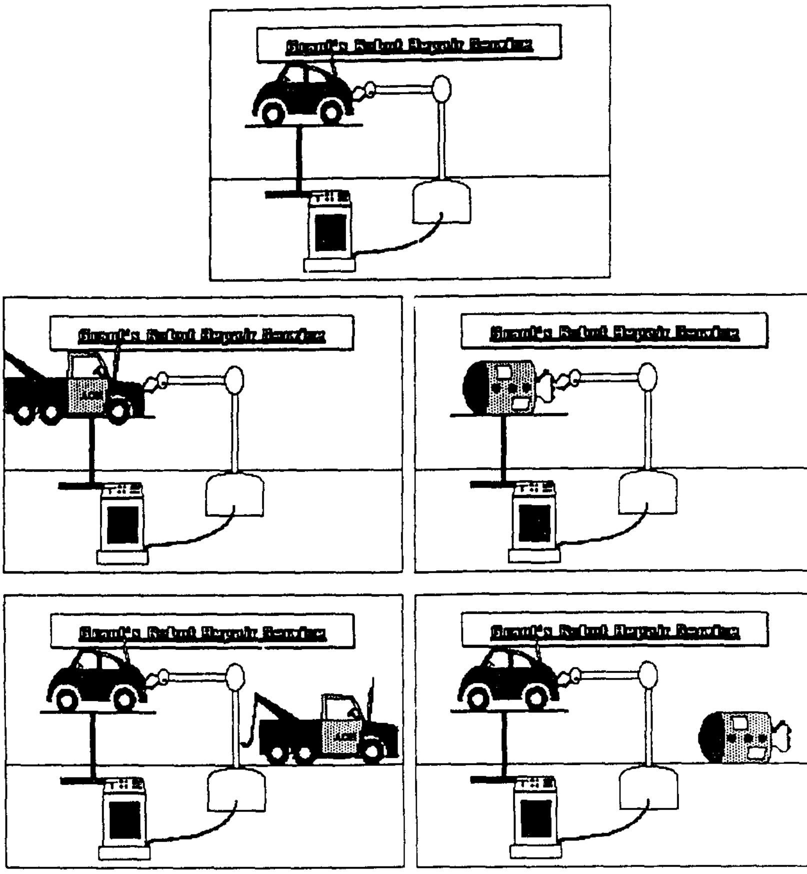

Figure 5. Sample stimuli from a similarity comparison task used by Markman and Gentner (1996)

Gentner (1997) had participants rate the similarity of pairs of scenes. In these scenes, some objects were alignable differences and some were nonalignable differences.[8] After a 30-min delay, participants were shown an object that was an alignable difference or an object that was a nonalignable difference and were asked to recall as many aspects of the picture from which the object came as possible. Participants remembered much more when the recall cue was an alignable difference than when it was a nonalignable difference, suggesting that they focused primarily on alignable differences when judging similarity.

This pattern may seem a bit paradoxical. We are suggesting that the greater the similarity of two items, the greater the number of alignable differences. But because alignable differences are psychologically salient, the more similar two things become, the more salient their differences become. We suggest that this pattern is reasonable. The world is full of differences to which we *could* pay attention: A chair has no seeds, an apple has no seat. The comparison mechanism seems to highlight differences that are informative—differences that exist against a web of similarity.

These findings are difficult to accommodate in nonstructural models of similarity. They are particularly problematic for mental distance models, in which difference is simply the inverse of similarity. In such models, similarity and difference both are represented by mental distance. The closer (the more similar) two things are, the less far apart (the less different) they are. It is hard to see how such models can capture the finding that higher similarity typically is accompanied by greater numbers of salient differences. Featural views fare somewhat better. They can readily accommodate the separation of features into commonalities and differences and, as Tversky (1977) showed, can even accommodate cases in which a given pair (e.g., Canada/United States) is considered both more similar and more different than another pair about which less is known (e.g., Ceylon/Nepal). However, an independent feature account has no way to distinguish between alignable and nonalignable differences or to predict the positive relationship between the number of alignable differences and the number of commonalities. Thus the patterns of findings for differences are strong evidence of structural alignment.

Connectivity and inferential pattern completion. According to structure-mapping theory, systematicity not only determines which commonalities and differences are important but also governs which candidate inferences are imported from base to target, thereby influencing how the target domain is extended by comparison to the base (Benchmark 4) (Falkenhainer et al., 1989). To make this more concrete, we show a sample base and target domain in Figure 6. The predicates and objects in this domain are represented by nodes of the graph, where the label on each node is denoted by the pattern that fills it. Correspondences between base and target are shown by arrows connecting a node in the base to a node in the target.

When there is matching structure in the base that is connected to matching facts in the target, and the new facts are structurally consistent with the existing match, then the novel information may be carried from base to target as a candidate inference. When information is carried from base to target, the relational structure is copied into the target domain. Substitutions are made for any representational element from the base for which there is already a matching element in the target (Falkenhainer et al., 1989; Gentner, 1982, 1983; Holyoak, Novick, & Melz, 1994). As shown in Figure 7, three elements can be carried over

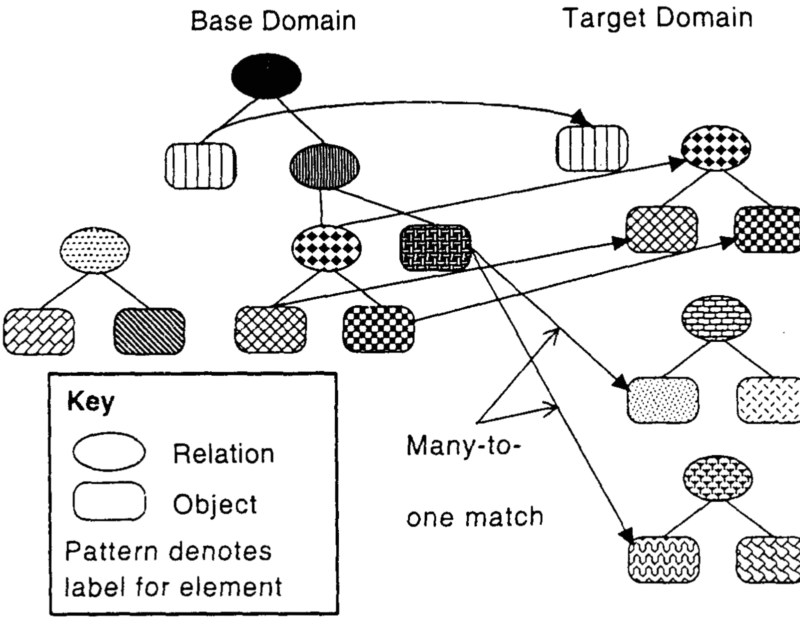

Figure 6. Relational structure with correspondences between the base and target

in this manner. In this example, the two relations carried over (shown in ovals) have no correspondence in the target and are carried over directly. The object (shown in the rounded rectangle) has two possible correspondences in the target, so one of the matching elements in the target must be substituted for this object. Leaving aside for the moment how to resolve which object should be substituted this analysis suggests that there should be a preference for making inferences that are shared system facts rather than nonshared system facts, similar to the greater weight given to shared system facts over nonshared system facts as commonalities, which we discussed earlier.

Studies of analogical inference have demonstrated that shared system facts are far more likely to be inferred than nonshared system facts (Clement & Gentner, 1991; Lassaline, 1996; Markman, 1997a; Spellman & Holyoak, 1996). In one study, Clement and Gentner (1991) took pairs of analogous stories like those described earlier and modified them by removing the two key matching lower-order facts from the target passages. Thus, for each story pair there were two key facts in the base that did not appear in the target. As before, one of these facts was connected to a matching causal antecedent, but the other was not. Participants given these story pairs and asked to make an inference of a new fact

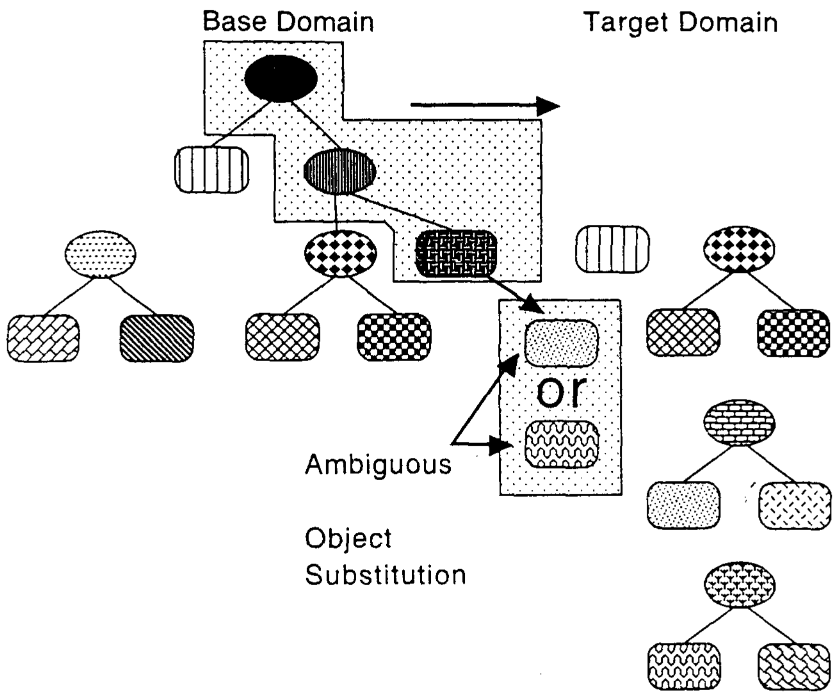

Figure 7. An inference that can be drawn from the relational structure in Figure 6

tended to infer the shared system fact rather than the nonshared system fact. A similar finding was obtained by Spellman and Holyoak (1996), who asked participants to extend the plot of one soap opera based on the plot of a second. Furthermore, they demonstrated that a pragmatic hint to focus selectively on a particular plot element was also effective at constraining the information that was inferred. Lassaline (1996) and Wu and Gentner (1998) extended these findings to category-based induction (e.g., "If robins have sesamoid bones, how likely is it that bluebirds have sesamoid bones?") and showed that these inductions also are governed by common relational structure.

A difficulty can arise in making substitutions in an inference when the match between base and target is many-to-one because it is not clear which target element should be substituted in the inference (Gentner, 1982; Markman, 1997a). In the example in Figure 7, there are two possible correspondences for the object in the base. Unless one-to-one mapping is strictly enforced (as it is in structure-mapping theory), it is possible that different target elements will be substituted into the inference for the same base element in different parts of the inference. Studies by Markman (1997a) presented participants with base and target domains with potential many-to-one mappings and asked them to infer facts about the target. Despite the potential for ambiguity in object substitutions in these studies, participants always maintained one-to-one mapping in their inferences. Thus, people resolve potential ambiguities in inference by maintaining one-to-one mapping when they are making inferences.

Gentner (1982) pointed out an additional danger in having many-to-one correspondences in the case where inferences are going to be made. If inferences are made on the basis of shared structure and if there are many-to-one mappings, then there will be a proliferation of candidate inferences because each combination of correspondences will give rise to a candidate inference. Because most analogical inferences are likely to be wrong, a mechanism that generates a lot of inferences is not very useful.

Systematicity also sheds light on the well-known phenomenon of directional asymmetry in comparison (e.g., the finding that people prefer "Hungary is like Russia" to "Russia is like Hungary"). Bowdle and Gentner (1997) suggested that such asymmetries typically arise when the matching system that forms the interpretation of a comparison is more systematic for one item than for the other. This proposal of systematicity imbalance derives from two sources: the structure mapping claim that inferences are projected from more systematic bases to less systematic targets and the pragmatic point that comparisons are meant to be informative. Gleitman et al. (1996) made the related point that informativity considerations partly drive the choice of which object is chosen as the base or ground for a comparison. Thus, given that two comparison items are alignable, the more systematic and coherent item should be preferred as the base to maximize the degree to which information can be mapped to the target.

This hypothesis was tested by presenting participants with pairs of literally similar passages in which one passage was designed so that the central relational structure contained more systematically interconnected relational structure than the other (Bowdle & Gentner, 1997). When asked to give their preferred order of comparison, participants consistently preferred comparisons in which the more coherent passage was the base and the less coherent passage was the target. Furthermore, when asked to generate inferences from one passage to the other, participants drew inferences from the more coherent passage to the less coherent one, consistent with the preferred direction of comparison. Thus, the preferred direction of a comparison is the one that permits the more coherent domain, which better serves as a source of inferences, to be the base.

Semantic similarity in mapping. The final key aspect of systems of analogical reasoning is the role of semantic similarity, not only psychologically but computationally. Computationally, carrying out an analogical match entails finding a good (partial) alignment between two working memory representations. As discussed earlier, the problem of matching two graphs is computationally intractable. Models of analogy must therefore make simplifying assumptions to explain how analogies are processed. Obviously, the closer these simplifications come to those used in human processing, the better. The simplifying assumption made by structure-mapping theory is that analogical matches are rooted in relational semantic similarities between domains (Benchmark 1). The early identities found between semantic components constrain the possible graph matches. However, although the claim that semantic similarity is important in analogical mapping might seem obvious, it has generated a fair amount of controversy. We begin by laying out the issues and challenges. Then we give evidence for the role of relational similarity. Finally, we discuss some current challenges in modeling representational similarity.

The importance of semantic similarity. Structure-mapping theory (and its implementation in SME) is committed to the claim that analogous domains must have some identical relational elements.[9] In contrast, models such as LISA (Hummel & Holyoak, 1997), ACME (Holyoak & Thagard, 1989), and IAM (Keane & Brayshaw, 1988) permit content-free correspondences that are similar only on the basis of the graph structure of the match. To make this issue more concrete, we consider comparisons of three sentences. First, the sentence

The Nazis invaded France, causing people to flee France. (4)

can be seen as similar to

The rats infested the apartment, causing the people to leave the apartment. (5)

because both have semantically similar relations in an identical syntactic structure. However, 4 and 6 also share a graph match

The game show host kissed the contestant, inviting the audience to applaud the contestant. (6)

In sentences 4 and 6, none of the relations that hold between the objects are the same. However, each sentence has an abstract relational structure that could be written as

Higher-order Relation1_2 [Relation1 (X, Y), Relation2 (Z, Y)]. (7)

Based on this abstract relational structure, the Nazis and the game show host can be placed in correspondence (both correspond to X in the abstract structure shown in sentence 7). Likewise, the other objects in the sentences could likewise be aligned based on the graph structure of the relations (e.g., France–contestant, people–audience). But this purely syntactic match leaves most people unmoved; intuitively, these two sentences are not analogous.

Structure-mapping theory conforms to this intuition. According to structure mapping, analogy requires that at least some of the relations be conceptually identical; thus sentences 4 and 5 form an analogy, but sentences 4 and 6 do not. In the comparison of 4 and 5, the Nazis correspond to rats because both play the role of invaders. Likewise, France corresponds to the apartment, because both are invaded (and fled from). As discussed earlier, objects that are not identical may be placed in correspondence if they play the same role within a matching relational structure; relational matches legitimize nonidentical object correspondences. Structure-mapping theory thus requires there to be similar semantic relations between base and target. Furthermore, the process model (SME) uses initial filtering by semantic similarity as a way to avoid the intractable general graph-matching problem.

An alternate position, most clearly articulated by Keane (1997), is that analogical matches can occur without semantic overlap (see also Holyoak & Thagard, 1989; Keane & Brayshaw, 1988). In a set of recent studies, Keane (1997) gave participants sets of sentences like those in Figure 8a and asked them to indicate which sentence pairs correspond by drawing lines between the sentences on the left and their corresponding sentences on the right (see Holyoak & Thagard, 1989, for similar materials). The sentences shared no semantic overlap, but they did have syntactic overlap. As shown in Figure 8b, each sentence on the left has a syntactically corresponding sentence on the right. The dependent measure was time to arrive at the correct sentence-matching solution. Keane's results suggest that people find corresponding sentences in a task like this by mapping the sentences in a sequential fashion. For example, the underlined sentences in Figure 8b are singleton sentences because the object in each sentence is a unique item. Thus, finding a singleton pair sentence specifies one object match, reducing the number of solutions that must be considered for the remaining matches. Consistent with this reasoning, Keane found that having singleton sentences that appeared in the same relative location in the two sentence sets greatly eased the task of finding the corresponding sentences.

Keane (1997) argued that people's performance on materials like these can be used to constrain models of analogy (although he conceded that these materials represent difficult, borderline cases of analogy). We disagree with Keane's assertion that these materials are processed as analogies, even as borderline analogies. The matter is not merely

(A)

List A	List B
Jim kisses Mary	Ruth motivates Debra
Jim loves Mary	Ruth knows Debra
Bill loves Mary	Laura motivates Debra
Bill is jealous of Jim	Laura waves to Ruth

(B)

List A	List B
Jim kisses Mary	Ruth motivates Debra
Jim loves Mary	Ruth knows Debra
Bill loves Mary	Laura motivates Debra
Bill is jealous of Jim	Laura waves to Ruth

(C)

List A	List B
Jim kisses Mary	Ruth adores Sam
Jim loves Mary	Ruth hugs Sam
Bill loves Mary	Laura adores Sam
Bill is jealous of Jim	Laura envies Ruth

(D)

List A	List B
Jim kisses Mary	Ruth motivates Sam
Jim loves Mary	Ruth knows Sam
Bill loves Mary	Laura motivates Sam
Mary works with Bill	_______________
Bill is jealous of Jim	

(E)

List A	List B
Jim kisses Mary	Ruth hugs Sam
Jim loves Mary	Ruth, adores Sam
Bill loves Mary	Laura adores Sam
Mary works with Bill	_______________
Bill is jealous of Jim	

Figure 8. Stimuli from a sentence-mapping task

terminological but bears on the processing conclusions to be drawn about analogical processing. As discussed earlier, if there are no semantic constraints on analogy, then the mapping problem reduces to the matching of arbitrary graphs. In such cases, as in the sentence sets in Figure 8a, Keane probably is correct that the solution involves a sequential process that examines one sentence at a time and tries to draw out a purely structural solution, much as one would solve a logical problem on an intelligence test. Processing an analogical comparison such as "The atom is like the solar system" does not seem to involve this fragmentary processing, however; rather, one has the sense of a rapid rush of converging matches. Structure-mapping theory suggests that the difference between these types of comparisons is that the presence of identical predicates in the base and target allows an early parallel matching in which preliminary local matches are made. These local matches can then be coalesced into global structurally consistent matches. The important computational point here is that in this process, only a small set of object matches is ever considered (i.e., those mandated by predicate similarity). On this account, cases in which there are no semantic similarities are processed completely differently; the early local match process cannot be used to constrain the graph match.

The psychological question here is whether empty structural matches receive different processing from conceptual matches, as claimed in structure mapping, or whether both kinds of matches receive the same kind of processing, as claimed by Keane (1997). To test whether people treat these types of problems differently, we used a task like the one described by Keane (1997) but varied whether the relations in the sentence sets were conceptually similar (Gentner & Markman, in preparation). Thus we compared materials like those in Figure 8c, which shared semantic similarity, with materials that did not, as in Figures 8a and 8b. For the sentence sets that had only an empty syntactic match, our results were much like those obtained by Keane. Participants were faster and more accurate at finding the correct correspondences when the singleton matches were in the same relative location than when they were not.

In a further study, a pair of sentence sets was developed such that it had both a conceptual match and a syntax-only match. One story was about two companies, Fox and Time–Warner, that were competing for a resource and were also trying to hire away each others' employees. The second story was about two friends, Andy and John. Andy and John both liked baseball, and both of them were competing for a resource. Obviously, the stories were similar in that both involved a competition for resources. Furthermore, the syntactic structure of the relations involved in the businesses hiring employees and the friends enjoying baseball were the same, although the content of these relations was different. Finally, the story about the companies had more information in it than did the story about the friends. Participants were asked to make an inference about what might happen in the story about the friends given what was known about the companies. The inferences made were exclusively based on the conceptual match about the competition for resources. No inferences were made that were related to the identical syntactic match involving the companies' hiring practices and the friends' love of baseball. These results suggest that finding correspondences in sets with conceptual similarity is done quite differently from finding correspondences in sets with only syntactic matches.

Similarity and identicality. A commitment to semantic similarity leaves open exactly how that similarity should be implemented. As discussed earlier, SME requires that predicates be identical in order to be placed in correspondence. To place semantically similar but nonidentical predicates (such as give(x,y) and trade(x,y)) in correspondence, SME must decompose them into subcomponents, some of which will be identical. This might seem rather cumbersome; why not simply allow similar predicates to match, with the degree of match being determined by their similarity? This second method was used in Holyoak and Thagard's (1989) ACME; a similarity table was used to assess the degree of similarity between predicates, and this in turn influenced the strength of the correspondence between the two predicates. Similarly, in Hummel and Holyoak's (1997) LISA, predicates are represented with connectionist activation vectors; the similarity between any pair of vectors is predefined by the user. Nonidentical predicates that have similar activation vectors can be placed in correspondence during mapping.

Although the direct use of semantic similarity is intuitively appealing, it has some potential costs. First, there are many semantically similar predicates that normally should not be placed in correspondence. For example, taller(x,y) and shorter(x,y) are very similar in meaning—in that both convey relative height—but if John is taller than Scott, then Scott is shorter than John. Simply allowing two similar predicates to be placed in correspondence misses this difference. Second, the differences between predicates may actually be important in reasoning. For example, if buy(x,y) and trade(x,y) are placed on correspondence based on their overall similarity, then the comparison process will gloss over the fact that buying involves transfer of cash for the exchange, whereas trading involves transfer of a good (Gentner, 1975, 1981). Third, using similarity tables misses the possibility of common abstractions that can emerge from the comparison, as discussed later.

If we assume that semantic similarity is cashed out in terms of identical subpredicate correspondences, then when nonidentical predicates are encountered, they are placed in correspondence only if they can be re-represented as identical (or at least partially identical; Clement, Mawby, & Giles, 1994; Forbus et al., 1995; Gentner & Rattermann, 1991). For example, nonidentical predicates might be decomposed into partially identical structures. In the case of trade(x,y) and buy(x,y), this decomposition would reveal commonalities such as the transfer of a good from one person or a reciprocal exchange. These commonalities, we suggest, may emerge from the comparison as common abstractions.

Often, as with *trade* and *buy*, these decompositions are routine. But sometimes re-representation involves conceptual insight. Nonidentical predicates that originally could not be matched come to be seen as similar after comparison experience. Gentner, Rattermann, Markman, and Kotovsky (1995) gave an example of this type of re-representation in children. They described studies in which children compared simple stimuli like those in Figure 9. In Figure 9a, the match is straightforward. The largest circle on the top corresponds to the largest square on the bottom, the medium-sized circle with the medium-sized square, and the smallest circle with the smallest square. Children as young as 4 can make these comparisons (Kotovsky & Gentner, 1996). Children find the comparison in Figure 9b more difficult. Here, the largest circle corresponds to the darkest square, the medium-sized circle with the medium-darkness square, and the smallest circle with the lightest square. Children do not make these cross-dimensional comparisons reliably until they are 8 years old unless they are given training. If the squares are labeled (e.g., "This is the Daddy square, this is the Mommy square and this is the Baby square") or if the children are given extensive practice on within-dimension comparisons like the one in Figure 9a for both relevant dimensions, then even 4-year-old children can make cross-dimensional comparisons. Gentner et al. (1995) suggested that children are learning to re-represent predicates like taller(x,y) and darker(x,y), that are initially tied specifically to a dimension, as greater[height(x),height(y)] and greater[darkness(x),darkness(y)], which makes the relational commonality in these predicates manifest. This type of re-representation is difficult, but it may promote consistency in representation across domains (Gentner & Rattermann, 1991).

Routine re-representation serves the same function as the similarity metrics in other models of analogy. It provides a quick, effortless way to re-represent nonidentical predicates as being partially identical. Like the similarity tables, it relies on stored structure. Decomposition requires that the way in which predicates can be decomposed is already known. Minimal ascension requires an existing abstraction hierarchy. Conceptual re-representation is the only truly creative form of re-representation. Unfortunately, this type of re-representation is also the one

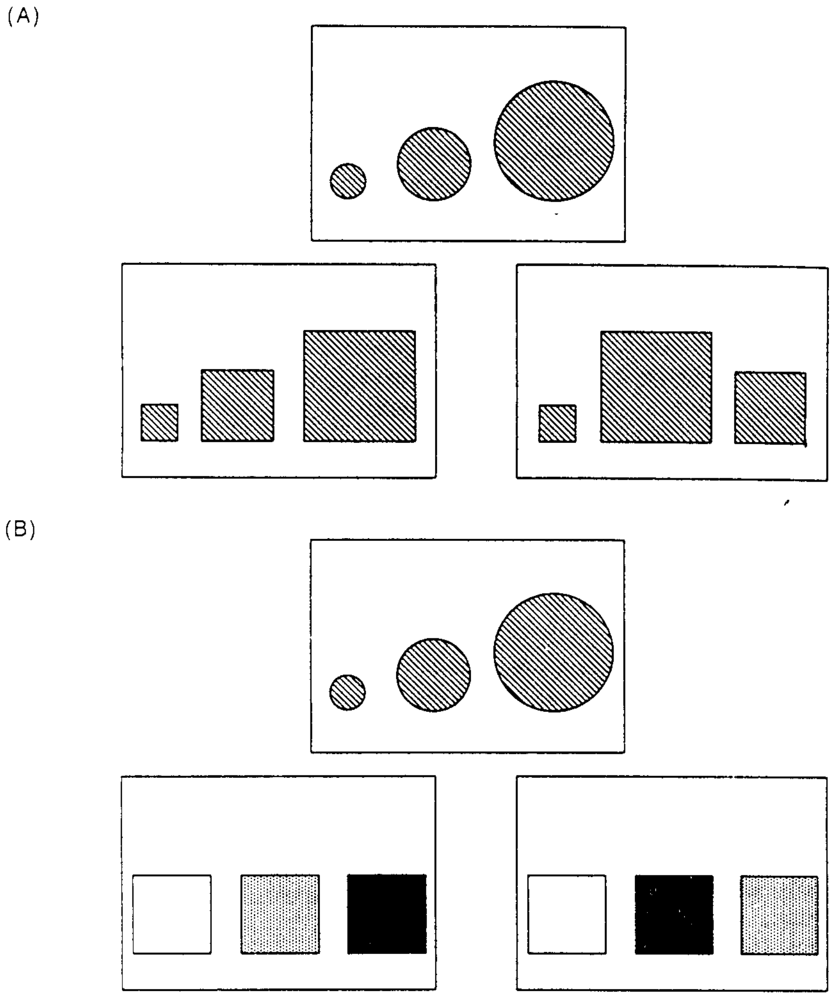

Figure 9. A two alternative forced-choice tasks given to children

we know the least about. However, examples such as the one described above make us think that it is an important mechanism of conceptual change.

Both re-representation and the use of a similarity metric between predicates are designed to allow nonidentical predicates to match. It is not clear whether both of these processes can operate within a single model. Further research must examine how re-representation processes can be integrated into the comparison process.

Conclusions

Research on analogy and similarity has benefited from the research interactions promoted by the cognitive science movement. Psychological theory has been developed and extended through the use of computational modeling. These computational models use sophisticated symbolic and connectionist techniques. The psychological implications of these models have been tested, leading to changes in the psychological theories and computational models.

In this article we have summarized structure-mapping theory and have focused on three implications of the theory: first, that mental representations are structured; second, that connectivity plays a key role in determining what information in a comparison is important and what information should be carried from one domain to another as an inference; and third, that analogical matches are based on semantic similarities.

Limits of structure. Although structured representations are critical for modeling similarity judgments, structure is not manifest in all cognitive tasks that involve similarity. Computational models of mapping suggest that the process of computing matches between pairs of structured representations is computationally intensive. Consistent with this suggestion, there is evidence that very rapid similarity comparisons may fail to show structural sensitivity. For example, Ratcliff and McKoon (1989) asked participants to study a set of simple subject–verb–object sentences (e.g., "John kissed Susan"). Later, the participants were given a memory test in which they had to respond "new" or "old" to sentences under a deadline. Participants had no difficulty in quickly rejecting sentences that contained words that did not appear in the training set (e.g., "Reginald waved to Jennifer"). However, at short latencies they were not able to reject sentences that contained words that did appear in the training set but were bound differently than in the original sentences (e.g., "Susan kissed John"). Only at longer response latencies could participants correctly reject both types of sentences. This finding suggests that people need time to represent and compare pairs of items that have relations and arguments. Likewise, Goldstone and Medin (1994b) showed participants pairs of objects with perceptual relations that contained either cross-mapped local correspondences or consistent global correspondences. At rapid deadlines, participants were prone to incorrectly accept the local cross-mappings as identical; in contrast, at longer deadlines, global correspondences were more important in their responding. These findings suggest that early processing involves local correspondences and that global matches emerge gradually from the local matches.

Another case in which structure is given short shrift is similarity-based retrieval from long-term memory. Studies of retrieval have found that retrieval is more likely when the cue shares attributes with the original material than when it shares relational structure (Gentner et al., 1993; Gick & Holyoak, 1980, 1983; Holyoak & Koh, 1987; Ross, 1987, 1989). Partly for this reason, some models of similarity-based retrieval involve two stages (e.g., Forbus et al., 1995; Thagard, Holyoak, Nelson, & Gochfeld, 1990). In the initial stage, the cue is compared to the contents of memory using a process that finds common representational elements without regard to how those elements are bound together. Only after the initial stage are the bindings taken into account.

The idea that structure, although normally crucial, can sometimes be neglected in comparisons may connect with Sloman's (1996) proposal of a two-system model of cognition. In this model, a system for rule-based judgments coexists with an associative system for storing and retrieving associations and similarity relations. The rule-based system is responsible for deliberative, typically accurate judgments, whereas the associative system carries out fast but error-prone judgments. The findings just discussed suggest that similarity comparisons may participate in both rapid error-prone cognition and structure-sensitive cognition.

To summarize, a central tenet of models of analogy is that representations are structured. This assumption is supported by evidence about the way people make comparisons, even when these comparisons involve simple stimuli such as geometric configurations. Although structure sometimes is disregarded in speeded comparisons, processes, or memory retrieval, comparisons normally involve alignment of structured representations.

A key test of any psychological model of similarity is its ability to shed light on key issues about other cognitive processes that involve comparison. For example, similarity often is given a central role in categorization (Hampton, 1995; Rosch, 1975; Smith & Medin, 1981). However, Rips (1989) and others have found cases in which rated similarity and probability of category membership are disassociated. Furthermore, a strong case has been made that category membership judgments are heavily theory-based (Keil, 1989; Murphy & Medin, 1985). For example, Keil (1989) demonstrated that 6-year-old children will categorize a cat who has been painted black with a white stripe and has had a pouch of smelly oil sewn into it as a cat rather than as a skunk. Gentner and Namy (1999) showed that alignment processes can facilitate 4-year-olds' categorizing according to deep relational commonalities rather than perceptual commonalities. All this raises the question of the precise role of similarity in relation to explanation-based or theory-based reasoning about concepts.

Goldstone (1994a) suggested that thinking about similarity as a process of comparison rather than as a relation that holds between pairs of items is crucial for integrating similarity-based and theory-based explanations of categorization (see also Gentner & Medina, 1997; Markman, 1997b). Many similarity-based models of categorization focus on simple representations of categories, like featural representations of exemplars (e.g., Medin & Schaffer, 1978; Nosofsky, 1986, 1987). If instead it is assumed that item representations include information about relations between features as well as the features themselves, then at least some of the theory-based knowledge could be included in the exemplar representations. In this case, comparisons of exemplars would reveal relational commonalities, which often reflect deeper theoretical knowledge about concepts. The inclusion of this relational information might be able to account for some systematic patterns of within-category variation. For example, Wattenmaker, Dewey, Murphy, and Medin (1986) showed that people can learn a nonlinearly separable pair of categories more easily than a linearly separable pair if they are given relations connecting the features within the categories. In a sense, the addition of a relation transforms the problem into a linearly separable distinction between two relational structures.

Structural alignment also sheds light on the processes underlying choice behavior. Decision makers must find the features of the options relevant to them. Some research suggests that alignable differences are given more weight in choice situations than are nonalignable differences (Lindemann & Markman, 1996; Markman & Medin, 1995; Slovic & MacPhillamy, 1974). In one study, participants made choices between video games and later justified their choices. Their justifications were more likely to contain alignable differences than nonalignable differences. As another example, Kahneman and Tversky (1984) described to participants a hypothetical store in which a jacket could be bought for $125 and a calculator for $15. They offered them the opportunity to go to another store and thereby save $5 on the total purchase. Participants who were offered a jacket for $125 and a calculator for $10 were more willing to spend the effort to go to another store than those offered a jacket for $120 and a calculator for $15. Participants were influenced by the alignable difference comparisons and failed to notice that the expected values were the same in both situations.

The focus on alignable differences in decision making also can influence what people learn about the choice options. In one set of studies, Zhang and Markman (1998) taught participants about three brands of microwave popcorn; one brand was learned in an initial session, and the other two brands were learned in a session 2–5 days later. The distinctive properties of the brands were either alignable differences (e.g., pops in a special bag versus pops in a microwave-safe bowl) or nonalignable differences (e.g., does not stick in teeth versus low in corn flavor). Participants were far better able to recall the distinctive attributes of the brands presented in the second session when they were alignable differences than when they were nonalignable differences. Participants' ability to recall properties about the brands also was related to their preference for the brands. To the degree that alignable differences are important in choice, it means that decision making is influenced by the extent to which properties of options are interconnected.

In sum, we view structure mapping not just as a convenient way to characterize how people process analogical comparisons but as a fundamental cognitive process that influences diverse aspects of higher cognitive processing. The ability to compare pairs of structured representations is critical for permitting rule-based reasoning, finding analogies between domains, constructing preferences, and solving complex problems. Furthermore, deep conceptual change can take place in service of making representations of distant domains more comparable, thereby promoting consistency in representational structure. Future research must examine how the alignment process interacts with other fundamental cognitive processes, such as working memory and attention, to provide a more unified view of cognition.

Notes

This research was supported by NSF CAREER award SBR-95-10924 to Arthur Markman, as well as NSF grant SBR-9511757 and ONR grant N00014-92-J-1098 to Dedre Gentner. The authors thank Keith Holyoak, John Hummel, Mark Keane, the Similarity and Analogy Group, and two anonymous reviewers for helpful comments and discussions.

Correspondence about this article should be addressed to Arthur B. Markman, Department of Psychology, University of Texas, Mezes Hall 330, Austin, TX 78712 (e-mail: markman@psy.utexas.edu). Received for publication March 17, 1997; revision received May 4, 1999.

1. Hummel and Holyoak (1997) proposed a similar list that includes, in addition to the phenomena listed by Gentner and Markman (1995), pragmatic information in mapping and phenomena relating to analogical retrieval.

2. To say that the relations must be identical (or at least partially identical) for the analogy to hold does not mean that the actual surface words used to express the relation must be identical. Rather, the relational concepts must be partially overlapping (see Gentner & Clement, 1988).

3. Similar algorithms have been incorporated into other computational models of analogy, although none is identical to SME (Burstein, 1986; Goldstone, 1994b; Goldstone & Medin, 1994a; Holyoak & Thagard, 1989; Keane, Ledgeway, & Duff, 1994).

4. SME uses a greedy algorithm (Forbus & Oblinger, 1990) that normally, but not invariably, finds the maximal interpretation (i.e., the largest and deepest interpretation, with the highest evaluation score).

5. Across comparisons, different aspects of a given item are used in a comparison, consistent with Benchmark 6. Furthermore, SME permits nonidentical elements to be placed in correspondence when they play the same role in a matching relational structure, thereby calculating alignable differences (Benchmark 5). Thus, SME can account for all eight benchmarks in Table 1.

6. This relationship is expected for all pairs except for identical pairs, which should have many commonalities and no differences.

7. To avoid bias in categorizing the differences, in these studies we used a purely formal criterion. A difference was scored as alignable if the same predicate or dimension is applied to both terms (with different values); for example, there is a single bed in the dorm room in Figure 1 but a bunk bed in the jail cell. A difference was scored as nonalignable if an assertion is made about one item and denied about the other (e.g., there is a mug of beer in the dorm room but not in the jail cell).

8. The materials were counterbalanced so that an object that was an alignable difference for some comparisons was a nonalignable difference for another. This design controlled for the salience of the individual objects.

9. Ironically, because of the emphasis on structural parallels in structure-mapping theory, sometimes it has been assumed that structure-mapping theory ignores semantic similarity among corresponding relations. In fact, the opposite is the case.

References

Barnden, J. A. (1994). On using analogy to reconcile connections and symbols. In D. S. Levine, & M. Aparicio (Eds.), *Neural networks for knowledge representation and inference.* Mahwah, NJ: Erlbaum.

Bowdle, B. F., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology, 34,* 244–286.

Burstein, M. H. (1986). Concept formation by incremental analogical reasoning and debugging. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach.* Los Altos, CA: Morgan Kaufmann.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science, 15,* 89–132.

Clement, C. A., Mawby, R., & Giles, D. E. (1994). The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language, 33,* 396–420.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. In *Proceedings AAAI-86.* Philadelphia: AAAI.

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

Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure-mapping. In A. Ram & K. Eiselt (Eds.), *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Erlbaum.

Forbus, K. D., & Gentner, D. (1989). Structural evaluation of analogies: What counts? In *The Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Erlbaum.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19*(2), 141–205.

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Erlbaum.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition.* San Francisco: W. H. Freeman.

Gentner, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology, 13,* 56–83.

Gentner, D. (1982). Are scientific analogies metaphors? In D. S. Miall (Ed.), *Metaphor: Problems and perspectives.* London: Harvester.

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

Gentner, D. (1988). Metaphor as structure-mapping: The relational shift. *Child Development, 59,* 47–59.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). New York: Cambridge University Press.

Gentner, D., & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation.* San Francisco: Academic Press.

Gentner, D., & Imai, M. (1992). Is the future always ahead? Evidence for system-mappings in understanding space–time metaphors. *Proceedings of the Fourteenth Annual Meeting of the Cognitive Science Society* (pp. 510–515). Mahwah, NJ: Erlbaum.

Gentner, D., & Markman, A. B. (1993). Analogy: Watershed or Waterloo? Structural alignment and the development of connectionist models of analogy. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in neural information processing systems 5* (pp. 855–862). San Mateo, CA: Morgan Kaufmann.

Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science, 5*(3), 152–158.

Gentner, D., & Markman, A. B. (1995). Analogy-based reasoning. In M. Arbib (Ed.), *The handbook of brain theory and neural networks.* Cambridge, MA: MIT Press.

Gentner, D., & Markman, A. B. (1997). Structural alignment in analogy and similarity. *American Psychologist, 52*(1), 45–56.

Gentner, D., & Markman, A. B. (in preparation). Why syntactic graph matches are not analogies.

Gentner, D., & Medina, J. (1997). Comparison and the development of cognition and language. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society, 4*(1), 112–149.

Gentner, D., & Namy, L. L. (1999). Comparison in the development of categories. *Cognitive Development, 14,* 487–513.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development.* New York: Cambridge University Press.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology, 25*(4), 524–575.

Gentner, D., Rattermann, M. J., Markman, A. B., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modelling* (pp. 263–313). Mahwah, NJ: Erlbaum.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10,* 277–300.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12,* 306–355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*(1), 1–38.

Gleitman, L. R., Gleitman, H., Miller, C., & Ostrin, R. (1996). Similar, and similar concepts. *Cognition, 58,* 321–376.

Goldstone, R. L. (1994a). The role of similarity in categorization: Providing a groundwork. *Cognition, 52,* 125–157.

Goldstone, R. L. (1994b). Similarity, interactive-activation and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(1), 3–28.

Goldstone, R. L., & Medin, D. L. (1994a). Similarity, interactive-activation and mapping: An overview. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections.* Norwood, NJ: Ablex.

Goldstone, R. L., & Medin, D. L. (1994b). The time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(1), 29–50.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology, 23,* 222–262.

Greiner, R. (1988). Learning by understanding analogies. In A. Prieditis (Ed.), *Analogica.* Los Altos, CA: Morgan Kaufmann.

Gunn, V., & Gentner, D. (in preparation). Structural alignment facilitates the noticing of differences.

Halford, G. S., Wilson, W. H., Guo, J., Gayler, R. W., Wiles, J., & Stewart, J. E. M. (1995). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections* (pp. 363–415). Norwood, NJ: Ablex.

Hampton, J. A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language, 34,* 686–708.

Hofstadter, D. R., & Mitchell, M. (1994). An overview of the Copycat project.

In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol 2. Analogical connections* (pp. 31–112). Norwood, NJ: Ablex.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory.* New York: Academic Press.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15*(4), 332–340.

Holyoak, K. J., Novick, L. R., & Melz, E. R. (1994). Component processes in analogical transfer: Mapping, pattern completion, and adaptation. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol 2: Analogical connections* (pp. 113–180). Norwood, NJ: Ablex.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*(3), 295–355.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104*(3), 427–466.

James, W. (1985). *Psychology: The briefer course.* Notre Dame, IN: University of Notre Dame Press.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist, 39*(4), 341–350.

Keane, M. T. (1997). What makes an analogy difficult? The effects of order and causal structure on analogical mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(4), 946–967.

Keane, M. T., & Brayshaw, M. (1988). The incremental analogy machine: A computational model of analogy. In D. Sleeman (Ed.), *Third European Working Session on Machine Learning* (pp. 53–62). San Mateo, CA: Morgan Kaufmann.

Keane, M. T., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science, 18,* 387–438.

Keil, F. C. (1989). *Concepts, kinds and cognitive development.* Cambridge, MA: MIT Press.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67,* 2797–2822.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(3), 754–770.

Lindemann, P. G., & Markman, A. B. (1996). Alignability and attribute importance in choice. In G. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society.* Mahwah, NJ: Erlbaum.

Markman, A. B. (1996). Structural alignment in similarity and difference judgments. *Psychonomic Bulletin and Review, 3*(2), 227–230.

Markman, A. B. (1997a). Constraints on analogical inference. *Cognitive Science, 21*(4), 373–418.

Markman, A. B. (1997b). Structural alignment in similarity and its influence on category structure. *Cognitive Studies, 4*(4), 19–37.

Markman, A. B., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language, 32*(4), 517–535.

Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology, 25*(4), 431–467.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition, 24*(2), 235–249.

Markman, A. B., & Gentner, D. (1997). The effects of alignability on memory. *Psychological Science, 8*(5), 363–367.

Markman, A. B., & Medin, D. L. (1995). Similarity and alignment in choice. *Organizational Behavior and Human Decision Processes, 63*(2), 117–130.

Markman, A. B., & Wisniewski, E. J. (1997). Similar and different: The differentiation of basic level categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(1), 54–70.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1990). Similarity involving attributes and relations: Judgments of similarity and difference are not inverses. *Psychological Science, 1*, 64–69.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100*(2), 254–278.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification. *Psychological Review, 85*(3), 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*(3), 289–315.

Nosofsky, R. M. (1986). Attention, similarity and the identification–categorization relationship. *Journal of Experimental Psychology: General, 115*(1), 39–57.

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(1), 87–108.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology, 21*(2), 139–155.

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21–59). New York: Cambridge University Press.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General, 104*(3), 192–233.

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*(4), 629–639.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(3), 456–468.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*(1), 3–22.

Slovic, P., & MacPhillamy, D. (1974). Dimensional commensurability and cue utilization in comparative judgment. *Organizational Behavior and Human Performance, 11,* 172–194.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology, 31,* 307–346.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence, 46,* 259–310.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*(4), 327–352.

Wattenmaker, W. D., Dewey, G. I., Murphy, T. D., & Medin, D. L. (1986). Linear separability and concept learning: Context, relational properties and concept naturalness. *Cognitive Psychology, 18,* 158–194.

Wu, M., & Gentner, D. (1998). Structure in category-based induction. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1154–1158). Mahwah, NJ: Erlbaum.

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