

Informativity and Asymmetry in Comparisons

Brian F. Bowdle and Dedre Gentner

Northwestern University

We propose an account of comparison asymmetries based on viewing comparison as a process of structural alignment and mapping. Specifically, we hypothesize that (1) comparison asymmetries result from directional differences in informativity, and that (2) asymmetries can therefore be predicted from the relative degree of systematicity or conceptual coherence of the items being compared. In Experiment 1, we found a clear preference for placing the more systematic of two passages in the base position of a comparison. Experiments 1 and 2 further showed that structural alignability is crucial in obtaining such asymmetries. In Experiment 3, we found that asymmetries are predicted by the relative systematicity of the comparisons items rather than by the relative size of the distinctive feature sets. These results are inconsistent with accounts of asymmetry based on feature contrast or stimulus bias. In Experiments 4 through 6, we tested the functional implications of our account by examining inference projection and perceived informativity across asymmetric comparisons. Critically, comparisons having the more systematic item as the base were more likely to result in inference projection and other forms of target modification and were rated as more informative than reverse comparisons. We conclude by demonstrating that this account can explain comparison asymmetries without positing underlying asymmetries in subjective similarity, and that it offers a unified approach to the directionality of literal comparisons, analogies, and metaphors. © 1997 Academic Press

When comparing two items, people often prefer one direction of comparison over the other. For example, subjects prefer the comparison *North Korea is similar to Red China* over the reverse comparison *Red China is similar to North Korea* (Tversky, 1977; Tversky & Gati, 1978). As Tversky (1977) pointed out, such comparison asymmetries are inconsistent with traditional geometric models of similarity (e.g., Carroll & Wish, 1974; Shepard, 1974), which treat the similarity of two items as an inverse function of their distance in a multidimensional space. Because these models inherit the axioms of

This research was supported by NSF Grant SBR-95-11757 and ONR Grant N00014-92-J-1098, awarded to the second author. We thank Keith Holyoak, Art Markman, Doug Medin, Kelly Mix, Gregory Ward, Phillip Wolff, and the entire Similarity and Analogy group at Northwestern for their comments and suggestions. We also thank Mike Bell, Jon Berkowitz, Keith Burton, Ted Kunkel, Jason Thompson, and Steve Whiteside for their help with data collection and coding.

Address reprint requests to Brian F. Bowdle, Department of Psychology, Indiana University, 10th and Walnut Grove, Bloomington, IN 47405.

metric distance, similarity is assumed to be a symmetric relation between two items. The fact that comparisons often behave asymmetrically suggests that this may not be the case.

Nevertheless, the notion that similarity is a symmetric relation seems highly intuitive: if two items are similar, then they should remain similar regardless of the direction of comparison. Why, then, should people prefer one direction of comparison over another? Clearly, the answer to this question is important to any psychologically plausible model of comparison. We begin our examination of this issue by briefly reviewing three existing approaches to comparison asymmetries.

PREVIOUS ACCOUNTS OF COMPARISON ASYMMETRIES

Feature Contrast

Perhaps the most common psychological explanation of comparison asymmetry is that similarity is an asymmetric relation. Comparison asymmetries can therefore be seen as reflecting underlying asymmetries in the similarity of the comparison items. Thus, people prefer the comparison *North Korea is similar to China* over the reverse comparison precisely because North Korea is perceived as being more similar to China than the reverse. In support of this position, Tversky (1977; Tversky & Gati, 1978) found that the preferred direction of similarity comparison was commonly the one with the greater degree of judged similarity. Such asymmetries in similarity judgments have been obtained in areas as diverse as music perception (Bartlett & Dowling, 1988) and self-other similarity judgments (Catrambone, Beike, & Niedenthal, 1996; Holyoak & Gordon, 1983; Srull & Gaelick, 1983).

The position that similarity is an asymmetric relation was first made explicit by Tversky's (1977) *feature contrast model*. According to this model, the similarity of two items increases as a function of their common features and decreases as a function of their distinctive features. Asymmetries in similarity are explained by the *focusing hypothesis*: because the target of a directional comparison is the focus of attention, the distinctive features of the target are weighted more heavily than those of the base.¹ Thus, similarity will be maximized by placing the more complex or salient item in the base position of a comparison. For example, the comparison *North Korea is similar to Red China* is preferred over the comparison *Red China is similar to North Korea* because more distinctive information is included in the representation of Red China, making North Korea more similar to Red China than the reverse. In the feature contrast model, comparison asymmetries are predicted by the relative size and salience of the distinctive feature sets.

¹ The a-term of a comparison is often referred to as the *tenor*, *figure* or *topic*; and the b-term as the *vehicle*, *ground*, or *source*. We shall refer to the a-term as the *target* and the b-term as the *base*.

Stimulus Bias

Recent multidimensional scaling models of similarity have sought to account for the asymmetric behavior of comparisons without abandoning the intuitive assumption of symmetric similarity relations. In *stimulus-bias models*, asymmetric proximities can be derived from a symmetric similarity function operating between the comparison items, plus a differential bias associated with these items (e.g., Holman, 1979; Krumhansl, 1978; Nosofsky, 1991). Nosofsky (1991) reviews a number of such potential stimulus biases, including item density in the surrounding space, frequency of stimulus instantiation, prototypicality, feature loss or gain, and pattern goodness. Assuming that the target of a directional comparison is the focus of attention, and therefore is weighted more heavily than the base, the perceived distance between two items will be minimized by placing the item with the larger bias in the base position. Thus, the item with the larger bias will be the preferred base in similarity comparisons. For example, the comparison *North Korea is similar to Red China* is preferred over the comparison *Red China is similar to North Korea* because Red China is associated with a larger bias than is North Korea. In stimulus-bias models, comparison asymmetries are not predicted by any properties of the common or distinctive feature sets, but rather by biases existing independently of the particular comparison (Nosofsky, 1991).

Cognitive Reference Points

A third approach to the directionality of similarity comparisons uses the general notion of cognitive reference points to explain comparison asymmetries. According to *reference-point models*, asymmetries arise when (1) the comparison items can be seen as members of a common category, and (2) one of the items acts as a more natural reference point or landmark for members of that category (e.g., Gleitman, Gleitman, Miller, & Ostrin, 1996; Huttenlocher, Hedges, & Duncan, 1991; Rosch, 1975; Shen, 1989). Under these circumstances, the reference item will appear as the preferred base in similarity comparisons. Unlike either the feature contrast model or stimulus-bias models, reference-point models do not rely on underlying asymmetries in similarity or proximity to predict the preferred order of comparison statements; rather, they invoke general grammatical principles that require placing the deviant item in the subject (or figure) position of a sentence, and the reference item in the object (or ground) position (see Talmy, 1978, 1983). However, reference-point models can account for asymmetries in similarity judgments. As claimed by many of the above authors, deviant items are more easily assimilated to, and therefore are perceived as more similar to, reference items than vice versa.

THE PRAGMATICS OF ASYMMETRY

In contrast to the above accounts, we propose that comparison asymmetries are largely due to the combined influence of two general pragmatic principles:

Grice's (1975) *maxim of informativity*, and Clark and Haviland's (1977) *given-new contract*. According to the maxim of informativity, people expect any given utterance to be informative. This implies that if one direction of comparison is more informative than the other, then it should be preferred. But what makes one direction of comparison more informative? In directional comparisons, new information is projected from the base domain to the target domain, rather than from the target to the base or bidirectionally. This follows from the given-new contract, according to which given information precedes new information in an utterance. Comparison targets, which occupy the first (subject) position, represent *given* information. As such, they may involve concepts that are problematic due to vague, incomplete, or conflicting information. Comparison bases, which occupy the second (object) position, represent *new* information that can contribute to the understanding of the target. For example, the statement *Saddam Hussein is similar to Adolf Hitler* asserts that properties typically associated with Hitler (e.g., *genocidal tendencies*) may also be applicable to Hussein.

Taken together, the maxim of informativity and the given-new principle suggest that people should prefer the direction of comparison between two items that maximizes the amount of information projected from the base to the target. That is, the preferred direction of comparison is better able to suggest new commonalities. We refer to this as the *directional informativity hypothesis*. Unlike many existing models of similarity, the directional informativity hypothesis suggests a functional and dynamic approach to comparison asymmetries.

STRUCTURE-MAPPING THEORY

To the above pragmatic principles operating over comparison statements, we add a process model of comparison that accounts for the projection of novel information from base to target: Gentner's (1983, 1989) *structure-mapping theory*. As we will see, comparison asymmetry is a natural outcome of the constraints imposed by this model on inference projection. According to structure-mapping theory, drawing inferences from a comparison involves a prior process of structural alignment. The alignment process operates to create a maximal structurally consistent match between two representations that observes *one-to-one mapping* and *parallel connectivity* (Falkenhainer, Forbus, & Gentner, 1989; Gentner, Rattermann, & Forbus, 1993; Markman & Gentner, 1993a, 1993b). That is, each element of one representation can be placed in correspondence with at most one element of the other representation, and arguments of aligned predicates are themselves aligned. Mappings that form interconnected structures, in which higher-order relations constrain lower-order relations, are preferred to less structured mappings (see also Clement & Gentner, 1991; Spellman & Holyoak, 1992).

There is considerable evidence that similarity comparisons involve the structural alignment of representations (e.g., Gentner & Markman, 1994; Gentner, Rattermann, & Forbus, 1993; Goldstone, 1994; Goldstone & Medin,

1994; Goldstone, Medin, & Gentner, 1991; Markman & Gentner, 1993a, 1993b; Medin, Goldstone, & Gentner, 1993). For example, Markman and Gentner (1993a) created pairs of cross-mapped pictures (see Gentner & Toupin, 1986), such that perceptually similar objects played different roles in like relational structures. Markman and Gentner asked subjects to indicate which object in one picture was the best match to a target item in another picture, and found that subjects generally chose the most perceptually similar object. However, when subjects first rated the similarity of the two pictures, they were more likely to override local object similarity and map the objects on the basis of relational roles, suggesting that the process of performing similarity comparisons promotes structural alignment.

Once a structurally consistent alignment between the target and base domains has been found, further predicates from the base domain that are connected to the common system (but that are not initially present in the target domain) are mapped to the target as *candidate inferences*. Clement and Gentner (1991) demonstrated that such inference selection can be predicted on the basis of connectivity. Subjects were presented with pairs of analogous passages in which the base passage contained two facts that could readily map to the target, only one of which was connected to a matching causal event. When asked to make predictions about the target passages based on the analogies, subjects imported whichever fact was linked to the common connected system, neglecting the equally available isolated fact. Similarly, Gentner, Rattermann, and Forbus (1993) found that subjects rated the inferential soundness of analogical comparisons more highly when the comparisons exhibited structural commonalities. Thus, structural alignment provides a basis for the selection and evaluation of candidate inferences (see also Holyoak & Koh, 1987; Holyoak & Thagard, 1989; Keane, 1988; Read, 1984, 1987; Ross, 1987, 1989).

As implied by the above discussion, structure-mapping theory assumes structured representations. The importance of structured representations has long been recognized in cognitive studies of frames (Minsky, 1977), scripts (Schank & Abelson, 1977), and schemas (Rumelhart & Ortony, 1978). However, the types of conceptual structures assumed by structure-mapping theory go beyond simple relations such as feature correlations, semantic associations, and temporal order; rather, representations are treated as possessing some degree of *systematicity*, in which higher-order relations govern or constrain lower-order relations. Systematic representations are characteristic of the kinds of conceptual structures that occur in mental models (e.g., Gentner & Stevens, 1983; Johnson-Laird, 1983; Lakoff, 1987) and intuitive theories (e.g., Carey, 1985; Keil, 1989; Murphy & Medin, 1985). Further, because systematicity is a relatively formal notion, it is applicable to a broad range of relational types and forms, including those contributing to perceptual coherence (e.g., monotonicity, symmetry—Palmer, 1977, 1978) and text coherence (e.g., causal networks, goal hierarchies—Johnson & Mandler, 1980; Kintsch & van Dijk, 1978; Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985).

Structure-Mapping Theory and Asymmetry

Given the process model of comparison reviewed above, we can now lay out the properties of the target and base that lead to comparison asymmetries. According to the directional informativity hypothesis, asymmetries should favor the direction of comparison providing the greatest number of candidate inferences to the target. According to structure-mapping theory, the process of deriving target predications is essentially one of structural completion (Gentner & Markman, 1993; Markman & Gentner, 1993a). Thus, given the base-to-target direction of inference projections, the base representation should possess a greater degree of systematicity than the target representation. Only when the aligned structure is more systematic in the base will the comparison be able to lend additional structure to the target. Our central prediction is that comparison asymmetries result from a *systematicity imbalance* between the comparison items. When two items *A* and *B* are compared, people will prefer *A is similar to B* over *B is similar to A* to the extent that *B* is more systematic than *A* with respect to the common system. We will refer to this as the *base systematicity advantage*.

The directional informativity hypothesis makes three additional predictions about comparison asymmetries, none of which are easily captured by the feature contrast model or stimulus-bias models. First, because only those properties that are connected to the aligned structure in the base may be projected as candidate inferences, asymmetries should only be obtained over alignable representations. Neither direction of comparison will be informative if the representations are not alignable. This means that asymmetries cannot be due to conceptual properties that are independent of the particular frame of comparison. Second, because systematicity cannot be reduced to number or salience of instantiated features, differences in systematicity should be a better predictor of asymmetries than differences in the number or salience of distinctive features. Third, the preferred direction of comparison should (a) be judged as more informative than the non-preferred direction, and (b) should support more candidate inferences than the non-preferred order.

EXPERIMENT 1

In Experiment 1, we tested the central prediction of a base systematicity advantage in directional comparisons. We focused on a particular type of systematicity, namely, text coherence, which emphasizes causal and explanatory relations (see Johnson & Mandler, 1980; Kintsch & van Dijk, 1978; Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985). Subjects were presented with pairs of similar stories designed to vary in degree of relative systematicity, and asked to indicate which direction of comparison they preferred. According to the directional informativity hypothesis, subjects should prefer comparing the less coherent story to the more coherent story over the reverse, as the former direction maximizes comparison informativity. We shall refer to comparisons in this direction as *forward*

TABLE 1

Sample Story Set (Standard plus Three Variants) from Experiment 1 (The Causal Pivot in the Standard and the Substitute Statements in the Deletion and Replacement Variants Are Italicized)

STANDARD. Peter was the most faithful and devoted of all the younger monks in the monastery. He was as skilled at illuminating texts as he was at tending the gardens. His simple passion for God was admired by all. One day at dinner, he spilled his food onto the floor. *Peter was dismayed, as he viewed the wasting of food as a sin.* He went to his cubical and began to fast in silence. His elders became worried, but they decided to say nothing. When a week had passed, Peter felt that he had repented, and resumed his life as usual.

DELETION VARIANT. Luke was the most faithful and devoted of all the younger monks in the monastery. He was as skilled at chanting hymns as he was at copying texts. His simple passion for God was admired by all. *His elders were particularly impressed by him, and expected that Luke would become a prominent figure.* One day at dinner, he spilled his food onto the floor. He went to his cubical and began to fast in silence. His elders became worried, but they decided to say nothing. When a week had passed, Luke felt that he had repented, and resumed his life as usual.

REPLACEMENT VARIANT. Luke was the most faithful and devoted of all the younger monks in the monastery. He was as skilled at chanting hymns as he was at copying texts. His simple passion for God was admired by all. One day at dinner, he spilled his food onto the floor. *Luke was dismayed, as the bean porridge they were serving that night was his favorite.* He went to his cubical and began to fast in silence. His elders became worried, but they decided to say nothing. When a week had passed, Luke felt that he had repented, and resumed his life as usual.

SHUFFLED VARIANT. Luke was the most faithful and devoted of all the younger monks in the monastery. He was as skilled at chanting hymns as he was at copying texts. His simple passion for God was admired by all. Luke was dismayed. One day at dinner, he began to fast in silence, as he viewed the wasting of food as a sin. When a week had passed, Luke felt that he had repented, and spilled his food on the floor. He went to his cubical and resumed his life as usual. His elders became worried, but they decided to say nothing.

comparisons, and to comparisons in the opposite direction as *reverse comparisons*.

We designed sets of brief (one paragraph) stories, each set containing a causally coherent standard story and three less systematic variants. We used three different methods of creating this systematicity imbalance, as shown in Table 1. *Deletion* variants were created by eliminating a primary causal pivot in the standard story and adding a distinctive noncausal sentence elsewhere in the story. *Replacement* variants were created by replacing the causal pivot with a substitute pivot that was inconsistent with the story plot. Finally, *shuffled* variants were created by rearranging the order of propositions in the standard story (while maintaining surface cohesion). Aside from differences in names and surface descriptions, the characters, locations, and actions contained in the variants were similar to those of the standard stories. Thus, each of the variants was at least locally alignable with the standard.

Subjects received pairs consisting of a standard story and one of its variants—that is, deletion pairs, replacement pairs, and shuffled pairs—and were asked to say which direction of comparison they preferred. We predicted an asymmetry in all three variant conditions favoring the forward direction of

comparison. We also expected the degree of asymmetry across variant conditions to reflect the degree of systematicity imbalance, as described below.

Method

Subjects

Forty-eight Northwestern University undergraduates participated in partial fulfillment of a course requirement.

Materials and Design

Subjects received booklets containing 18 story pairs, one from each story set constructed. Variant condition was manipulated within subjects, so that each subject received six exemplars of each pair type: deletion, replacement, and shuffled (see Table 1). Thus, each subject read only one comparison from each story set, but read comparisons in all three variant conditions. Story pair assignment was counterbalanced across subjects. Each story pair was presented on a separate page, with the top story labeled "Story A" and the bottom one "Story B." The order of standard-variant presentation (i.e., which was Story A and which was Story B) was counterbalanced across subjects and story pairs. The order of story pair presentation was randomized. The design was 3 (variant condition: deletion, replacement, shuffled) \times 2 (order of presentation: standard first, variant first), both within subjects.

To check the degree of systematicity imbalance of the story pairs, we presented the materials to a separate group of 48 Northwestern University undergraduates in the manner described above. Subjects read the pairs of stories, and for each pair they rated the level of coherence of the individual stories on a scale of 1 (low coherence) to 7 (high coherence). Coherence was loosely defined as "the extent to which a story makes sense or 'hangs together.'" Separate coherence scales were provided for each story at the bottom of each page. We calculated the systematicity imbalance for each pair by subtracting the coherence ratings of the variant stories from those of the standard stories. As expected, the systematicity imbalance ($M = 2.96$) was significantly greater than zero, $t_1(17) = 21.45$, $p < .001$. This imbalance held within each variant condition: the difference for deletion ($M = 1.99$), replacement ($M = 2.70$), and shuffled ($M = 4.19$) were all significantly greater than zero, $t_1(17) = 8.76$, $p < .001$, $t_1(17) = 12.05$, $p < .001$, and $t_1(17) = 29.52$, $p < .001$, respectively. A one-way (variant condition: deletion, replacement, shuffled) repeated-measures analysis of variance (ANOVA) was performed on the systematicity imbalance scores. As expected, a main effect was obtained for variant condition, $F_1(2,34) = 39.00$, $p < .001$. Planned comparisons confirmed a lower degree of systematicity imbalance for that the deletion pairs showed a lower degree of systematicity imbalance for deletion pairs than for replacement pairs, $t_1(17) = 2.90$, $p < .01$, and for replacement pairs than for shuffled pairs, $t_1(17) = 6.24$, $p < .005$. Thus, the standard was rated as more systematic than the variant in each variant condition, with the smallest difference for the deletion pairs and the greatest difference for the shuffled pairs.

To check the alignability of the story pairs, we presented the materials to an additional group of 48 Northwestern University undergraduates in the manner described above. Subjects read the pairs of stories, and for each pair they rated the alignability of the stories on a scale of 1 (low alignability) to 7 (high alignability). Alignability was defined as "the extent to which certain aspects of one story correspond to, or 'line up with,' certain aspects of the other story." The mean alignability ratings were 5.29 for the deletion pairs, 5.09 for the replacement pairs, and 4.09 for the shuffled pairs. A one-way (variant condition: deletion, replacement, shuffled) repeated-measures ANOVA was performed on the alignability ratings. A main effect was obtained for variant condition, $F_1(2,34) = 35.00$, $p < .001$. Planned comparisons indicated that the shuffled pairs were less alignable than either the deletion pairs, $t_1(17) = 7.00$, $p < .001$, or the replacement pairs, $t_1(17) = 6.40$, $p < .001$. There was no difference between the deletion and replacement pairs.

Procedure

Subjects read the 18 pairs of stories, and for each story pair they indicated which of the following statements they preferred: "Story A is like Story B," or "Story B is like Story A."

Results and Discussion

As expected, there was an overall asymmetry favoring forward comparisons. The proportion of responses in this direction ($M = .76$) was significantly greater than chance, $t_s(47) = 10.52, p < .001$ and $t_l(17) = 16.44, p < .001$. This asymmetry held within each variant condition: the proportions of forward comparisons for deletion ($M = .68$), replacement ($M = .79$), and shuffled ($M = .81$) were all significantly greater than chance (all p 's $< .001$ across both subjects and items). Thus, the predicted base systematicity advantage was obtained in all three variant conditions.

A one-way (variant condition: deletion, replacement, shuffled) repeated-measures ANOVA was performed on the proportion of forward comparisons. As expected, a main effect of variant condition was obtained across both subjects and items, $F_s(2,94) = 7.14, p < .005$ and $F_l(2,34) = 9.06, p < .001$. Planned comparisons confirmed that the deletion pairs were less asymmetric than either the replacement pairs, $t_s(47) = 2.97, p < .005$ and $t_l(17) = 3.15, p < .01$, or the shuffled pairs, $t_s(47) = 2.96, p < .005$ and $t_l(17) = 3.77, p < .005$. However, there was no difference between the replacement and shuffled pairs.

The major prediction of the directional informativity hypothesis was borne out: subjects overwhelmingly preferred forward over reverse comparisons. However, our prediction that the degree of asymmetry would reflect the degree of systematicity imbalance was only partially borne out. The deletion pairs were less asymmetric than the other two variant conditions, but contrary to expectation, the replacement and shuffled pairs did not differ in asymmetry, despite the fact that the shuffled pairs showed the greatest degree of systematicity imbalance in the stimulus pretest. One potential explanation for this result is that the stories in the shuffled condition were rated as less alignable than those in the other two conditions. Indeed, the stories in the shuffled pairs are only locally alignable by design: the objects and individual events contained in the stories are similar, but the higher-order relational structure characterizing the standard story plots was virtually eliminated in the variants. In contrast, the stories in the deletion and replacement pairs are both locally and globally alignable: with the exception of the primary causal pivots, most of the higher-order structure of the standards is preserved in the variants.

Why should alignability matter in predicting comparison asymmetries? According to structure-mapping theory, structural alignment precedes and constrains inference generation. In particular, only information connected to the common system is likely to be projected from base to target as candidate inferences (Clement & Gentner, 1991), thereby contributing to the informativity of the comparison. Thus, a reduction in alignability should result in a reduction in asymmetry. In terms of the present results, this means that the

lower alignability of the shuffled pairs could have mitigated the predicted effects of systematicity imbalance.

Regression Analysis

In order to test this hypothesis, we performed a multiple regression analysis using the systematicity imbalance scores and alignability ratings obtained in the stimulus pretest as independent variables, and the proportions of forward comparisons obtained in Experiment 1 as the dependent variable. We also included the interaction between systematicity imbalance and alignability as an independent variable. Because directional informativity is predicted by the systematicity imbalance of aligned structures only, we expect that this interaction should be a significant predictor of the degree of asymmetry.

The analysis resulted in the following regression equation:

$$\text{Asymmetry} = -1.04(\text{Systematicity Imbalance}) \\ - 0.48(\text{Alignability}) + 1.32(\text{SI} \times \text{A})$$

This equation possessed an overall R^2 of .42, $F(3,50) = 12.28$, $p < .001$. Neither the coefficient for systematicity imbalance nor that for alignability was significantly greater than zero. However, the interaction between these two variables was a significant predictor of the degree of asymmetry, $t = 2.71$, $p < .01$. This result supports the prediction that only information connected to the common relational structure should influence comparison directionality, as only such information can be projected from the base to the target in a constrained manner, thereby making the comparison informative. In other words, there can be no asymmetry without alignability.

EXPERIMENT 2

The results of Experiment 1 bear out the claim that both systematicity imbalance and alignability are necessary for asymmetry. However, the story pairs were all highly alignable. A stronger test of the claim that directional informativity (and therefore asymmetry) depends on the alignment of conceptual structures would involve comparisons between alignable and non-alignable story pairs.

In Experiment 2, subjects were again presented with pairs of stories and asked to say which direction of comparison they preferred. In this experiment, only two of the three conditions—deletion and shuffled—were maintained. The critical manipulation was that for half of these comparisons, the standard and variant stories were drawn from the same story set and were therefore alignable, whereas for the other half, they were drawn from different story sets and were not alignable.

Our predictions are as follows: First, for the alignable comparisons, a base systematicity advantage should be obtained. For the non-alignable comparisons, however, there should be no significant base systematicity advantage, as these comparisons cannot convey new information about their targets.

Second, when the comparison is between alignable stories, the shuffled pairs should be more asymmetric than the deletion pairs, reflecting the relative degrees of systematicity imbalance of these two conditions. In contrast, there is no reason to expect that these two variant conditions should differ in degree of asymmetry for the non-alignable comparisons.

Method

Subjects

Thirty-two Northwestern University undergraduates participated in partial fulfillment of a course requirement. None had previously participated in a similar experiment.

Materials and Design

Twelve of the 18 original story sets were used, each containing a standard plus two variants (the deletion and shuffled pairs, as shown in Table 1). Subjects received booklets containing 12 story pairs, one from each story set. Variant condition was manipulated within subjects, so that each subject received six exemplars of each pair type. Thus, each subject read only one comparison from each story set, but read comparisons in both variant conditions. Story pair assignment was counterbalanced across subjects. For half of these story pairs, the standard and variant were from the same story set, while for the other half, they were drawn from different story sets. Each story pair was presented on a separate page, with the top story labeled "Story A" and the bottom one "Story B." Story pair alignability (i.e., standard and variant from the same or different story sets) and order of standard-variant presentation (i.e., which was Story A and which was Story B) were counterbalanced across subjects and story pairs. The order of story pair presentation was randomized. The design was 2 (variant condition: deletion, shuffled) \times 2 (alignability: alignable, nonalignable), both within subjects.

Procedure

The procedure was identical to that used in Experiment 1.

Results and Discussion

The results are summarized in Fig. 1. There was a significant asymmetry favoring forward comparisons for both alignable and non-alignable story pairs. Averaging across variant conditions, the proportion of forward comparisons was significantly greater than chance for the alignable pairs ($M = .78$), $t_s(31) = 8.90$, $p < .001$ and $t_t(11) = 10.22$, $p < .001$, and for the non-alignable pairs ($M = .58$), $t_s(31) = 2.44$, $p < .05$, but no across items. These results are contrary to our prediction that there should be no base systematicity advantage for non-alignable comparisons. However, the picture changes when the two variant conditions are analyzed separately. Both the deletion pairs ($M = .68$) and the shuffled pairs ($M = .88$) showed significant directional preferences when the stories were alignable (all p 's $< .05$ across both subjects and items). In contrast, when the stories were non-alignable, the deletion pairs ($M = .57$) showed no base systematicity advantage, and the shuffled pairs ($M = .59$) only showed a base systematicity advantage across items, $t_t(11) = 2.69$, $p < .05$. Thus, consistent comparison asymmetries were only obtained for alignable comparisons.

A 2 (variant condition: deletion, shuffled) \times 2 (alignability: alignable, non-




FIG. 1. Proportion of forward comparisons as a function of alignability and variant condition in Experiment 2.

alignable) repeated-measures ANOVA was performed on the proportion of forward comparisons. As expected, there was a main effect of alignability: alignable comparisons were more asymmetric than non-alignable comparisons across both subjects and items, $F_{S(1,31)} = 20.65, p < .001$ and $F_{I(1,11)} = 13.80, p < .005$. A main effect of variant condition was also obtained: averaging across alignability assignments, the directional preference was higher for shuffled pairs ($M = .71$) than for deletion pairs ($M = .63$) across subjects, $F_{S(1,31)} = 5.26, p < .05$, but not across items. Finally, there was an interaction between alignability and variant condition. As predicted, the difference in directional preference between the two variant conditions was greater for the alignable pairs ($M = .20$) than for the non-alignable pairs ($M = .02$) across subjects, $F_{S(1,31)} = 4.24, p < .05$, but not across items.

For alignable pairs, subjects showed a strong preference for forward comparisons. In contrast, for non-alignable pairs, subjects showed only a marginal preference for placing the more systematic item in the base position. This bears out the structure-mapping prediction that alignability is necessary for asymmetry. These results are problematic for the feature contrast model of similarity. Because this model derives directionality from the size and salience of the distinctive feature sets, the degree of comparison asymmetry should not depend on the number of common (alignable) features. The present results also pose a challenge to stimulus-bias models, which posit that the item associated with the larger bias will appear as the preferred base in similarity comparisons. Because such biases exist independently of any particular com-

parison, asymmetries resulting from differential stimulus bias should be independent of alignability. Of course, it might be objected that the non-alignable pairs in this experiment represent an extreme case: to the extent that two items possess no common features, or are maximally distant in a multidimensional space, the effects of distinctive feature sets or stimulus biases may be minimized. However, this objection does not apply to Experiment 1, in which all story pairs were sufficiently similar to permit meaningful comparison: here, too, the interaction of systematicity imbalance and alignability was found to be a significant predictor of asymmetry. In sum, structural alignment is a prerequisite for directional informativity, and therefore for comparison asymmetry.

EXPERIMENTS 3A AND 3B

The results of Experiments 1 and 2 support the structure-mapping account: to the extent that comparison items are alignable, the item possessing greater systematicity is the preferred base. However, it is possible that the standard stories in the above experiments, along with being more systematic than the variants, also possessed a greater number of distinctive features due to the increased number of higher-order relations. Leaving aside the issue of alignability, this would allow the feature contrast model of similarity to account for the observed base systematicity advantage. In Experiment 3, we addressed this issue. Subjects were again presented with pairs of stories and asked to say which direction of comparison they preferred. As in Experiment 2, only the deletion and shuffled conditions were used. The critical manipulation in this experiment was that for one-third of the comparisons, two distinctive statements were added to the standard story, and for another one-third of the comparisons, two such statements were added to the variant story. The remaining one-third of the comparisons were the same as in Experiment 1. These additional statements were designed so that they did not increase the systematicity of the stories they were added to. Specifically, in Experiment 3a, the additional statements were *structurally irrelevant* to the stories, in that they were unrelated to the story plots; and in Experiment 3b, the additional statements were *structurally inconsistent* with the stories, in that they interfered with the stories' causal flow (like the substitute pivots of the replacement variants in Experiment 1). Examples in which structurally irrelevant and structurally inconsistent statements were added to the standard are given in Table 2 and Table 3, respectively.

Structure-mapping theory and the feature contrast model make different predictions as to the effects of adding these structurally irrelevant or inconsistent statements. The feature contrast model predicts that adding either of these types of distinctive features to a story will increase the preference for placing that story in the base position of a comparison, where they will count less against the similarity of the pair. In contrast, structure-mapping theory predicts that adding structurally irrelevant statements to a story will have no effect on comparison asymmetries, as they do not affect the systematicity of the

TABLE 2
Example of Irrelevant Statement Additions to a Standard Story from Experiment 3a
(Irrelevant Statements Are Italicized)

Walter was a poor woodcutter who lived on the outskirts of Portshire. He had a wife and two children, but he was hardly able to support them. *Twice every month, Walter journeyed to the other side of the forest to visit his parents.* Now, the King of Portshire cared very much about his subjects, and would often travel in disguise to see how they were faring. One winter, a beggar came knocking at Walter's door, asking for lodging. Despite his wife's protests, Walter took him in. He gave the beggar as much food as he could spare, and made him a bed by the fire. *Because the fire was so warm, the beggar slept without any blanket.* In the morning, Walter found the beggar gone, and a sack of gold on the mantle.

aligned structure in that story. Adding structurally inconsistent statements to a story should actually *decrease* the preference for placing that story in the base position of a comparison, as they lower the systematicity of the aligned structure in that story. Thus, our predictions are as follows: In Experiment 3a, adding structurally irrelevant statements to either the standard or the variant should not significantly influence comparison asymmetries. In Experiment 3b, adding structurally inconsistent statements to the standard should lower the preference for forward comparisons, whereas adding structurally inconsistent statements to the variant should increase this preference, relative to the original pairs. In addition, the overall asymmetry should be higher for the shuffled than the deletion pairs in both Experiments 3a and 3b, replicating the results of Experiments 1 and 2.

Method

Subjects

Ninety-six Northwestern University undergraduates participated in partial fulfillment of a course requirement: 48 in Experiment 3a and 48 in Experiment 3b. None had previously participated in a similar experiment.

Materials and Design

Experiment 3a. Twelve of the 18 original story sets were used, each containing a standard plus two variants (the deletion and shuffled pairs, as shown in Table 1). Subjects received

TABLE 3
Example of Inconsistent Statement Additions to a Standard Story from Experiment 3b
(Inconsistent Statements Are Italicized)

Before long, the pirate ship had overtaken the Duke's vessel. After firing a shot across her bow, the ruffian sailors swung across from their deck to the royal clipper. *The pirates fought valiantly, but they were badly outnumbered by the Duke's men.* When all resistance had been subdued, Jack, the pirate captain, ran down to the hold to seize the finest gold and jewels for himself. Hiding there amidst the chests of treasure was the Duchess. *Upon seeing her, Jack could hardly contain his loathing.* She was the most beautiful woman that he had ever seen. He promptly forgot all about the treasure and took the Duchess as his bounty instead.




FIG. 2. Proportion of forward comparisons as a function of irrelevant statement addition and variant condition in Experiment 3a.

booklets containing 12 story pairs, one from each story set. Variant condition was manipulated within subjects, so that each subject received six exemplars of each pair type. Thus, each subject read only one comparison for each story set, but read comparisons in both variant conditions. For one-third of these story pairs, two structurally irrelevant statements were added to the standard story. For another one-third, two structurally irrelevant statements were added to the variant story (as illustrated in Table 2). For the final one-third, no statements were added to either story. Within each story set, the statements added to the standard and variant were the same. Each story pair was presented on a separate page, with the top story labeled "Story A" and the bottom one "Story B." Statement addition (i.e., standard, neither, or variant) and order of standard-variant assignment (i.e., which was Story A and which was Story B) were counterbalanced across subjects and story pairs. The order of story pair presentation was randomized. The design was 3 (statement addition: standard, neither, variant) $\times 2$ (variant condition: deletion, shuffled), both within subjects.

Experiment 3b. The materials and design were identical to Experiment 3a, with the following exceptions: For one-third of the story pairs, two structurally inconsistent statements were added to the standard story. For another one-third, two structurally inconsistent statements were added to the variant story (as illustrated in Table 3). For the final one-third, no statements were added to either story.

Procedure

The procedure was identical to that used in Experiment 1.

Results and Discussion

Experiment 3a: Irrelevant Statement Additions

The results are summarized in Fig. 2. A 3 (statement assignment: standard, neither, variant) $\times 2$ (variant condition: deletion, shuffled) repeated-measures




FIG. 3. Proportion of forward comparisons as a function of inconsistent statement addition and variant condition in Experiment 3b.

ANOVA was performed on the proportion forward comparisons. As expected, there was no effect of statement assignment on comparison asymmetries. As in Experiments 1 and 2, however, there was a main effect of variant condition: averaging across statement assignments, the directional preference was higher for shuffled pairs ($M = .80$) than for deletion pairs ($M = .67$) across both subjects and items, $F_s(1,47) = 14.03$, $p < .001$ and $F_t(1,11) = 11.83$, $p < .01$. There was no interaction between statement assignment and variant condition.

Experiment 3b: Inconsistent Statement Additions

The results are summarized in Fig. 3. A 3 (statement assignment: standard, neither, variant) \times 2 (variant condition: deletion, shuffled) repeated-measures analysis of variance was performed on the proportion forward comparisons. As expected, there was a main effect of statement addition across both subjects and items, $F_s(2,94) = 6.79$, $p < .005$ and $F_t(2,22) = 10.52$, $p < .001$. Planned comparisons indicated a lower preference for forward comparisons when the structurally inconsistent statements were added to the standard ($M = .54$) than when the story pairs were left unmodified ($M = .71$), $t_s(47) = 3.37$, $p < .005$ and $t_t(11) = 4.08$, $p < .005$, or when structurally inconsistent statements were added to the target ($M = .73$), $t_s(47) = 2.87$, $p < .01$ and $t_t(11) = 4.33$, $p < .005$. However, there was no significant difference in asymmetry between the latter two conditions. The predicted main effect of variant condition was also obtained: averaging across statement assignments, the directional preference was higher for shuffled pairs ($M = .74$) than for deletion pairs ($M = .58$) across both subjects and items, $F_s(1,47) = 17.94$, $p < .001$ and $F_t(1,11) = 20.96$, $p < .001$.

Finally, there was an interaction between statement addition and variant condition. The effects of statement addition were larger for the deletion pairs than for the shuffled pairs across subjects, $F_s(2,94) = 3.35$, $p < .05$, but not across items.

The results bear out the predictions of structure-mapping theory and the directional informativity hypothesis. In Experiment 3a, the addition of structurally irrelevant statements to either story of the story pairs had no effect on comparison asymmetries, and in Experiment 3b, the addition of structurally inconsistent statements to the standard story reduced the suitability of that story as a comparison base. Indeed, in the deletion condition, adding structurally inconsistent statements to the standard story actually reversed the directional preference ($M = .40$). These results are contrary to the predictions of the feature contrast model. Because both inconsistent and irrelevant statements were presumably highly salient additions to the stories, the influence of systematicity on asymmetry cannot be explained in terms of sheer number of distinctive features. Rather, the effect of distinctive features depends on their relation to the aligned system.

FROM SYSTEMATICITY IMBALANCE TO INFORMATIVITY: A COMPUTATIONAL SIMULATION

The results so far show that (a) subjects prefer placing the more systematic item in the base position of a comparison, (b) systematicity is a better predictor of asymmetry than sheer number of conceptual features, and (c) alignability is a prerequisite for asymmetry. These results specify the properties of comparison items that support directional informativity, and favor the structure-mapping account of asymmetry over feature contrast or stimulus-bias accounts. However, we have not directly examined the hypothesized relationship between asymmetry and the process of projecting novel information from the base to the target.

According to the directional informativity hypothesis, asymmetries in similarity reflect a base systematicity advantage. Speakers construct their comparisons to have the less systematic item as given information, and the more systematic item as new information, as this ordering maximizes the perceived informativity of the comparison. This hypothesis depends on the structure-mapping claim that subjects derive inferences from more systematic to less systematic items. To make this claim specific, we used the *Structure-mapping Engine* (SME), a computational model of the structure-mapping process outlined above (for details, see Falkenhainer, Forbus, & Gentner, 1989; Forbus, Gentner, & Law, 1995; Gentner, Falkenhainer, & Skorstad, 1987; Gentner, Rattermann, & Forbus, 1993). Given propositional representations of the target and base of a comparison, SME begins by placing all pairs of identical predicates in correspondence without regard to structural consistency (freely violating one-to-one correspondence and parallel connectivity).² Next, SME

² The representations given to SME consist of *entities* (logical individuals) and three types of predicates: *attributes* (one-argument predicates describing properties of entities), *functions* (typi-

collects these local matches into structurally consistent clusters (kernels). These kernels are combined into the two or three largest possible mutually consistent systems, which constitute the interpretations of the comparison. Each interpretation is then given a structural evaluation, using a type of cascade algorithm that favors deep interconnected relational structures over shallow or fragmentary structures. The interpretation with the highest structural evaluation is preferred. Given such an interpretation, candidate inferences can be projected from base to target. These are further predicates from the base domain that are connected to the common relational system but that are not initially present in the target domain. It is precisely this process of structural completion that motivates our prediction of a base systematicity advantage. By placing the more systematic of two alignable representations in the base position, the speaker signals the hearer as to the more richly inferential direction of mapping. As discussed earlier, this signalling follows from the given-new contract (Clark & Havilland, 1977).

Inference projection is not the only way in which comparisons can be informative. Structure-mapping theory allows comparisons to modify the target representation in (at least) the following additional ways: *structural highlighting*, *restructuring*, and *rerepresentation* (Gentner, Brem, Ferguson, Wolff, Markman, & Forbus, in press; Gentner & Wolff, in press). *Structural highlighting* is akin to schema abstraction (Gick & Holyoak, 1983), and involves increasing the salience of the common relational system in the target (or in both terms). Structural highlighting can occur when either (1) the common relational system is less prominent in the target; or (2) multiple possible structures can be entertained for the target and/or the base, only one of which is highly alignable. *Restructuring* involves changing the relational bindings of common elements in the target to match their bindings in the base, and can occur when these common elements are less mutually constrained in the target than in the base. Finally, in *rerepresentation*, nonidentical predicates in analogous relational roles are recast to permit discovering partial identities, thereby creating more coherent structural alignments (e.g., Clement, Mawby, & Giles, 1994; Gentner, 1989; Gentner & Rattermann, 1991; Gentner, Rattermann, Markman, & Kotovsky, 1995; Kotovsky & Gentner, 1996). This may involve representing the predicates as instances of a more general concept in an abstraction hierarchy (Falkenhainer, 1990), or decomposing the predicates until identical subpredicates are located (Burstein, 1983). (For related proposals, see Karmiloff-Smith, 1991; Kass, 1994; Novick & Holyoak, 1991.)

Currently, SME is extremely fluent in both inference projection and structural highlighting, has some proficiency in rerepresentation, and does virtually

cally representing dimensional information), and *relations* (predicates representing events, comparisons, and states applying to two or more arguments). First-order relations apply between entities, and higher-order relations apply between predicates or other lower-order relations.

no restructuring. Thus, we used SME to check the computational sufficiency of our claims concerning asymmetry and inference projection. We encoded the deletion pair from Table 1 (the "Peter the Monk" story set) into propositional form, as shown in Fig. 4, and submitted it to SME for both forward and reverse comparisons. In creating these narrative representations, we assumed that subjects formed a discourse model for each of the propositional macrostructures of the original texts (see Kintch & van Dijk, 1978).³ The representations reflect the fact that the standard (Fig. 4a) was designed to be more relationally systematic than deletion variant (Fig. 4b). In particular, the deletion variant contains no causal link between the spilling event and the fasting event.

Different interpretations were produced by SME depending on the direction of the comparison. In the forward direction (i.e., variant-to-standard), SME produced the interpretation shown in Fig. 5, as well as the following system of candidate inferences:



```
(REASON
  (CAUSE
    (AND (SPILL(luke, food, floor), BELIEVE(luke, SIN(SPILL)))),
    (DISMAY(luke)))
  (REPENT(luke, transgression)))
```

These inferences assert that the reason that Luke was dismayed was that he believed that wasting food—in this case, by spilling his food on the floor—is a sin. Further, it is inferred that this belief led Luke to repent, as it did with Peter. Thus, when the more coherent standard story is placed in the base position, the comparison is highly informative.


When the direction of comparison was reversed (i.e., standard-to-variant), SME obtained the same core interpretation (see Fig. 5), reflecting the fact that this common system is the maximal structurally consistent alignment. However, no candidate inferences resulted from this comparison. In this direction, the distinctive information in the base (the Luke story) is not connected to the common structure, so it cannot be projected to the target. Thus, the reverse direction of comparison is not informative. In sum, structural alignment in similarity comparisons is a symmetric process, but inference projection is directional. Further, inference projection depends not only on structural alignment, but also on the direction of comparison.

This simulation highlights an advantage of the present account over alternative accounts of comparison asymmetry: it provides an explicit mechanism for comparisons to convey new inferences about their targets.

³ We additionally assumed that subjects will include certain background knowledge concerning causal and explanatory relations in their representations, to the extent that it is invoked by the text. Thus, both the standard and variant representations in these simulations included propositions stating that fasting is done to repent for some perceived moral transgression.



FAITHFUL(peter)
 DEVOTED(peter)
 SKILLED AT(peter, illuminating texts)
 SKILLED AT(peter, tending gardens)
 ADMIRER BY(POSSESSION(peter, passion for god), all)



SPILL(luke, food, floor)
 FAITHFUL(luke)
 DEVOTED(luke)
 SKILLED AT(luke, chanting hymns)
 SKILLED AT(luke, copying texts)
 ADMIRER BY(POSSESSION(luke, passion for god), all)
 EXPECT(elders, BECOME(PROMINENT(luke)))

FIG. 4. (a) SME representation of the "Peter the Monk" standard story. (b) SME representation of the "Luke the Monk" deletion variant.

In the feature contrast model, for example, the most obvious candidate for inferences in literal comparisons is the set of distinctive features in the base. However, no mechanism is provided for selecting any particular predication from this potentially large set of non-matching information. Stimulus-bias models also cannot explain comparison informativity in a constrained man-

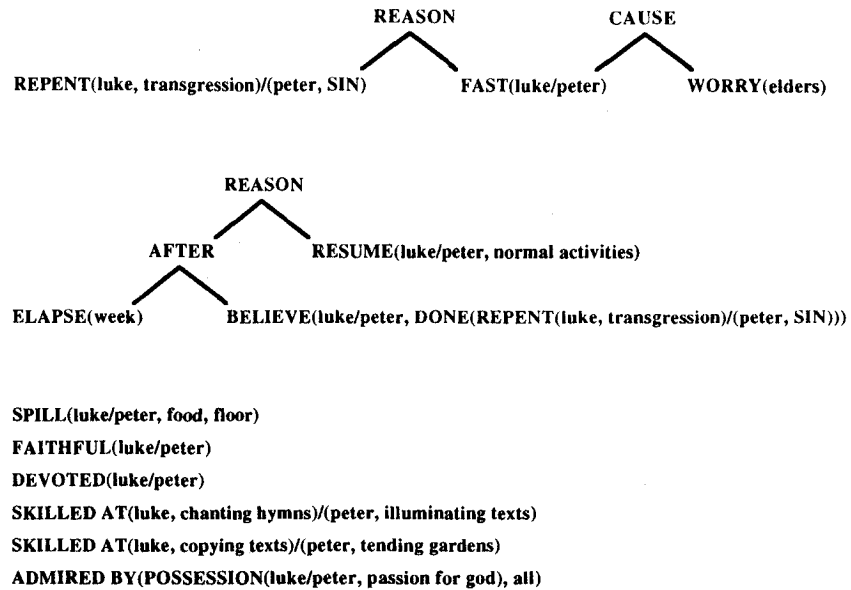


FIG. 5. SME interpretation of the comparison between the standard story and the deletion variant. Corresponding entities and predicates are separated by slashes.

ner. Indeed, to the extent that the preferred base is the item with the larger stimulus bias, the most obvious way to model informativity would be to increase the bias associated with the target. But many presumed sources of stimulus bias, such as prototypicality, frequency of instantiation, and item density in the surrounding space, do not seem readily amenable to change in this manner.

Of the approaches to asymmetry previously reviewed, the reference-point models seem closest in spirit to the present account. According to these models, once a common category for the comparison items has been established, the degree to which one item is preferentially seen in relation to or assimilated to the other item dictates the preferred order of comparison, with the deviant item as the target, and the reference item as the base. However, existing reference-point models are limited in several respects. For one thing, they do not explain how a common category is established if one does not already exist. For another, the mechanism by which deviant items may be assimilated to reference items is usually left unspecified (but see Huttenlocher, Hedges, & Duncan, 1991), and none of these models explains how new information may be transferred from the base to the target. Our account addresses these limitations as follows. First, structural alignment can give rise to the induction of relational schemas (Gick & Holyoak, 1983; see also Turner, 1988), which in turn can provide a common category for the comparison items. Second, assimilation of deviant items to reference items can be accomplished by inference projection, as well as structural highlighting, restructuring, or rerepresentation. On our view, highly systematic or coherent items make good reference points.

EXPERIMENT 4

In Experiment 4, we tested these computational predictions concerning inference projection by asking subjects to produce inferences for the story pairs used in Experiment 1. We predicted (a) that more inferences would be generated from standard to variant stories than in the reverse direction, and (b) that in this forward order, the inferences would consist of information initially present only in the base but connected to the common relational structure of the two stories. A second question of interest concerned the fine structure of these inference asymmetries. It seemed likely that inference asymmetry would be greater for the deletion and replacement conditions than for the shuffled condition for two reasons. First, the shuffled pairs were rated as least alignable in Experiment 1; and second, the shuffled variants lack a single aligned gap or anomalous point analogous to those found in the deletion and replacement variants. Either of these factors could act to increase the difficulty of drawing inferences in the shuffled condition.

We were also interested in other kinds of responses that might be generated for the story pairs. In pilot work, we found that subjects occasionally generated "inferences" that were either abstractions over the story pairs, or differences between the stories. Such responses—particularly those listing differences between stories—can further illuminate the process of structural alignment in comparison. For example, Markman and Gentner (1993b; Gentner & Markman, 1994) distinguish between differences connected to the common system, or *alignable differences*, and differences not connected to the common system, or *non-alignable differences*. Markman and Gentner (1993b) found that subjects listed more alignable differences for pairs of similar items (e.g., Brooms clean dry things whereas mops clean wet things), and more non-alignable differences for pairs of dissimilar items (e.g., Phone books are arranged alphabetically whereas lampshades are not). In a separate study, subjects found it easier to list differences for similar pairs than for dissimilar pairs, suggesting that alignable differences are more salient than non-alignable differences (Gentner & Markman, 1994). In other words, structural alignment not only determines how two things are similar, but also how they are different, and we therefore expected that any differences generated by subjects in this experiment would be predominantly alignable.

Method

Subjects

Forty-eight Northwestern University undergraduates participated in partial fulfillment of a course requirement. None had previously participated in a similar experiment.

Materials and Design

The materials and design were identical to those used in Experiment 1.

Procedure

Subjects read 18 pairs of stories, and for each story pair they made an inference about one of the stories based on the content of the other. They were told that the inference could be "anything

at all that you decide might be true of the story." Subjects were asked to (1) indicate which story they were making their inference about and (2) write out their inference in the space provided at the bottom of the page.

Results and Discussion

Inference Asymmetry

As predicted, there was an overall asymmetry favoring forward (standard-to-variant) inferences. The proportion of responses in this direction ($M = .80$) was significantly greater than chance, $t_s(47) = 12.08$, $p < .001$ and $t_l(17) = 20.20$, $p < .001$. This asymmetry held within each variant condition: the proportions of forward inferences for deletion ($M = .85$), replacement ($M = .83$) and shuffled ($M = .73$) were all significantly greater than chance (all p 's $< .001$ across both subjects and items). Thus, subjects preferred to use the more systematic story as the comparison base when generating inferences.

A one-way (variant condition: deletion, replacement, shuffled) repeated-measures ANOVA on the proportion of forward inferences indicated a main effect of variant condition across both subjects and items, $F_s(2,94) = 7.98$, $p < .001$ and $F_l(2,34) = 13.30$, $p < .001$. Planned comparisons confirmed that the shuffled pairs showed a lower degree of inference asymmetry than either the deletion pairs, $t_s(47) = 3.79$, $p < .001$ and $t_l(17) = 4.41$, $p < .001$, or the replacement pairs, $t_s(47) = 2.59$, $p < .025$ and $t_l(17) = 3.85$, $p < .005$. As expected, subjects were less consistent in drawing forward inferences in the shuffled condition.

Inference Coding

In addition to examining the degree of inferential asymmetry across variant conditions, we analyzed the kinds of inferences generated by subjects. The responses were coded by two blind judges into six inference categories: (a) *projected inferences* derived from the content of the other story; (b) *non-projected inferences* not based on the other story; (c) *alignable differences* (differences connected to common entities or relations); (d) *non-alignable differences* (differences not connected to common entities or relations); (e) *abstractions* of information contained in both stories; and (f) *other* responses. The coders could indicate more than one inference category for a response. Inter-coder reliability was 90%.⁴

The results of the inference coding are summarized in Table 4. As predicted, the predominant category for forward (standard-to-variant) responses across all three variant conditions was *projected inferences* (although *non-projected inferences* were almost as frequent for shuffled pairs). In contrast, the predominant category for reverse (variant-to-standard) responses across all three conditions was *non-projected inferences*, demonstrating the difficulty of finding

⁴ Reliability ratings were generated by taking the proportion of initial codings for each response that matched the final coding, and averaging these proportions over all responses.

TABLE 4
Proportion of Responses in Each Inference Category as a Function of Variant Condition
and Inference Direction in Experiment 4

Inference category	Deletion		Replacement		Shuffled	
	Forward	Reverse	Forward	Reverse	Forward	Reverse
Projected inference	.81	.24	.71	.28	.42	.12
Non-projected inference	.13	.55	.23	.48	.39	.65
Alignable difference	.04	.12	.06	.18	.05	.18
Non-alignable difference	—	—	—	—	—	—
Abstraction	.02	.10	.04	.08	.01	.04
Other	.02	.02	.03	.08	.20	.06

an acceptable inferential mapping from base to target in this direction. Additionally, both *alignable differences* and *abstractions* were more prevalent for reverse than forward responses across all three variant conditions. It appears that when subjects made reverse comparisons, they were often unable to locate an obvious inference to project to the target. In this case, they (1) interpolated from the target itself (*non-projected inferences*), (2) noted that two corresponding aspects of the stories were different (*alignable differences*), or (3) listed an explicit or implicit commonality between the two stories (*abstractions*). When subjects made forward comparisons, however, they were readily able to project new inferences from the base to the target. This inference asymmetry is consistent with SME's performance on forward and reverse comparisons given above.

None of the responses given by subjects were *non-alignable differences*, regardless of variant condition or direction. This finding is particularly informative for the deletion pairs, which contained an explicit non-alignable difference by design: the variant contained a distinctive non-causal statement not found in the standard. Thus, the finding that no subjects listed *non-alignable differences* for the deletion pairs supports the structure-mapping prediction that differences not connected to the common relational structure of comparison items should not be particularly salient (Gentner & Markman, 1994; Markman & Gentner, 1993b).

In this experiment, the degree of inference asymmetry favoring the more systematic item as the base was highest for the deletion and replacement pairs, and lowest for the shuffled pairs. Further, forward responses were less likely to involve projected inferences (and more likely to involve non-projected inferences) in the shuffled condition than in either the deletion or replacement conditions. Given that (1) the shuffled variants were less globally alignable with the standard stories than either the deletion or replacement variants, and that (2) the shuffled variants lacked a single aligned gap or anomalous point like those found in the deletion and replacement variants, these findings seem reasonable. However, this pattern of inference asymmet-

ries contrasts with the magnitude ordering of the comparison asymmetries in Experiment 1. There, the deletion pairs showed the lowest asymmetry, and the replacement and shuffled pairs the highest asymmetry. This leaves us with a puzzle: if the preferred direction of comparison maximizes the degree of perceived informativity, then why do these two measures not show the same ordering? In particular, why are the shuffled pairs strongly asymmetric in terms of preferred direction of comparison, but not in terms of inference projection? To attack this problem, we must look beyond inference projection to other means of making comparisons informative.

EXPERIMENT 5

Inference projection is a common mode of conceptual change in theories of literal similarity, analogy, and metaphor (e.g., Glucksberg & Keysar, 1990; Keane, 1988; Ortony, 1979). Nevertheless, there are other ways in which a comparison can change its target (Gentner, Brem, Ferguson, Wolff, Markman, & Forbus, 1997; Gentner and Wolff, in press). In our discussion of SME, we noted three such means: structural highlighting, restructuring, and rerepresentation. Although the shuffled story pairs in Experiment 4 may have been resistant to predicate carryover, restructuring is an obvious candidate for interpreting variant-to-standard comparisons in this condition. Indeed, further inspection of the forward responses in Experiment 4 revealed that the projected inferences generated for the shuffled pairs tended to involve restoring causal connections between existing predicates rather than importing novel predicates. At the same time, instructing the subjects in Experiment 4 to draw a single inference about one of the stories in each pair could have discouraged them from generating restructurings, and instead biased them to search for a single projectable predicate in the other story. This would explain why the shuffled condition showed the lowest degree of inference asymmetry in Experiment 4, despite having the highest degree of comparison asymmetry in Experiment 1. In Experiment 5, therefore, we broadened the range of allowable mappings by asked subjects to "change or modify" one story based on the content of the other. The term "inference" was removed from the instructions altogether.

According to the directional informativity hypothesis, (a) more modifications should be generated from standard to variant stories than in the reverse direction, and (b) in this forward order, the modifications should consist of information initially present only in the base but associated with the common relational structure of the two stories. This pattern of results would replicate Experiment 4. We further expected that, unlike in the previous experiment, the degree of this modification asymmetry across conditions would mirror the pattern of comparison asymmetries obtained in Experiment 1. That is, the preference for forward modifications should be lowest for the deletion pairs, and highest for the replacement and shuffled pairs.

Method

Subjects

Forty-eight Northwestern University undergraduates participated in partial fulfillment of a course requirement. None had previously participated in a similar experiment.

TABLE 5
Proportion of Responses in Each Modification Category as a Function of Variant Condition and Modification Direction in Experiment 5

Modification category	Deletion		Replacement		Shuffled	
	Forward	Reverse	Forward	Reverse	Forward	Reverse
Projected inference	.86	.47	.82	.45	—	.24
Non-projected inference	.11	.31	.12	.28	.08	.32
Projected subtraction	.20	.37	.69	.48	—	.16
Non-projected subtraction	.04	.12	.04	.03	.13	.24
Projected restructuring	—	—	—	—	.84	.12
Non-projected restructuring	.01	—	—	—	.04	.12
Other	.02	.12	.02	.24	.05	.24

Materials and Design

The materials and design were identical to those used in Experiments 1 and 4.

Procedure

Subjects read 18 pairs of stories, and for each story pair they made one or more changes or modifications to one of the stories based on the content of the other. Subjects were asked to (1) indicate which story they were modifying and (2) write out their modifications in the space provided at the bottom of the page.

Results and Discussion

Modification Asymmetry

As predicted, there was an overall asymmetry favoring forward (standard-to-variant) modifications. The proportion of responses in this direction ($M = .87$) was significantly greater than chance, $t_s(47) = 15.40$, $p < .001$ and $t_l(17) = 18.45$, $p < .001$. This asymmetry held within each variant condition: the proportions of forward modifications for deletion ($M = .80$), replacement ($M = .90$), and shuffled ($M = .91$) were all significantly greater than chance (all p 's $< .001$ across both subjects and items). Thus, subjects preferred to use the more systematic story as the comparison base when generating modifications.

A one-way (variant condition: deletion, replacement, shuffled) repeated-measures ANOVA on the proportion of forward modifications indicated a main effect of variant condition across both subjects and items, $F_s(2,94) = 10.30$, $p < .001$ and $F_l(2,34) = 6.82$, $p < .005$. Planned comparisons confirmed that the deletion pairs showed a lower degree of modification asymmetry than either the replacement pairs, $t_s(47) = 3.62$, $p < .001$ and $t_l(17) = 2.58$, $p < .025$, or the shuffled pairs, $t_s(47) = 3.55$, $p < .001$ and $t_l(17) = 3.24$, $p < .005$. However, there was no difference between the replacement and shuffled pairs. This pattern of modification asymmetries mirrors precisely the pattern of comparison asymmetries obtained in Experiment 1. Indeed,

there was a significant correlation between the modification asymmetries from this experiment and the comparison asymmetries from Experiment 1, $r = .61$, $p < .001$. This is strong evidence for the directional informativity hypothesis.

Modification Coding

A preliminary examination of the subjects' responses indicated a set of response categories that was radically different from that obtained in Experiment 4. First, in addition to making inferences, subjects often suggested deleting material from a story, or completely reorganizing the structure of a story. Second, subjects almost never generated abstractions over or differences between the stories in a given pair. Based on these observations, the subjects' responses were coded by two blind judges into seven modification categories: (a) *projected inferences* (material added to one story that was present in the other story); (b) *non-projected inferences* (material added to one story that was not present in the other story); (c) *projected subtractions* (material deleted from one story that was not present in the other story); (d) *non-projected subtractions* (material deleted from one story that was present in the other story); (e) *projected restructurings* (reorganizations of material in one story based on the structure of other other story); (f) *non-projected restructurings* (reorganizations of material in one story not based on the structure of the other story); and (g) *other responses*.⁵ The coders could indicate more than modification category for a response. Inter-coder reliability was 81%.

As shown in Table 5, *projected* modifications showed an advantage over *non-projected* modifications for forward (standard-to-variant) responses, and the degree of this advantage was comparable across all three variant conditions. The advantage for projected modifications was noticeably weaker for reverse (variant-to-standard) responses in the deletion and replacement conditions, and was altogether absent for reverse modifications in the shuffled condition. Further, non-projected modifications were far more typical of reverse responses than forward responses across all three conditions, demonstrating the relative difficulty of finding an acceptable mapping from base to target in this direction. These results replicate the overall pattern of inferences found in Experiment 4. There, projected inferences was the predominant category for forward responses, whereas non-projected inferences was the predominant category for reverse responses.

Looking at the specific types of forward modifications generated by subjects, we found that *projected inferences* were prevalent for both the deletion and replacement conditions. However, *projected subtractions* were only prevalent for the replacement condition. This is likely due to the fact that, in order

⁵ Under this coding scheme, responses that involved introducing predicates to a story were coded as inferences, whereas responses that involved changing causal connections between existing predicates were coded as restructurings. In Experiment 4, both types of responses were coded as inferences.

to fully integrate the causal pivot of a standard story into a replacement variant, the inconsistent causal statement in the variant story would have to be eliminated. Finally, turning to one of our key predictions, *projected restructurings* was the predominant category for forward responses in the shuffled condition (in sharp contrast to the deletion and replacement conditions). Thus, by allowing subjects more flexibility in the types of target projections that could be made, Experiment 5 fills the gap between Experiments 1 and 4. These results suggest that inference projection may be common in well-formed comparisons, but other forms of target modification (such as restructuring) can also contribute to the informativity of comparisons.

EXPERIMENT 6

In Experiment 5, we found that subjects prefer target modifications that are projected from more systematic to less systematic items, and that the degree of this preference corresponds to the degree of comparison asymmetry exhibited by the items in Experiment 1. Taken together with our previous results, this bears out our hypothesis that the preferred direction of comparison maximizes comparison informativity. To complete the picture, in Experiment 6 we tested a final prediction of the directional informativity hypothesis: forward comparisons should be considered more informative than reverse comparisons. Evidence supporting this prediction would suggest that people do indeed consider modes of target modification such as inference projection and restructuring when asked to select which direction of comparison they prefer between two items.

We asked subjects to rate the informativity of forward and reverse comparisons. In keeping with our functional approach to asymmetry, we predicted that the forward direction of comparison would be rated as more informative than the reverse direction. We further predicted that this asymmetry in informativity would be greater for the replacement and shuffled pairs than for the deletion pairs. This result would explain the high comparison asymmetry of the replacement and shuffled pairs relative to the deletion pairs observed in Experiments 1 through 3, and would mirror the pattern of modification asymmetries observed in Experiment 5.

Method

Subjects

Forty-eight Northwestern University undergraduates participated in partial fulfillment of a course requirement. None had previously participated in a similar experiment.

Materials and Design

The materials and design were identical to those used in Experiments 1, 4, and 5.

Procedure

Subjects read 18 pairs of stories, and for each story pair they rated how informative they felt the comparison statement "Story A is like Story B" was as a statement about Story A. Subjects




FIG. 6. Informativity ratings as a function of comparison direction and variant condition in Experiment 6.

rated the informativity of each comparison by circling the appropriate number on a scale of 1 (very uninformative) to 7 (very informative) provided at the bottom of the page.

Results and Discussion

The results are summarized in Fig. 6. A 3 (variant condition: deletion, replacement, shuffled) \times 2 (direction: standard-to variant, variant-to-standard) repeated-measures ANOVA on the rated informativity of the comparisons confirmed a main effect of direction: forward comparisons ($M = 5.37$) were rated as more informative than reverse comparisons ($M = 2.49$) across both subjects and items, $F_S(1,47) = 343.02$, $p < .001$ and $F_I(1,17) = 383.99$, $p < .001$. There was also a main effect of variant condition across both subjects and items, $F_S(2,94) = 5.99$, $p < .005$ and $F_I(2,34) = 3.62$, $p < .05$. Planned comparisons indicated that, averaging across both comparison directions, the deletion pairs ($M = 4.17$) were rated as more informative than the replacement pairs ($M = 3.80$), $t_S(47) = 2.73$, $p < .01$ and $t_I(17) = 2.43$, $p < .05$, or the shuffled pairs ($M = 3.82$), $t_S(47) = 2.85$, $p < .01$ but not significant by items. Critically, the predicted interaction between direction and variant condition was obtained across both subjects and items, $F_S(2,94) = 26.43$, $p < .001$ and $F_I(2,34) = 14.13$, $p < .001$. Planned comparisons indicated that the difference in informativity between the two directions was greater for the shuffled pairs ($M = 3.70$) than for the replacement pairs ($M = 2.85$), $t_S(47) = 3.76$, $p < .001$ and $t_I(17) = 2.99$, $p < .01$, and for the replacement pairs than for the deletion pairs ($M = 2.08$), $t_S(47) = 3.63$, $p < .001$ and $t_I(17) = 2.27$, $p < .05$.

As predicted by the directional informativity hypothesis, subjects rated comparisons between aligned items as more informative when the item with greater systematicity was in the base position. Further, the magnitude ordering of this difference corresponded to the magnitude ordering of both comparison asymmetries (Experiment 1) and modification asymmetries (Experiment 5), with replacement and shuffled pairs exhibiting greater directionality than deletion pairs on all three measures. In fact, there was a significant correlation between the directional differences in informativity from this experiment and the comparison asymmetries from Experiment 1, $r = .63$, $p < .001$. These findings demonstrate the connection between systematicity imbalance and informativity, and suggest that people are indeed sensitive to the types of target modifications observed in Experiment 5 (e.g., inference projection, restructuring) when comparing two items. In sum, comparison asymmetries reflect the fact that people prefer informative comparisons over non-informative ones.

GENERAL DISCUSSION

Based on two general pragmatic principles—the maxim of informativity and the given-new contract—we have proposed that a major source of comparison asymmetry is directional informativity. Given the process model of comparison specified by structure-mapping theory, the directional informativity hypothesis makes four predictions. First, people will prefer the direction of comparison that places the more systematic item in the base position. Second, asymmetries should only be observed when the items are alignable. Third, systematicity imbalance should be a better predictor of asymmetries than differences in the number or salience of distinctive features. Fourth, the preferred direction of comparison should (a) be judged as more informative than the non-preferred direction, and (b) should support more candidate inferences and other forms of target modification than the non-preferred direction. All four of these predictions were borne out.

In Experiment 1, we found a clear preference for placing the more systematic of two passages in the base position of a comparison. Further, the degree of this preference increased with the degree of systematicity imbalance. Experiments 1 and 2 further showed that alignability is crucial in obtaining asymmetries. In particular, the base systematicity advantage disappeared in comparisons between non-alignable items in Experiment 2, despite substantial differences in the systematicity of the individual items. These results argue against potential explanations of our findings in terms of individual stimulus biases such as familiarity or typicality (see Nosofsky, 1991). These results are also not predicted by Tversky's (1977) feature contrast model of similarity, in which asymmetries result from an imbalance in the size or salience of distinctive feature sets defined independently of the aligned structure. Experiment 3 further tested the directional informativity hypothesis against the feature contrast model. According to this model, adding distinctive features to the preferred base of a comparison should increase the asymmetry, whereas

adding them to the preferred target should decrease or even reverse the asymmetry. Instead, consistent with structure-mapping theory, we found that the effect of distinctive features depends on their relation to the aligned system (Gentner & Markman, 1994; Markman & Gentner, 1993b). Specifically, adding structurally irrelevant features to preferred comparison targets or bases had no effect on asymmetries. Further, adding structurally inconsistent features to preferred comparison bases actually *decreased* asymmetries, contrary to the predictions of the feature contrast model. Thus, asymmetries are predictable in terms of systematicity imbalance, and not relative numbers of distinctive features.

Going to the heart of the matter, Experiments 4 and 5 examined how comparisons invite changes in the representation of the target. We found that people preferred generating modifications from more systematic to less systematic items over the reverse, and that these modifications most often involved information associated with the common relational structure. Further examination of the responses indicated that different kinds of target modifications (e.g., inference projection, restructuring) contributed to the directional informativity of different kinds of comparisons. Finally, Experiment 6 directly tested a central tenet of our functional approach: that asymmetric preferences in similarity stem from differences in perceived informativity. Put simply, our claim is that (a) people prefer informative comparisons, and (b) a comparison is more informative when the base is more systematic or coherent than the target. Consistent with this claim, comparisons having the more systematic item as the base were rated as more informative than reverse comparisons. Further, the degree of this difference in informativity across items corresponded to both the degree of comparison asymmetry (Experiment 1) and the degree of modification asymmetry (Experiment 5). In short, comparison asymmetries arise from the directionality of the mapping process.

Systematicity and Related Notions

A key notion in our discussion of the directional informativity hypothesis has been *systematicity*, or the degree to which a conceptual system consists of interconnected predicates governed by higher-order relations (e.g., Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1982, 1983, 1989; Gentner & Toupin, 1986). Because it is a relatively formal construct, systematicity is applicable to a range of relational types. Although we have focused primarily on causal and explanatory higher-order relations as determinants of systematicity, other types of relations can serve to provide structured representations, both conceptual (e.g., Barr & Caplan, 1987; Barsalou, 1992; Gentner & Rattermann, 1991; Gentner, Rattermann, Markman, & Kotovsky, 1995; Novick, 1988; Tversky & Hemmenway, 1984; Winston, Chaffin, & Herrmann, 1987) and perceptual (e.g., Biederman, 1987; Halford, 1992, 1993; Kotovsky & Gentner, 1996; Markman & Gentner, 1993a; Palmer, 1977, 1978; Posner & Keele, 1968).

The notion of systematicity is related to several psychological constructs

that have sometimes been called *coherence*.⁶ Systematicity is perhaps most clearly related to the sense of coherence used in theories of text processing, especially those that emphasize causal relations. Many of these theories stress the importance of sequential chains of causally connected events, actions, and states that lead from the beginning to the end of a narrative (e.g., Black & Bower, 1980; Schank & Abelson, 1977). Other theories have focused predominantly on causal networks (e.g., Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985) and relational hierarchies (e.g., Johnson & Mandler, 1980; Kintsch & van Dijk, 1978). These latter approaches stress that it is not the sheer number of causal relations that is crucial for a text to be coherent; rather, the degree of interconnectedness and mutual constraint between such relations is primary.⁷

Systematicity is also relevant to the structure of everyday concepts such as *horse* or *hammer*. A number of recent theories explain concept-level coherence in terms of inherently relational systems of world knowledge, such as mental models and intuitive theories (e.g., Carey, 1985; Gentner & Stevens, 1983; Johnson-Laird, 1983; Keil, 1989; Lakoff, 1987; Murphy & Medin, 1985). For example, Murphy and Medin (1985) suggest that concept-level coherence is derived both internally, from causal or functional relations between attributes, and externally, from the position of the concept in causal or explanatory networks existing between multiple concepts. In a similar vein, Barr and Caplan (1987) have argued for the role of external relations in defining concepts, and Barsalou (1991) suggests that the causal and functional information found in goal structures can serve to create coherent categories. If concept representations include the kinds of relational systems manipulated in our studies, then the directional informativity hypothesis and the notion of systematicity imbalance are applicable not only to comparisons between complex scenarios, but also to comparisons between individual concepts.

For example, consider again the comparison between North Korea and Red China, for which there exists a clear directional preference. We suggest that *North Korea is similar to Red China* is preferred over the reverse comparison because our representation of China possesses a greater degree of systematicity: our knowledge of the country (e.g., its history, culture, economy, and relationships with other countries) is richer and more highly structured than that of North Korea. The preferred direction of comparison maximizes informativity and enables us to make candidate inferences about North Korea, such as how the communist system might interact with traditional culture, how the country might treat student dissidents, and so on. This could explain why prototypical or reference point concepts are not only preferred *comparison* bases (e.g., Rosch, 1975; Tversky, 1977) but also preferred *inferential* bases (Rips, 1975): such concepts

⁶ We thank Keith Holyoak for a useful discussion of different notions of coherence.

⁷ Text *coherence* should not be confused with text *cohesion*, which refers to interrelations between the surface elements of a text. For a review of the differences between coherence and cohesion, see de Beaugrand and Dressler (1981).

tend to be more systematic in their internal structure and/or their relational role in larger external structures. The effects of systematicity imbalance should be especially powerful when one comparison item is informationally ambiguous relative to the other (see Herr, Sherman, & Fazio, 1983).

The above arguments can be extended to the domain of social concepts, particularly those of *self* and *other*. One consistent finding in the literature is that subjects rate others (both individuals and groups) and being more similar to the self than vice versa (e.g., Catrambone, Beike, & Niedenthal, 1996; Holyoak & Gordon, 1983; Srull & Gaelick, 1983). A second finding is that subjects are generally more willing to make inferences and predictions from the self to others than vice versa (e.g., Kunda & Nisbett, 1988; McFarland & Miller, 1990). Although these two lines of research have progressed independently, we suggest that both of these effects have a common source: a base coherence advantage favoring the self concept. The construct of systematicity imbalance may thus help to explain several related phenomena in social cognition, such as the false consensus effect, wherein people tend to overestimate the proportion of others sharing their own beliefs (Ross, Greene, & House, 1977).

As noted by Palmer (1977, 1978), the notion of hierarchical relational structures is relevant to perceptual as well as conceptual representations. In particular, Palmer suggests that the construct *goodness-of-form* is reducible to the degree of relational interconnectivity between perceptual elements. This suggests that systematicity imbalance may also be applicable to comparisons between perceptual stimuli. In support of this notion, Medin, Goldstone, and Gentner (1993) found that structurally ambiguous figures tend to be reinterpreted in terms of available non-ambiguous figures rather than the reverse. In fact, Tversky (1977) found that subjects preferred placing good forms in the base position of comparisons, despite the fact that these same stimuli were rated as being less complex (and hence, as having fewer distinctive features) than less good forms.

Of course, the notion of systematicity as embodied in structure-mapping theory and SME may yet be further refined, as shown by Thagard's (1989, 1992) discussion of *explanatory coherence*. Explanatory coherence is meant to apply scientific theories and hypotheses. According to Thagard, hypotheses gain explanatory coherence to the extent that they explain the evidence at hand, are themselves explained by higher-order hypotheses, and are analogous to accepted hypotheses. Thagard's notion of explanatory coherence can enrich the notion of systematicity in three ways. First, it suggests that the coherence of a given system of relations is influenced by its fit both with the external environment and with general background knowledge. Second, it suggests that not all higher-order relations contribute equally to coherence: some may explain the given evidence better than others. We further speculate that higher-order constraining relations should not be considered as merely present or absent, but rather as varying in their salience or availability. Future research should investigate whether greater specificity in characterizing higher-order relations could improve predictions of comparison asymmetry with regard to systematicity imbalance.

Alternative Explanations of Asymmetry

Systematicity and complexity. A possible challenge to our account of asymmetry that would salvage the feature contrast model would be to argue that the higher-order relations in the standard stories make these stories more informationally *complex* than the less systematic variants. On this account, the observed directional preferences would be derived from the standard stories' possessing more conceptual features than the variants. However, the results of Experiment 3 argue against this possibility. Adding structurally inconsistent statements to stories increased these stories' complexity, but at the same time decreased their likelihood of being placed in the base position of a comparison. This fits our claim that the influence of distinctive features on asymmetry depends on their relation to the common system. Further, recent research suggests that the coherence of a representation may be established at the expense of detailed information. For example, Barton and Sanford (1993) found that establishing global coherence of a passage often results in shallow processing of narrative details, such that subjects were unable to detect anomalies in passages that had been incorporated into a global model. This suggests that coherence may actually lead to a loss of surface information or detail. Again, these findings argue against an explanation of our results in terms of relative complexity or relative numbers of distinctive features.

Systematicity and salience. Our approach to asymmetry might also be challenged by arguing that the higher-order relations providing systematic structure to the standard stories were more *salient* than the causal gaps or anomalous sentences that replaced them in the variant stories. This possibility could allow the feature contrast model to explain several of our results by recasting the base systematicity advantage in terms of greater saliency of distinctive features in the base. However, evidence presented by Kemper (1982) suggests that information gaps leading to violations of the causal structure of texts are more easily detected than other kinds of gaps. Coherence-disrupting alterations have also been found to be particularly salient in text (Bharucha, Olney, & Schnurr, 1985). More generally, the unexpected and the atypical have been found to be highly salient in script-based stories (Bower, Black, & Turner, 1979; Graesser, Gordon, & Sawyer, 1979), tonal musical sequences (Bharucha, 1984), and person descriptions (Hastie & Kumar, 1979; Woll & Graesser, 1982). Taken together, these findings suggest that structurally inconsistent alterations are typically *more* salient than structurally consistent higher-order relations of the type found in our more coherent standard stories. We speculate that the high salience of causal gaps and disruptions may be useful, as it could alert the hearer to the possibility of importing structure from a more systematic base domain.

Other sources of asymmetry in comparisons. We would not wish to claim that systematicity is the only source of asymmetry in comparisons. Certain types of stimulus bias may also contribute to asymmetry (e.g., Holman, 1979; Krumhansl, 1978; Nosofsky, 1991). For example, a frequently encountered

item may become a preferred comparison base relative to a rarely encountered item. Because such biases exist independently of the particular comparison, they are not reducible to systematicity imbalance, and probably have little impact on comparison informativity. On the other hand, systematicity imbalance is not easily reduced to such constructs as feature contrast or stimulus bias, as we have argued above. More importantly, our results clearly point to sources of directional informativity that are not captured by other models of similarity. While systematicity imbalance may not be the only source of asymmetries, it is the only one that predicts asymmetries in both directional preference and informativity, and it follows naturally from an independently motivated process model of comparison.

Is Similarity Itself Asymmetric?

By combining the explicit processing assumptions of structure-mapping theory with consideration of the pragmatics of comparison statements, we can account for comparison asymmetries without needing to hypothesize underlying asymmetries in subjective similarity judgments. However, our account raises an important question: If comparisons are often asymmetric, then is similarity itself an asymmetric relation? At the risk of equivocation, our answer is both yes and no. Because the initial process of structural alignment is theoretically symmetric (as demonstrated in the SME simulation), structure-mapping theory suggests that the initial similarity of two items is also symmetric. However, inference projection and other forms of target modification are directional. Thus, asymmetries in subjective similarity could reflect differential informativity. Going further, if one direction of comparison suggests more candidate inferences than the other, then its target will end up being more alignable with—and therefore more similar to—its base as a result of the projected information (see also Lassaline, 1996). Thus it is possible to explain comparison asymmetry without abandoning the intuition that similarity is a symmetric relation.

This position is akin to that of stimulus-bias models, in which asymmetric proximities can be derived from a symmetric similarity function plus a differential stimulus bias (Nosofsky, 1991). It is also consistent with reference-point models (e.g., Gleitman, Gleitman, Miller, & Ostrin, 1996; Huttenlocher, Hedges, & Duncan, 1991; Rosch, 1975; Shen, 1989). For example, Gleitman et al. link predicates of comparison such as *similar to* and *like* with a larger set of symmetric predicates whose arguments can be reversed without a radical change in meaning. This set includes *equal to*, *near*, *meet*, *marry*, and so on. Despite their semantic symmetry, asymmetric preferences can be found over any of these predicates. According to Gleitman et al., such asymmetries are due to a grammatical preference for placing the more prominent item in the base position of the comparison, rather than to an underlying asymmetry in judgment. For example, we prefer to say *Sally met the Pope* rather than *The Pope met Sally*. But as Gleitman et al. note, this asymmetry cannot be said to derive from any underlying asymmetry in the sense of *meeting*: it makes little

sense to say that *Sally met the Pope* is more true than *The Pope met Sally*, or that the first event ranks higher on some underlying scale of "meetingness."

Gleitman et al. provide an elegant account of asymmetric preference across various predicate types. Their account captures an important sense in which similarity is a symmetric relation with asymmetric felicity conditions. But there is another side to similarity. As discussed above, similarity statements allow transfer of information from base to target. In this respect, similarity is quite different from other symmetric predicates such as *near*, *meet*, and *marry*: these predicates convey information, but of a different sort (e.g., Sally does not inherit papal traits by meeting with the Pope). In other words, similarity is more dynamic than other symmetric predicates. Similarity is a symmetric *relation*, but an asymmetric *process*.

Asymmetries in Analogical Transfer

Our approach to asymmetry in literal similarity is consistent with recent approaches to analogy (e.g., Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 1989; Gick & Holyoak, 1980, 1983; Halford, 1992, 1993; Holyoak, 1985; Holyoak & Thagard, 1989; Keane, 1988) and case-based reasoning (e.g., Kass, 1994; Kolonder, 1993; Schank, 1982) that have postulated complex system mappings to explain information transfer. Most of these models emphasize the importance of relational structure (see Clement & Gentner, 1991; Spellman & Holyoak, 1992). In particular, some computational models of analogy postulate that mapping consists of an alignment stage, in which systematic correspondences are established between the relational structures of the base and target analogs, followed by the projection of inferences from the base to the target (e.g., Falkenhainer, Forbus, & Gentner, 1989; Holyoak & Thagard, 1989). One implication of such models is that the directional informativity hypothesis should apply not only to literal similarity comparisons, but also to analogical comparisons. In fact, asymmetries in mapping have been reported in a number of studies of analogical reasoning (e.g., Bassok & Holyoak, 1989; Burns, 1996; Reed, Ernst, & Banerji, 1974). We suggest that such asymmetries are best explained in terms of systematicity imbalance: people prefer the more systematic or coherent representation as the base of an analogy.

Asymmetries in Metaphors

As noted by a number of authors, metaphors are often radically asymmetric (e.g., Glucksberg & Keysar, 1990; Glucksberg, McGlone, & Manfredi, 1997; Ortony, 1979; Ortony, Vondruska, Foss, & Jones, 1985). In contrast to literal similarity comparisons, many metaphors are anomalous when reversed. For example, although the literal comparison *North Korea is similar to China* is preferred over its reverse, both directions of comparison are interpretable. However, whereas the metaphor *Life is a journey* is an informative statement, *A journey is life* is nonsensical. These observations have been cited as evidence for different sources of directionality in literal versus metaphoric state-

ments. According to Ortony's (1979) *salience imbalance model*, metaphors differ from literal comparisons in that the common features of metaphors are far more salient in the base than in the target. For example, in *Time is a river*, the common feature *flows* is more salient for rivers than for time. On this account, metaphoric asymmetries stem from a requirement for maintaining this direction of salience imbalance. Glucksberg and Keysar's (1990) *class-inclusion model* also draws a sharp distinction between literal comparisons and metaphors. They state that metaphors are not comparison statements but rather categorization statements, in which the target is assigned to an ad hoc category generated by the base. For example, in *Time is a river*, *river* evokes a category for which it could serve as a prototype (e.g., *things that flow forward*). *Time* is then understood as being a member of this category. On this account, the concept that best provides a predicated category is the preferred base, and metaphors are nonreversible for the same reasons that literal categorization statements are nonreversible (for similar proposals, see Honeck, Kibler, & Firment, 1987; Kennedy, 1990; Shen, 1992).

However, we believe that metaphors share essentially the same processes of structural alignment and mapping as literal similarity comparisons and analogies (e.g., Gentner, 1982; Gentner & Clement, 1988; Gentner, Falkenhainer, & Skorstad, 1987; Gentner & Wolff, 1997, in press). Indeed, a number of influential approaches to metaphor comprehension have suggested processes that are broadly compatible with structure-mapping theory. For example, Verbrugge and McCarrell (1977) argue that metaphors involve a novel schematization of the target in terms of the relational structure of the base. Likewise, Lakoff and Johnson (1980) claim that metaphors allow abstract concepts to be understood in terms of more coherent and imageable domains of experience. These approaches to metaphor, in which structural information is projected from the base to the target, suggest that metaphoric asymmetries may best be predicted by systematicity imbalance. We propose that the extreme directionality of metaphors is due to greater levels of systematicity imbalance than are typically found in literal comparisons. This claim is supported by the common observation that, unlike literal comparisons, metaphors often set up mappings between abstract targets and concrete bases (e.g., Connor & Kogan, 1980; Katz, 1989; Lakoff & Johnson, 1980).

There are some commonalities between the salience, imbalance, and class-inclusion approaches to metaphoric asymmetry and our own approach. In particular, all three models stress that metaphors are directionally informative. The salience imbalance model allows for *predicate promotions* in metaphors, in which the salience of a common feature is raised in the target term due to its emphasis in the base (Ortony, 1979, 1993); the class-inclusion model claims that the ad hoc category generated by the base is used to project novel attributes to the target via *feature inheritance*; and structure-mapping theory postulates post-alignment processes of inference projection and other forms of target modification. However, whereas both the salience imbalance model and the class-inclusion model draw a qualitative distinction between metaphor

and literal similarity based on asymmetries, structure-mapping theory provides a unified framework for describing asymmetries in all types of comparisons, be they literal, analogical or metaphorical.

CONCLUSIONS

In this paper, we have used two general pragmatic principles—the maxim of informativity and the given-new contract—to propose that a major source of comparison asymmetries is directional informativity. Adopting the process model of comparison specified by structure-mapping theory, we further hypothesized that asymmetries could best be characterized in terms of systematicity imbalance. Across six experiments, we examined the representations, processes, and products of directional comparisons, and found robust evidence favoring the directional informativity hypothesis, and more generally for structure-mapping as a theory of similarity.

By viewing comparison as a dynamic process, we can account for factors influencing asymmetry that cannot be handled by more traditional approaches to similarity. These include the necessity of alignment for directionality, the importance of distinguishing between structurally consistent and inconsistent distinctive information, and the role played by target modification processes such as inference projection and restructuring. Further, structure-mapping theory allows us to explain directional preferences without recourse to underlying asymmetries in similarity judgments. Finally, we suggest that directional informativity provides a unified account of asymmetry across a range of comparison types, including literal similarity, analogy, and metaphor.

REFERENCES

- Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and Cognition*, **15**, 397–418.
- Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, vol. 27. New York: Academic Press.
- Barsalou, L. W. (1992). Frames, concepts, and conceptual fields. In A. Lehrer & E. F. Kittay (Eds.), *Frames, fields, and contrasts*. Hillsdale, NJ: Erlbaum.
- Bartlett, J. C., & Dowling, W. J. (1988). Scale structure and similarity of melodies. *Music Perception*, **5**, 285–315.
- Barton, S. B., & Sanford, A. J. (1993). A case-study of pragmatic anomaly detection: Relevance-driven cohesion patterns. *Memory and Cognition*, **21**, 477–487.
- Bassok, M., & Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **15**, 153–166.
- Bharucha, J. J. (1984). Anchoring effects in music: The resolution of dissonance. *Cognitive Psychology*, **16**, 485–518.
- Bharucha, J. J., Olney, K. L., & Schnurr, P. P. (1985). Detection of coherence-disrupting and coherence-conferring alterations in text. *Memory and Cognition*, **13**, 573–578.
- Biederman, I. (1987). Recognition-by-components: A theory of human-image understanding. *Psychological Review*, **94**, 115–147.
- Black, J. B., & Bower, G. H. (1980). Story understanding as problem solving. *Poetics*, **9**, 223–250.

- Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, **11**, 177–220.
- Burns, B. D. (1996). Meta-analogical transfer: Transfer between episodes of analogical reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**, 1032–1048.
- Burstein, M. H. (1983). 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*. Boston: Morgan Kaufmann.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT.
- Carroll, J. D., & Wish, M. (1974). Multidimensional perceptual models and measurement methods. In Carterette, E. C., & Friedman, M. P. (Eds.), *Handbook of perception*. New York: Academic Press.
- Catrabone, R., Beike, D., & Niedenthal, P. (1996). Is the self-concept a habitual referent in judgments of similarity? *Psychological Science*, 1–6.
- Clark, H. H., & Haviland, S. E. (1977). Comprehension and the given-new contract. In R. O. Freedle (Ed.), *Discourse production and comprehension*. Norwood, NJ: Ablex.
- 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 analogical retrieval. *Journal of Memory and Language*, **33**, 396–420.
- Connor, K., & Kogan, N. (1980). Topic-vehicle relations in metaphor: The issue of asymmetry. In Honeck, R. P., & Hoffman, R. R. (Eds.), *Cognition and figurative language*. Hillsdale, NJ: Erlbaum.
- de Beaugrande, R. A., & Dressler, W. U. (1981). *Introduction to text linguistics*. New York: Longman.
- Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, **41**, 1–63.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, **19**, 141–205.
- Gentner, D. (1982). Are scientific analogies metaphors? In D. S. Miall (Ed.), *Metaphor: Problems and perspectives*. Brighton, Sussex: Harvester.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, **7**, 155–170.
- Gentner, D. (1989). The mechanisms of analogical learning. In Vosniadou, S., & Ortony, A. (Eds.), *Similarity and analogical reasoning*. New York: Cambridge University.
- Gentner, D., Brem, S., Ferguson, R. W., Wolff, P., Markman, A. B., & Forbus, K. D. (1997). Analogy and creativity in the works of Johannes Kepler. In T. B. Ward, S. M. Smith & J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes*. Washington, DC: American Psychological Association.
- Gentner, D., & Clement, C. A., (1988). Evidence for relational selectivity in interpreting analogy and metaphor. In Bower, G. (Ed.), *The psychology of learning and motivation*. New York: Academic Press.
- Gentner, D., Falkenhainer, B., & Skorstad, J. (1987). Metaphor: The good, the bad, and the ugly. In D. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence, computer science, and philosophy*. Boston: Reidel.
- Gentner, D., & Markman, A. B. (1993). Analogy: Watershed or Waterloo? Structural alignment and the development of connectionist models of cognition. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in neural information processing systems 5*. San Mateo, CA: Morgan Kaufmann.
- Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science*, **5**, 153–158.
- Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In Gelman,

- S. A., & Byrnes, J. P. (Eds.), *Perspectives on language and thought*. New York: Cambridge University.
- Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, **25**, 524–575.
- Gentner, D., Rattermann, M. J., Markman, A. B., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In G. Halford & T. Simon (Eds.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Erlbaum.
- Gentner, D., & Stevens, A. L. (Eds.) (1983). *Mental models*. Hillsdale, NJ: Erlbaum.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, **10**, 277–300.
- Gentner, D. & Wolff, P. (1997). Alignment in the processing of metaphor. *Journal of Memory and Language*, **37**, 331–355.
- Gentner, D. & Wolff, P. (in press). Metaphor and knowledge change. In E. Dietrich & A. B. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines*. Cambridge, MA: MIT.
- 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**, 1–38.
- Gleitman, L., Gleitman, H., Miller, C., & Ostrin, R. (1997). Similar, and similar concepts. *Cognition*, **58**, 321–376.
- Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, **97**, 3–18.
- Glucksberg, S., McGlone, M. S., & Manfredi, D. (1997). Property attribution in metaphor comprehension. *Journal of Memory and Language*, **36**, 50–67.
- Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**, 3–28.
- Goldstone, R. L., & Medin, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**, 29–50.
- Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, **23**, 222–262.
- Graesser, A. C., Gordon, S. E., & Sawyer, J. D. (1979). Memory for typical and atypical actions in scripted activities: Test of a script pointer + tag hypothesis. *Journal of Verbal Learning and Verbal Behavior*, **18**, 319–332.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics*, vol. 3. New York: Academic Press.
- Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development*.
- Halford, G. S. (1993). *Children's understanding: The development of mental models*. Hillsdale, NJ: Erlbaum.
- Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, **37**, 25–38.
- Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology*, **19**, 323–340.
- Holman, E. W. (1979). Monotonic models for asymmetric proximities. *Journal of Mathematical Psychology*, **20**, 1–15.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In Bower, G. H. (Ed.), *The psychology of learning and motivation: Vol 19*. New York: Academic Press.
- Holyoak, K. J., & Gordon, P. C. (1983). Social reference points. *Journal of Personality and Social Psychology*, **44**, 881–887.
- Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory and Cognition*, **15**, 332–340.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, **13**, 295–355.

- Honeck, R. P., Kibler, C. T., & Firment, M. J. (1987). Figurative language and psychological views of categorization: Two ships in the night? In R. E. Haskell (Ed.), *Cognition and symbolic structures: The psychology of metaphoric transformation*. Norwood, NJ: Ablex.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, **98**, 352–376.
- Johnson, N. S., & Mandler, J. M. (1980). A tale of two structures: Underlying and surface forms in stories. *Poetics*, **9**, 51–86.
- Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University.
- Karmiloff-Smith, A. (1991). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT.
- Kass, A. (1994). Tweaker: Adapting old explanations to new situations. In Schank, Kass, & Riesbeck (Eds.), *Inside case-based reasoning*. Hillsdale, NJ: Erlbaum.
- Katz, A. N. (1989). On choosing the vehicles of metaphors: Referential concreteness, semantic distances, and individual differences. *Journal of Memory and Language*, **28**, 486–499.
- Keane, M. (1988). *Analogical problem solving*. Chichester, England: Ellis Horwood.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT.
- Kemper, S. (1982). Filling in the missing gaps. *Journal of Verbal Learning and Verbal Behavior*, **21**, 99–107.
- Kennedy, J. M. (1990). Metaphor—Its intellectual basis. *Metaphor and Symbolic Activity*, **5**, 115–123.
- Kintsch, W., & van Dijk, T. A. (1978). Towards a model of text comprehension and production. *Psychological Review*, **85**, 363–394.
- Kolonder, J. L. (1993). *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, **67**, 2797–2822.
- Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, **85**, 445–463.
- Kunda, Z., & Nisbett, R. E. (1988). Predicting individual evaluations from group evaluations and vice versa: Different patterns for self and other? *Personality and Social Psychology Bulletin*, **14**, 326–334.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago.
- Lassaline, M. E., (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**, 754–770.
- Markman, A. B., & Gentner, D. (1993a). Structural alignment during similarity comparisons. *Cognitive Psychology*, **25**, 431–467.
- Markman, A. B., & Gentner, D. (1993b). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, **32**, 517–535.
- McFarland, C., & Miller, D. T. (1990). Judgments of self-other similarity: Just like other people, only more so. *Personality and Social Psychology Bulletin*, **16**, 475–484.
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, **100**, 254–278.
- Minsky, M. L. (1977). A framework for representing knowledge. In Winston, P. H. (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, **92**, 289–316.
- Nisbett, R. E., & Kunda, Z. (1985). The perception of social distributions. *Journal of Personality and Social Psychology*, **48**, 297–311.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, **23**, 94–140.
- Novick, L. R., & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **17**, 398–415.
- Ortony, A. (1979). Beyond literal similarity. *Psychological Review*, **86**, 161–180.

- Ortony, A., Vondruska, R. J., Foss, M. A., & Jones, L. E. (1985). Salience, similes, and the asymmetry of similarity. *Journal of Memory and Language*, **24**, 569–594.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, **9**, 441–474.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In Rosch, E., & Lloyd, B. B. (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, **77**, 353–363.
- Read, S. J. (1984). Analogical reasoning in social judgment: The importance of causal theories. *Journal of Personality and Social Psychology*, **46**, 14–25.
- Read, S. J. (1987). Similarity and causality in the use of social analogies. *Journal of Experimental Social Psychology*, **23**, 189–207.
- Reed, S. K., Ernst, G. W., & Banerji, R. (1974). The role of analogy in transfer between similar problem states. *Cognitive Psychology*, **6**, 436–440.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, **14**, 665–681.
- Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology*, **7**, 532–547.
- Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **13**, 629–639.
- Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **15**, 456–468.
- Ross, L., Greene, D., & House, P. (1977). The false consensus phenomenon: An attributional bias in self-perception and social perceptual processes. *Journal of Experimental Social Psychology*, **13**, 279–301.
- Rumelhart, D. E., & Ortony, A. (1978). The representation of knowledge in memory. In Anderson, R. C., Spiro, R. J., & Montague, W. E. (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.
- Schank, R. (1982). *Dynamic memory*. New York: Cambridge University.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, goals, plans, and understanding*. Hillsdale, NJ: LEA.
- Shen, Y. (1989). Symmetric and asymmetric comparisons. *Poetics*, **18**, 517–536.
- Shen, Y. (1992). Metaphors and categories. *Poetics Today*, **13**, 771–794.
- Shepard, R. N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, **39**, 373–421.
- Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, **62**, 913–933.
- Srull, T. K., & Gaelick, L. (1983). General principles and individual differences in the self as a habitual reference point: An examination of self-other judgments of similarity. *Social Cognition*, **2**, 108–121.
- Talmy, L. (1978). Figure and ground in complex sentences. In J. Greenberg, C. Ferguson, & M. Moravcsik (Eds.), *Universals of human language*, vol. 4. Stanford: Stanford University.
- Talmy, L. (1983). How language structures space. In H. Pick & L. Acredolo (Eds.), *Spatial orientation: Theory, research, and application*. New York: Plenum.
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, **12**, 435–502.
- Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University.
- Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, **24**, 595–611.
- Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, **24**, 612–630.
- Turner, M. (1988). Categories and analogies. In D. H. Helman (Ed.), *Analogical reasoning*. Kluwer.

- Tversky, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
- Tversky, A., & Gati, I. (1978). Studies of similarity. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, **113**, 169-193.
- Verbrugge, R. R., & McCarrell, N. S. (1977). Metaphoric comprehension: Studies in reminding and resembling. *Cognitive Psychology*, **9**, 494-533.
- Winston, M. E., Chaffin, R., & Herrmann, D. (1987). A taxonomy of part-whole relations. *Cognitive Science*, **11**, 417-444.
- Woll, S. B., & Graesser, A. C. (1982). Memory discrimination for information typical or atypical of person schemata. *Social Cognition*, **4**, 287-310.

(Accepted June 24, 1997)