Gentner, D., & Wolff, P. (2000). Metaphor and knowledge change. In E. Dietrich & A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines* (pp. 295-342). Mahwah, NJ: Lawrence Erlbaum Associates.

11

Metaphor and Knowledge Change

Dedre Gentner
Phillip Wolff
Northwestern University, Evanston

The tiger roaring like a fire, the fire roaring like a tiger, are metaphors by which both fire and tiger are made clear.

—Litt

In a recent magazine article, the issue of government control of research agendas was explored by using the metaphor "Science is a flashlight." Science, like a flashlight, is bounded in its scope: Only areas directly under its beam are visible. So far, the metaphor may seem obvious, but a further implication is that because of this boundedness, decisions about the direction of the beam are crucial. Until recently, the article pointed out, scientists have had the major say in where the flashlight was aimed, but increasingly it is government agencies, not scientists, who control the direction of search. The notion of controlling the direction of science was presumably not a new issue to readers. Nevertheless, the metaphor contributed a new connection: The fact that vision is limited makes it crucial who controls its direction.

There is broad consensus that metaphors can lead to change of knowledge. In some cases, the change seems one of enrichment: New concepts, connections, or perspectives are added to the underlying representations. In other cases, the change involves the re-representation of old concepts and/or the restructuring of the conceptual systems.

Yet despite this agreement, there is little consensus on *how* such change might occur: that is, on what processes might bring about knowledge change. Black (1979) asserted that metaphor may lead to emphasizing, suppressing, or reorganizing features of the terms, especially those of the primary subject (the target) (see also Verbrugge & McCarrell, 1977). He suggested that these changes may follow from the construction of a "parallel implication-complex," but the details of this process were left unspecified. Comparison models figure largely in the current approach to metaphor, yet, as noted by Glucksberg and Keysar (1990), pure comparison models do not offer a means by which knowledge change can be achieved. For example, Ortony's (1979a) salience imbalance model is relatively explicit as to which existing features are incorporated into a metaphor's interpretation—namely, those that are high-salient in the base and low-salient in the target. However, although Ortony noted that metaphors sometimes convey that features from the base (vehicle) should be attributed to the target (tenor), the salience imbalance framework does not provide a mechanism by which attempted predications can come about. Glucksberg and Keysar (1990) proposed that the process of understanding a metaphor may prompt generation of a new category and from this category the attribution of novel inferences (see also Glucksberg & Keysar, 1990; Keysar, 1989; Shen 1992; Way, 1991).

Some accounts have explained metaphor comprehension in terms of mappings between dimensional spaces (as in Rumelhart & Abrahamson's [1973] model of analogy). For example, Tourangeau and Sternberg (1981, 1982) proposed that metaphor understanding follows from the identification of dimensional mappings that contribute to high in-domain similarity and low between-domain similarity. Kittay and Lehrer (1981) discussed the domain-interaction view in terms of semantic field theory. When lexical items from one semantic field are transferred to another semantic field, the donor field provides structure for the second field. For metaphors that are "alive," Kittay and Lehrer predicted changes in the semantic relations governing the fields because typically the fields in such metaphors are structured differently. Kelly and Keil (1987) suggested that change of knowledge may occur at the level of the domains or semantic fields from which the target and base terms are drawn. They found evidence that metaphor comprehension could affect domain organization. Pairs of concepts that could form appropriate metaphors increased in rated similarity after subjects read other metaphors relating their domains. But although domain-interaction accounts help to specify the conditions under which change may occur, they do not specify how it occurs: by what processes semantic fields are aligned, especially when the attributes of the target (tenor) and base (vehicle) do not match perfectly. Ortony (1979a) called this problem "domain incongruence." Furthermore, any model of metaphoric change of meaning must deal with the fact that not just any adaptation can occur: People are quite selective in interpreting nonliteral comparisons.

In this chapter, we set forth four mechanisms of change of knowledge: *knowledge selection, projection, re-representation,* and *restructuring*. We also discuss two specific representational outcomes that may follow from application of these mechanisms: *stored categories* and *stored mappings*. We discuss research relevant to these mechanisms and outcomes and offer our own approach to metaphor, one based on considering metaphor as akin to analogy. Our perspective draws on distinctions outlined in Falkenhainer, Forbus, and Gentner (1986, 1989), Gentner (1983, 1989), Gentner and Markman (1993, 1995), and Medin, Goldstone, and Gentner (1993), in which comparison is viewed as a process of alignment and mapping between pairs of structured representations. As we show, this view of metaphoric comparison allows us to escape many of the limitations inherent in pure comparison accounts.

Finally, although our proposals are aimed at addressing metaphorically driven change of knowledge, we think their scope may be broader. As noted by Pylyshyn (1979) and Fodor (1975), there is an inherent paradox in the cognitive approach to knowledge acquisition. If existing concepts and schemas are the medium of our thinking, then how can something new be expressed or comprehended? Does not what is needed for the acquisition of a new idea presume the prior presence of the idea itself? Nowhere is this paradox encountered more squarely than in the arena of metaphor and analogy.

THE STRUCTURE-MAPPING ACCOUNT: METAPHOR AS ALIGNMENT AND MAPPING

Representational Assumptions

We assume representations that include explicit labeled relations and arguments, including higher-order relations that can take whole assertions as their arguments. The representational system must also be able to capture rich perceptual information, including detailed object descriptions and dimensional information. Formally, the representational elements are objects (or *entities*), object descriptors (called *attributes*), functions (which express dimensional information), and *relations* between representational elements. Attributes and relations are predicates with truth values. Functions differ from predicates in that they map from a set of arguments onto values other than truth values. For example, the assertion that the ball is red can be represented by using *red* as an attribute, as in (1) or by using *color* as a relation, as in (2) or by using *color* as a function, as in (3):

(1) *Red* (ball).
(2) *Color* (ball, red).
(3) *Color* (ball) = red.

These three representations reflect different construals, which we take to be psychologically meaningful. In Representation 1, redness is expressed as an independent attributional property of the ball. Representation 2 focuses on the *color* relation—redness is expressed as one of a set of alternative arguments to the *color* relation. In Representation 3, *color* is represented as a dimension with *red* as one of its values. As noted previously, whereas 1 and 2 can take truth values, 3 cannot. Computationally, this functional notation is convenient for expressing statements when the exact value of the quantity is not of immediate interest. For example, color (ball 1) = color (ball 2) states that the balls are the same color without having to specify *which* color. Functions are useful for representing dimensions. Our psychological assumption here is that physical quantities like height and weight, numerical quantities, and eventually many abstract qualities—such as wealth or status—are often represented as dimensions.

Structure Mapping

In our account of metaphor comprehension, we make two fundamental assumptions: (a) that metaphor comprehension typically involves a comparison process, and (b) that the comparison process is structure sensitive. This second assumption is drawn from structure-mapping theory. We first lay out this framework as it applies to analogy and similarity and then discuss its application to metaphor.

A central characteristic of this account is that analogy and other comparison processes involve an alignment of relational structure (Gentner, 1982, 1983, 1989). This alignment process finds matches that are *structurally consistent*: that is, that observe *parallel connectivity* and *one-to-one correspondence*. *Parallel connectivity* means that if two predicates correspond then their arguments must also correspond. *One-to-one correspondence* means that any element in one representation can correspond to at most one matching element in the other representation (Falkenhainer, Forbus, & Gentner, 1986, 1989; Gentner, 1983, 1989; Holyoak & Thagard, 1989). For example, when comparing the atom (the *target* domain) to the solar system (the *base* domain), if *revolve* (planet, sun) is matched to *revolve* (electron, nucleus), then by parallel connectivity the sun must correspond to the nucleus and the planet to the electron, because they play the same role in the common relational structure. By one-to-one correspondence, these object bindings must be unique.[1]

A central claim is that in understanding analogies we seek matching *connected systems of relations.* A matching set of relations interconnected by higher-order constraining relations makes a better analogical match than do an equal number of matching relations that are unconnected to one another. To put it another way, an individual match matters more, and is more likely to be included in the final interpretation of the comparison, if it is connected via higher-order relations to other relations that also match. We call this tendency to align and map interconnected systems of predicates the *systematicity principle* (Gentner, 1983, 1989). This preference for connected systems is also what drives new inferences from a comparison. When the alignment process has resulted in a "best" interpretation,[2] then any predicates that exist in the base but not the target and that are *connected to the common system* that constitutes the comparison's interpretation can be imported into the target as *candidate inferences.* The systematicity principle thus represents a tacit preference for coherence and inferential potential in interpreting comparisons.

We suggest that structure mapping provides a framework for other comparison types as well as for analogy. Indeed, process models of structural alignment and mapping have proved fruitful in the study of overall similarity (literal similarity) comparisons (Bowdle & Gentner, 1997; Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1989; Gentner & Markman, 1993, 1994, 1995, 1997; Goldstone, 1994; Goldstone, Gentner, & Medin, 1989; Goldstone & Medin, 1994; Goldstone, Medin, & Gentner, 1991; Markman & Gentner, 1993a, 1993b, 1996; Medin, Goldstone, & Gentner, 1990, 1993). Analogy differs from overall similarity in that in literal similarity, the corresponding objects are similar to one another. In contrast, analogy is characterized by *relational focus*: Objects correspond by virtue of playing like roles in the relational structure rather than by any inherent object similarity. For instance, the nucleus need not be hot and gaseous like the sun. A particularly striking example of this sort of structural dominance is a *cross-mapping* (Gentner & Toupin, 1986), in which similar objects play different relational roles in two analogous scenarios (see also Gentner & Rattermann, 1991; Ross, 1989), for instance, grandmother is to mother as mother is to daughter.

---

[1]The nine planets, because they are relationally equivalent, can be treated as one generic planet.

[2]The best interpretation is determined by a number of factors, including whether it is structurally consistent and maximal in size and depth (Falkenhainer, Forbus, & Gentner, 1989; Gentner, Rattermann, & Forbus, 1993) and whether it is contextually relevant (Forbus & Oblinger, 1990; Holyoak & Thagard, 1989; Spellman & Holyoak, 1996).

There is considerable evidence supporting the general assumptions of he structure-mapping theory of comparison, as well as its application to netaphor. A central idea of this model is that processing comparisons nvolves the matching of structured representations (as opposed to lists f independent features), an assumption that is in line with a substantial mount of research in cognitive psychology (e.g., Biederman, 1987; Genter & Stevens, 1983; Johnson-Laird, 1983; Lassaline & Murphy, 1996; Lockead & King, 1977; Markman, 1999; Murphy, 1996; Murphy & Medin, 985; Norman, Rumelhart, & the LNR Group, 1975; Palmer, 1977, 1978; 'omerantz, Sager, & Stoever, 1977; Rumelhart & Ortony, 1977; Schank & belson, 1977). One way to observe the effects of structure in comparison s to use cross-mappings to put structural commonalities in conflict with ther kinds of similarities (e.g., of features or objects; Gentner & Toupin, 986; Goldstone & Medin, 1994; Markman & Gentner, 1993b; Rattermann, entner, & DeLoache, 1989, 1999; Ross, 1987). For example, Markman nd Gentner (1993b) showed people two scenes. In one, a women was hown *giving* food to a squirrel; in the other, the women was shown *ceiving* food from a man. One group of subjects rated how similar the wo scenes were to each other, while another group simply rated the two cenes' aesthetic value (to control for time spent looking at the pictures). ll subjects were then asked to say which thing in the second picture est corresponded to the *woman* in the first picture. Subjects who first ated the similarity of the scenes made significantly more relational mapings (i.e., woman to squirrel) (M = 69%) than subjects who did not (M 42%). These findings suggest that the act of carrying out a similarity omparison induced a structural alignment and increased people's likehood of making matches on the basis of shared relational roles over mple object similarities. (See Fig. 11.1.)

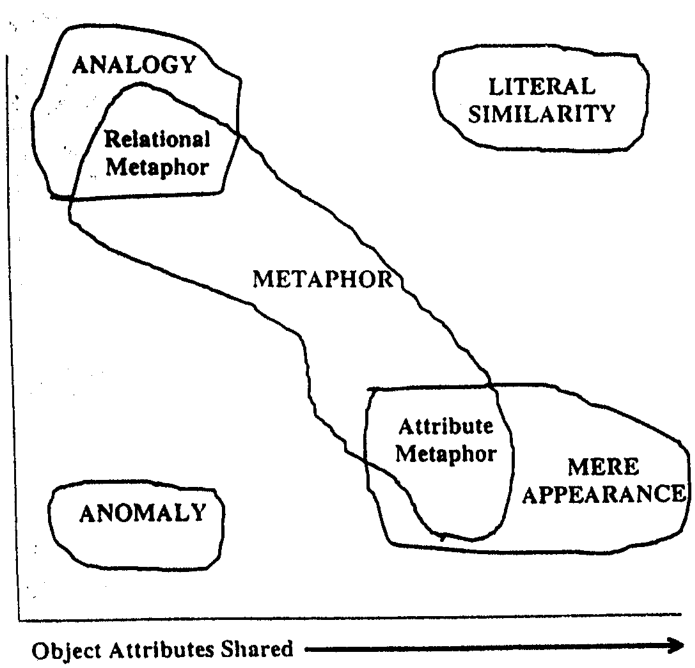

FIG. 11.1. Similarity space.

Metaphor

Now let us apply structure-mapping theory to metaphor. It must be admitted at the outset that metaphor is a protean phenomenon. Metaphors, like analogies, are nonliteral comparisons, but whereas analogies are used for explanatory-predictive purposes, metaphors can also be used in expressive-affective contexts (for longer discussions, see Boyd, 1979; Fauconnier, 1990; Gentner, 1982; Gentner, Falkenhainer & Skorstad, 1988; Steen, 1991). Figure 11.1 illustrates a similarity space showing the continuum from analogy (common relational structure) to literal similarity (common relational structure plus common object descriptions) and from literal similarity to mere-appearance matches (common object descriptions). Metaphor spans the range from attributional comparisons like (1) through relational comparisons like (2):

1. For the black bat, night, has flown (Tennyson).
2. A novel is a mirror carried along a main road (Stendhal).

Indeed, metaphor is even more polymorphous than Fig. 11.1 can portray, for in certain literary contexts metaphors can escape the constraints of structural consistency, as in Example 3.

3. On a star of faith pure as the drifting bread, /As the food and flames of the snow (Dylan Thomas).

This permissiveness follows from the communicative emphasis on capturing affect in metaphor. An unresolvable mapping is irritating in an analogy, because the comparison is responsible for conveying clear inferences based on common structure. It can be pleasing in metaphor, where the sense of rich intermeshing can be part of the experience.

The bulk of the metaphors used in psychology experiments are explanatory-predictive metaphors and are structurally well behaved, if somewhat on the lifeless side. There are occasional attribute comparisons such as:

4. Her hair was spaghetti.

However, relational comparisons (analogies) like (5) form by far the most frequent category:

5. Cigarettes are time bombs.

We will focus our discussion on explanatory-predictive metaphors. These are almost always relational, and most can be analyzed in the same manner as analogies. For example, consider A. E. Housman's metaphorical comparison: "A poet can no more define poetry than a terrier can define a rat." Although not a cross-mapping, this metaphor involves a set of unlikely object correspondences. Clearly, the correspondences between *poet* and *terrier* and between *poetry* and *rat* are not meant to reflect pairwise object similarity. The poet is not seen as a dog, much less poetry as vermin, except with respect to the relation between them—the unthinking avidity of the pursuit.[3]

In this metaphor, the intended meaning conveys a highly specific relational structure. Some metaphors go even further and invite the application of an extended domain mapping. For example, consider Shakespeare's metaphor:

I have ventured,
Like little wanton boys that swim on bladders,
This many summers in a sea of glory;
But far beyond my depth: my high-blown pride
At length broke under me; and now has left me,
Weary and old with service, to the mercy
Of a rude stream, that must forever hide me.

This extended mapping between swimming boys and the arena of political intrigue preserves perfect structural consistency throughout, deepening as the further implications of the parallel are developed. Here the very incongruity of the object-level correspondence between the adventuring boys and the aging and defeated man is part of its effectiveness.

As these examples suggest, much of metaphoric comprehension can be seen as analogical mapping.[4] As we review in the following sections, there is considerable evidence that the process of comparison is sensitive to relational structure (Clement & Gentner, 1991; Gentner & Clement, 1988; Goldstone, Gentner, & Medin, 1989; Markman & Gentner, 1993a, 1993b) and that comprehension of metaphors involves an alignment process (Gentner, Imai, & Boroditsky, 1999; Gentner & Wolff, 1997; Wolff & Gentner, 1992, in press). In the following sections we amplify the structural

[3]Consider the hopelessness of trying to capture this meaning if one were restricted to parameter space representations. By mapping a multidimensional space of animal dimensions onto the multidimensional space or spaces for poets and poetry, one can convey that poets are more fierce than poetry, but not the specific relation of pursuit between them.

[4]We do not attempt to deal here with metaphors that are radical departures from structural consistency. Fauconnier's (1990) research on complex metaphor and blending processes provides a good framework for these extensions.

alignment model and provide more evidence focusing on the issue of how metaphors and analogies can lead to change of knowledge.

MECHANISMS OF CHANGE

Knowledge Selection

In the richness of our representations, knowledge can often get buried. When two concepts are compared, aspects that typically remain unconsidered can be picked out and made the focus of attention. For example, in thinking about televisions, certain properties easily come to mind: that they are medium-sized appliances, that they display pictures, that they have antennae, and so on. Given the metaphor "Television is bubble gum for the mind," however, the shared idea of mindless activity comes to the fore. The power of knowledge selection is perhaps most evident when this *highlighting* mechanism is used to pick out knowledge that is not normally salient. Highlighting or knowledge selection is one of the major ways that analogies and metaphors illuminate their targets (Black, 1962, 1979; Clement & Gentner, 1991; Elio & Anderson, 1981; Hayes-Roth & McDermott, 1978; Indurkhya, 1991; Kuhn, 1979; Way, 1991; Winner, 1988). (See Fig. 11.2.)

A process model of knowledge selection and highlighting must account for the generativity of metaphor comprehension. For example, consider the metaphor: "If we do not plant knowledge when young, it will give us no shade when we are old" (Chesterfield). You were probably able to infer the base domain of a growing tree and the intended image of something that begins small and grows slowly but that, properly cared for, eventually becomes an immense and rewarding presence. Yet you had no way of anticipating this meaning from the foregoing text; it had to emerge from the metaphorical comparison. This example demonstrates that modeling metaphor comprehension as the process of verifying expected patterns (whether derived from external context or from the person's current goal state) is overly restrictive (See Carbonell, 1981, 1982; Holyoak, 1985, for arguments on the other side). Instead, we need a

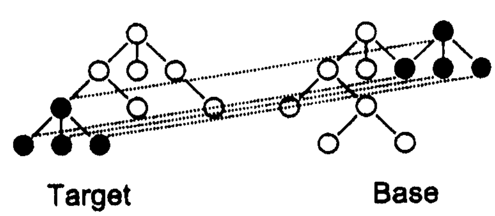

FIG. 11.2. Knowledge selection.

process model that can derive a plausible meaning de novo from the juxtaposition of two representations.

In the structure-mapping engine (SME; Falkenhainer, Forbus, & Gentner, 1989), structural alignment is performed by using a local-to-global algorithm that can arrive at unanticipated matches.[5] Initially, matches are made between all pairs of identical elements.[6] This initial set of local matches is typically inconsistent and many to one. In the next phase, SME imposes the structural constraints of *one-to-one correspondence* and *parallel connectivity* to coalesce the local matches into (typically) several structurally consistent subsystems (kernels), which in turn are joined into one or a few maximal structurally consistent interpretations. SME produces a structural evaluation of each of its possible final interpretations by using an algorithm that favors systems that are large (i.e., have many matching predicates) and *deep* (i.e., with higher-order relational connections, rather than large sets of independent matches).[7] Finally, predicates connected to the common structure in the base, but not initially present in the target, are proposed as *candidate inferences* in the target. As we discuss in the next section, in this way structural alignment and mapping can lead to the projection of unplanned inferences.

The systematicity assumption is crucial to our account of metaphor comprehension. The assumption that people prefer to match predicates belonging to interconnected systems of knowledge rather than isolated independent components has several implications. First, not all matching predicates should enter into an interpretation: only information tied to the maximal common structure. Thus the structural alignment model does not fall prey to the problem that besets simple comparison models, that of predicting inclusion of all matching information (Camac & Glucksberg, 1984; Glucksberg & Keysar, 1990; Goodman, 1970; Rips, 1989; Tourangeau & Rips, 1991; Tourangeau & Sternberg, 1981). Evidence for the selective power of systematicity was found in an experiment by Clement and Gentner (1991). In their study, subjects read two analogous passages. The target and base contained two clearly matched pairs of facts. In one case, the key matching facts were connected to larger causal systems that also

[5]The initial representations must contain between them the relations that are matched, but the process is nontrivial all the same, because the representations can be sufficiently rich that the outcome of a mapping process cannot be determined by simply importing "the" base structure to the target.

[6]This initial search for identities means that SME does not need to solve the (computationally intractable) general graph-matching problem, contrary to Hummel and Holyoak's (1997) claim. Semantic similarity between predicates is captured through a decomposition into partial identities; we return to this later.

[7]See Forbus and Gentner (1989) and Forbus and Oblinger (1990) for the details of how the smaller kernal structures are combined into larger structures and evaluated.

matched between base and target. Specifically, each fact was the consequent of another matching pair of facts. In the other case, the pair of key facts was matched equally well locally, but each was the consequent of a different (nonmatching) antecedent. Thus, the two did not belong to matching systems, even though they matched perfectly well locally. According to the systematicity principle, only facts that are connected to *corresponding relational systems* should get included in the interpretation. Consistent with this prediction, when asked which matching pair contributed more to the analogy, subjects selected the pair that was connected to a matching antecedent. Thus, the feature selection problem is dealt with by using systematicity as a selection constraint.

A second implication of the systematicity assumption is that it predicts that, on the whole, people should have a preference for relational information over attribute information in their interpretations. This result is predicted because relations, to a greater degree than attributes, serve to make knowledge more connected and systematic. Studies by Gentner (1988) and Gentner and Clement (1988) are consistent with this prediction: Adult subjects' interpretations of metaphors were found to contain predominantly relational rather than attributional information. In contrast, their descriptions of the individual concepts that entered into the metaphor contained approximately the same amount of relational and attributional information. Furthermore, subjects' aptness ratings were correlated with the relationality, but not the attributionality, of their interpretations (Gentner, 1988). Tourangeau and Rips (1991) found a similar emphasis on relational commonalties in a metaphor-interpretation task. Their subjects rated the degree to which interpretations (from another group of subjects) specified relations rather than simple attributes. They found that relationality ratings for assertions used in interpretations were higher than for those not included in the interpretation but included in the target and base descriptions.

A final example of metaphoric highlighting comes from an unpublished study by Gentner and Koenig. They asked whether comparison processes would induce an abstract schema that would permit subsequent relational retrieval, instead of surface-based retrieval as is typically found in reminding studies (Gentner, Rattermann, & Forbus, 1993; Gick & Holyoak, 1983; Holyoak & Koh, 1987; Keane, 1988; Ross, 1987). Subjects rated the similarity of pairs of proverbs. The pairs were either relationally similar—"You can't tell a book by its cover" and "All that glitters is not gold"—or object similar—"Don't look a gift horse in the mouth" and "You can lead a horse to water but you can't make it drink." Subjects who rated the similarity of relationally related pairs were much better able to recall the original items when given another relationally similar proverb than were subjects who rated the similarity of dissimilar or surface-similar pairs. In contrast, although surface-based retrieval (i.e., retrieval given a surface-similar cue) was generally high, it was not much improved by rating the similarity of surface-similar pairs. This result suggests that the similarity comparison highlighted the common relational structure.

A second study assessed whether the schemas derived would carry forward. Subjects rated the similarity of pairs of proverbs and afterward wrote out interpretations for new proverbs, some of which were relationally similar to the originals. Subjects who gave similarity ratings for relationally similar pairs wrote abstract interpretations of the new proverbs by using the schema consistent with that embedded in the original pair. In contrast, subjects who rated surface similar pairs and control subjects who merely rated the importance of the original proverbs instead of comparing them tended to write concrete, idiosyncratic interpretations of the new proverbs. These results suggest that the act of comparison led to a highlighting of common structure, resulting in a more abstract representation of the proverbs' meaning.

Knowledge selection is an important aspect of metaphorical insight. It is true that by itself it suffers from the classic problem of traditional comparison models: It is limited to information contained in the initial representations of the terms (Glucksberg & Keysar, 1990; Tourangeau & Rips, 1991; Way, 1991). Nonetheless, if we assume that human knowledge is typically rich and situationally embedded, then the metaphoric selection and extraction of smaller subsystems from the thicket of knowledge serve a valuable function. Furthermore, when alignment and highlighting identify common structure, they provide a basis for processes that add or alter knowledge, such as *projection*.

Projection

Linguists and rhetoricians have often asserted that metaphor involves a transfer of meaning from the base to the target. The Greek ancestor of the term *metaphor* means "to transfer or carry over" (Verbrugge & McCarrell, 1977; Wheelwright, 1962). We refer to this sort of transfer as *projection of candidate inferences*. It is also called *property introduction* (Glucksberg & Keysar, 1990; Ortony, 1979a), or *attribution*.

This process of transferring inferences from one domain to another is well illustrated in the history of scientific discovery (Boyd, 1979; Dreistadt, 1968; Gentner, 1982; Gentner & Jeziorski, 1993; Hesse, 1966; Kuhn, 1979; Nersessian, 1992; Oppenheimer, 1956; Thagard & Holyoak, 1985). The discovery of mesons offers an apt example (Oppenheimer, 1956). In the late 1940s, the Japanese physicist Hideki Yukawa proposed that the electrical and nuclear forces might be analogous (Yukawa, 1982). It was already known that for the electric force, interactions between electrically charged bodies were mediated by electrical fields. Using arguments from relativity and quantum theory, Yukawa speculated that corresponding fields and particles might exist for nuclear forces. These particles—mesons—were eventually found in cosmic rays. In this instance, as in most instances of scientific invention, the projection depended on a previous partial alignment. A comparison and partial alignment led to a further inference.

A similar kind of reasoning may occur in more mundane metaphors. For example, "My surgeon is a butcher" suggests a clumsy, brutal surgeon. Common structure between surgeons and butchers emerges in this juxtaposition: Both cut flesh with specialized implements. The normal manner and purpose of the cutting are quite different for surgeons and butchers, however. In this metaphor, the cutting structures are easily aligned, permitting a transfer of information from the base term (*butcher*) to the target term (*surgeon*). The resulting candidate inference is that the surgeon cuts in the manner of a butcher, crudely and without regard for the well-being of the organism. Thus, in the structure-mapping framework, projection is a structural completion process that follows the initial structural alignment (see Fig. 11.3). As discussed earlier, candidate inferences are formed by mapping across parts of the relational structure of the base that are connected to the base's matching structure, but for which there is not yet corresponding structure in the target. Once such potential inferences are identified, they are brought over from the base and inserted into the target structure, subject to verification of their validity in the target domain.

Projection represents an important way in which metaphor can lead to knowledge change: by the carryover of information from one concept to another. Because the structure-mapping account involves both alignment and projection, it does not suffer from the criticism leveled at pure comparison models, namely that they are inherently incapable of explaining how information found in only one of the terms can enter into an interpretation (Glucksberg & Keysar, 1990; Way, 1991). In addition, projection offers a way of explaining the phenomenon of *metaphorical directionality:* that is, for people's preference for one ordering of the terms over another. Because projection occurs directionally, from the more systematic

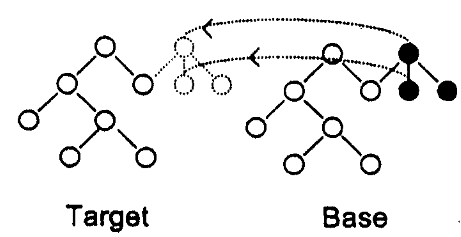

FIG. 11.3. Projection.

of the two aligned structures—which, in a felicitous metaphor, will be the base domain—people implicitly prefer orderings that result in the greatest amount of transfer (Bowdle & Gentner, 1997). Bowdle and Gentner found that asymmetries in the preferred direction of a comparison can be predicted by the degree to which one term possesses more systematic relational structure than the other. People prefer to place the more relationally coherent structure in the base position and are more likely to draw inferences in this direction (Bowdle & Gentner, 1997; Gentner & Bowdle, 1994).

Findings from Clement and Gentner (1991) supported the view that projection is constrained by common relational structure. Subjects read two analogous passages and then made predictions about one of the passages (the target) on the basis of information contained in the other (the base). The base passage contained two facts for which there were no corresponding facts in the target. One of these facts was the consequent of a fact that matched a target fact: That is, it was part of a relational structure in the base that matched relational structure in the target. The other fact was the consequent of a fact that did not match anything in the target. As predicted, subjects' willingness to make a particular inference depended on whether it was tied to the common relational structure.

Evidence for projection has also been found in children. In a study by Gentner (1977), children were asked to say where on a mountain or tree a body part, such as a knee, would be found. Importantly, the place on a tree where a knee might be expected to be was not marked; the task thus required projection on the basis of a partial mapping of the objects. Four-year-old children performed very well. Their high level of performance held even when the trees or mountains were put on their sides or upside down. Chen and Daehler (1989) showed that 6-year-olds can transfer a solution from a story to a physical apparatus, provided that the elements are easily aligned. In another study, Gentner and Toupin (1986) investigated the role of higher-order relations in children's mappings. They asked children 4 to 6 and 8 to 10 years of age to transfer a story plot from one group of characters to another Half the children were given simple sequential plots; the other half (the systematic condition) received the same plots but with added beginning and ending sentences that expressed a causal or moral summary. In addition, the transparency of the mapping was manipulated by varying the similarity of corresponding characters. For both ages, transfer accuracy was nearly perfect with highly similar corresponding characters (e.g., *chipmunk* → *squirrel* and *moose* → *elk*), lower when corresponding characters were quite different (e.g., *chipmunk* → *lion* and *moose* → *trout*), and lower still in the cross-mapped condition in which similar looking characters played different roles (e.g., *chipmunk* → *elk* and *moose* → *squirrel*). For the older group, but not the younger group, systematicity also had strong effects: 9-year-olds could map virtually perfectly even in cross-mappings as long as they had systematic relational structure. However, when given stories that lacked systematicity, they were at the mercy of the object similarities; their accuracy was low in the cross-mapped condition. These results show that processes of alignment and mapping are present from early in development and that they become more fluent and more independent of surface similarity as children acquire higher-order connecting relations to guide their projections.

We have argued that in general alignment precedes and guides projection. However, there are exceptions. As discussed below, highly conventional 'stock' metaphoric bases have stored metaphorical senses: for instance, 'jail' as a confining institution). Further, some metaphors involve a more complex interplay of alignment with other processes of blending (Fauconnier, 1990). For example, consider Alexander Pope's couplet:

> "Satire or sense, alas! can Sporus feel?
> Who breaks a butterfly upon a wheel."

The reader is invited to imagine stretching a butterfly on a rack; the very difficulty of doing so invites the image of one so insubstantial as to be unworthy of torture.

In summary, the mechanism of projection provides a way of importing knowledge that is initially present in the base but not in the target. In this case, change of knowledge occurs in the target. We now consider cases in which change of knowledge applies to *both* domains. In the next section we discuss mechanisms for re-representing initially mismatching predicates to reveal common structure.

Predicate Re-Representation

The evolution of plants and animals has been compared to the growth pattern of a great tree. In *On the Origin of Species,* Darwin extended this metaphor in several interesting ways (Dreistadt, 1968). Just as the competition between spatially close twigs can be especially intense, so can competition between animal species at the same ecological niche. The winning twigs may grow into great branches that spread out and bear other branches and twigs. Likewise, animal species that survive can become the progenitors of other species.

The Great Tree metaphor is grounded in similarities between two large systems of relations. However, the precise way in which competition occurs among animals is markedly different from the way in which it occurs among branches in a tree. Animals often compete by physically fighting for food, territory, and mates, whereas twigs compete in a less dramatic manner, by gaining or losing resources rather than by aiming direct injury at one another. As another example, when it is said that one animal species *comes from* another, the relation is one of genealogy over time, whereas when the same thing is said of twigs, the relation seems to be spatial in addition to (or instead of) temporal. The question is, how do we align such nonidentical structures? Competition differs across animals and plants, but not so much so that the similarities cannot be perceived.

The issue of nonidentical correspondences is an important problem for models based on similarity (Black, 1962, 1979; Hesse, 1966; Miller, 1979; Ortony, 1979a, 1979b; Ortony, Vondruska, Foss, & Jones, 1985; Tourangeau & Sternberg, 1981, 1982; Way, 1991). If meaning components (predicates) that should correspond in order to fit a larger mapping do not match exactly, by what mechanism can their correspondences be known?

In computational models of analogy, several proposals have been made as to how this problem might be solved. The *taxonomic re-representation* approach employs abstraction hierarchies. According to Burstein (1986), when the relations in the base and target differ, a "virtual relation" can be formed in the target that is a sibling or ancestor of the corresponding base relation. In a similar fashion, Winston (1980) suggested that ancestors for predicates like KILL can be found with subroutines that generate predicates like HURT or HAS-CONFLICT-WITH. Falkenhainer (1990) incorporated a notion of *minimal ascension* into his contextual structure-mapping engine (*Phineas*). In cases where two consequents match identically but the antecedents do not, Phineas attempts to match the antecedents by climbing the taxonomic hierarchy until the minimal common ancestor is found: DESTROY and STAB might have the common superordinate of HARM. This approach seems psychologically plausible in cases where a firm taxonomic hierarchy can be assumed. However, this assumption may be unwarranted for verbs (Gentner, 1981; Graesser & Hopkinton, 1987; Huttenlocher & Lui, 1979).

Another way in which nonidentical predicates can be matched is by performing *decompositional re-representation* (Gentner & Rattermann, 1991; Gentner, Rattermann, Markman, & Kotovsky, 1995; Kotovsky & Gentner, 1996). (See Fig. 11.4.) In re-representation, predicates are decomposed into subpredicate structures, much as in lexical decomposition. This allows similarities between predicates to be manifest as identities in (some of the) subcomponents. For example, in a metaphor like "The hotter the anger the sooner quenched," there is an implicit comparison between anger and

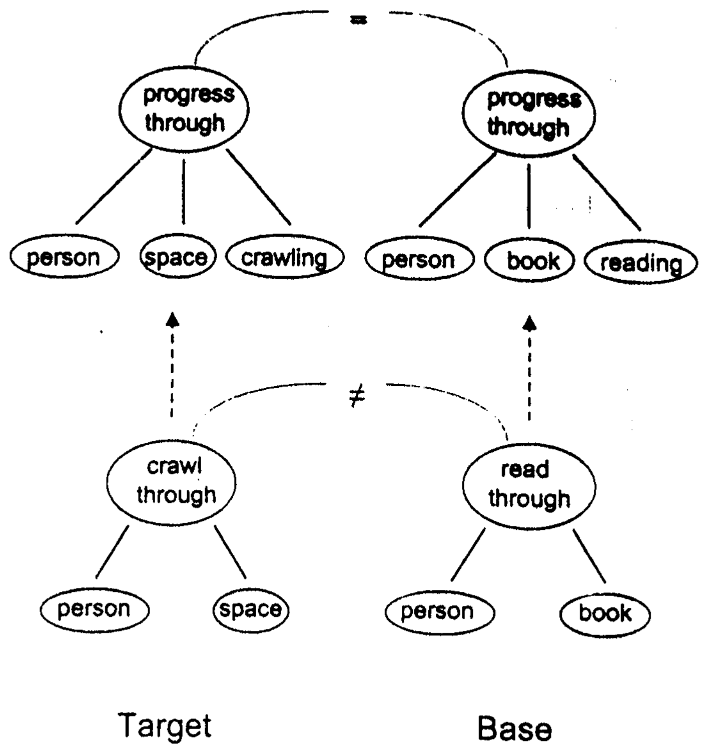

FIG. 11.4. Re-representation. *He crawled his way through the book.*

fire. To align these structures, the initial representation[8] of the assertion HOTTER-THAN (f_1, f_2) can be re-represented to form the equivalent subpredicate structure GREATER-THAN (*temperature* f_1, *temperature* f_2). Similarly, the relational predicate ANGRIER-THAN (s_1, s_2) can be re-represented as GREATER-THAN (*anger* s_1, *anger* s_2). The mapping problem now reduces to placing the nonidentical dimensional functions of *anger* and *temperature* in correspondence.[9] Once the original comparative assertions are re-represented, their similarity becomes apparent: Both involve a notion of comparative magnitude, albeit along different dimensions.

Re-representation provides a means by which matches between nonidentical predicates can be made, but unrestrained re-representation would be computationally expensive and, worse, could lead to profligate matching. The decision to re-represent must be constrained. One way to accomplish this is by limiting re-representation to just those cases in which neighboring predicates already match. This approach to the nonidenticality problem is the one used in Falkenhainer's (1990) Phineas and in Keane

[8]These representations are assumed to be conceptual rather than specifically verbal. That is, we assume a conceptual level of representation that is not modality specific.

[9]SME can match nonidentical functions and entities but requires identity to set up a correspondence between relations and attributes (truth-bearing predicates). Thus re-representing specific relations in terms of abstract relations over specific functions allows us to capture the sense that the same structures hold across different specific domains.

and Brayshaw's (1988) incremental analogy machine (IAM) model of analogical reasoning.

A second issue is whether there is a set of standard processes for re-representation and whether this is influenced by experience. One plausible scenario is that people first check an abstraction hierarchy before matching nonidentical features. If a common abstraction does not exist, as would be the case when a particular pair of nonidentical predicates is encountered for the first time, people may then perform re-representation. If such a re-representation is used repeatedly, it may come to be added to the abstraction hierarchy. Another possibility is that the underlying predicate formats become more available so that in future matches their similarity can be more easily identified. Over the course of many comparisons, this process of format change can have the effect of making mental representations more unified and therefore easier to work with (Gentner & Rattermann, 1991). Such a process may lead to the kind of *representational redescriptions* envisioned by Karmiloff-Smith (1991). She argued that in conceptual development children move through stages of understanding in which representations in one phase are redescribed in the next (e.g., procedural to declarative), with the result that the child's representations become increasingly more flexible and context independent. We suggest that alignment and re-representation mechanisms may contribute to this redescription process.

The process of re-representation is important in cross-domain metaphor, in which people must match descriptions across disparate dimensions. For example, Asch (1955) explored how predicates used to describe physical objects can be used to describe qualities about people. Like physical objects, people can be described as deep or shallow, narrow or wide, hard or soft, bright or dull. In several instances, Asch found that many metaphorical usages had the same meaning across cultures (see also Greenberg, 1966; Lakoff & Johnson, 1980; Osgood, 1949). For instance, the morpheme for "straight" is used to designate honesty, righteousness, and correct understanding. However, in some cases, there is considerable cultural variation: For example, "hot" can stand for a wide range of meanings, including wrath (Hebrew), enthusiasm (Chinese), sexual arousal (Thai), worry (Thai), energy (Hausa), and nervousness (Shilba). In some sense, these are all alike in being positive ends of some generalized energy dimension, but the variety of specific dimensions used suggests that these dimensional relations are at least partly culturally selected or constructed.

How do children learn such systems of dimensional correspondences? Kotovsky and Gentner (1996) asked whether re-representation might contribute to children's learning about cross-dimensional matches. Their study focused on children's ability to perceive similarity solely on the basis of common perceptual higher-order relations such as monotonicity and symmetry. They gave children a forced-choice triads task in which they were shown a standard embodying some relational structure—say, *symmetry* (e.g., XoX)—and asked to say which of two other figures it was most similar to: another instance of symmetry (HiH) or a second figure that lacked the symmetry relation (iHH). One of the choices (HiH) preserved the relational structure of the standard while changing the specific object attributes (e.g., shape). The other choice (iHH) was created by rearranging the components of the relational choice. Thus, both choices were equally dissimilar to the standard in terms of object attributes, but only the relational choice shared the higher-order relational structure of the standard. The key variable was the degree of concrete lower-order similarity between the standard and the relational choice. In some instances, the relational choice had the same dimension of change (either size or darkness) as the standard. In other instances, the dimension of change was different.

When children were given mixed sets of these similarity triads (without feedback), 4-year-olds chose randomly except in the close-similarity case when both the dimension of change (size or shading) and the polarity of change were the same. They could match XoX/HiH but not xOx/HiH. Kotovsky and Gentner then investigated ways of teaching 4-year-olds to perceive the cross-dimensional match. In one study, they attempted to induce re-representation by teaching 4-year-olds names for the higher-order relations: for instance, *even* for symmetry. Learning to label and sort cards according to these higher-order labels improved children's performance on subsequent cross-dimensional similarity matches. Perhaps more surprisingly, simple juxtaposition of several concrete "easy" matches also seemed to induce re-representation. When 4-year-olds received a set of in-dimension pairs first, followed by a set of cross-dimensional pairs, they performed much better on the cross-dimensional pairs than children who received the same set in mixed order. A similar benefit occurred for children who received blocks of same-polarity matches—such as xOx/iVi—before blocks of opposite-polarity matches—such as xOx/IvI. Kotovsky and Gentner interpreted these findings as indicating that initial concrete matches between the standard and relational choice helped children form more abstract representations of monotonic increase and decrease (cf. Gentner et al., 1995). This finding that close literal similarity matches facilitate subsequent analogical matches that embody the same relational structure appears to be quite general in learning and development (Gentner & Medina, 1998).

Why should close similarity matches facilitate subsequent abstract matches involving the same relational structure? Gentner and Kotovsky suggested a mechanism of *progressive alignment*. Four specific assumptions were made. First, children initially represent the relations in a dimension-specific manner (e.g., *darker than* and *bigger than*). That is, for young children, the representation of a difference in magnitude is bound up with the dimension of difference. Second, close matches are easy to perceive—in a sense, they are automatically aligned. Third, alignment results in a slight highlighting of the common relational structure. After repeated such alignments, the higher-order relational structure—such as *symmetry* or *monotonic change*—is strong enough so that a partial match can be made even in a cross-dimensional pairing. Fourth, this partial match invites a re-representation: a decomposition that further brings out and clarifies the likeness. In this case, a re-represention that separates the *greater-than* relation from the specific dimension—for instance, *greater* [*darkness* (a), *darkness* (b)])—will reveal an identical system of relations expressed across different dimensions. The common higher-order symmetry pattern can then be perceived.

Research by Smith and Sera (1992) provides another possible example of re-representation over the course of learning and development. In their experiments, children were asked to say "which is *more*" for pairs of mice of different size, loudness, or darkness. Even 2-year-olds showed a consistent mapping between *more* and *bigger*. However, the mapping between *more* and *louder* was not firmly present at 2 years and became more consistent with age and experience. The mapping between magnitude and darkness showed yet another pattern. Two-year-olds had a consistent mapping from *more* to *darker* and from *bigger* to *darker*: Big mice were paired with dark mice and little mice with light mice. However, 4-year-olds were random on the mapping from size to darkness. Adults were split: Half assigned *large* to *dark*; the other half, *large* to *light*. Smith and Sera speculated that the shift from consistent to random to split mappings between size and darkness may be related to the onset of words for the darkness dimension. The English language is ambiguous as to the polarity of the *light/dark* dimension. Both *light* and *dark* can occupy the positive pole in different contexts. Perhaps an initial consistent perceptual mapping is set aside as children try to integrate language and perception.

A study by Vosniadou, Ortony, Reynolds, and Wilson (1984) examined re-representation developmentally. In one of their experiments, children listened to seven short stories that ended with a concluding sentence that was either literal or metaphoric. A metaphoric ending was something like "Sally was a bird flying to her nest"; a literal counterpart to this was "Sally was a girl running to her home." Comprehension was measured by children's ability to act out the stories with toys. First graders tended to act out the test sentence literally. For example, in demonstrating how "Sally flew to her nest," they moved the doll through the air. Third graders, on the other hand, reinterpreted (re-represented) the verb to mean "run" and thus had the doll move quickly across the ground. These children apparently could re-represent two events to reveal common structure. This re-representation potential contributes to making nonliteral comparison a potent developmental force. However, the results also show that re-representation is not automatic. Juxtaposition does not guarantee an illuminating alignment.

One result of re-representation can be an increase in the uniformity with which the relations are encoded. We hypothesize that such representational uniformity can increase the likelihood of relational remindings (Forbus, Gentner, & Law, 1995). Clement, Mawby, and Giles (1994) explored the effect of relational predicate similarity on analogical access and mapping between disparate domains. They used materials in which similar relations were expressed either in terms of the same predicates (so that the likeness was manifest) or else in terms of merely similar predicates (so that the likeness was latent, requiring re-representation to be aligned). Analogical access (that is, being reminded of a past situation given a current analogous situation) was better with manifest similarity. In contrast, analogical mapping (that is, the alignment and projection between two current situations) was relatively unaffected by the latent-manifest distinction. In analogical reminding, with only the current situation in working memory, success depends on the degree of match of the pre-existing representations, whereas during mapping, with both situations present in working memory, there is opportunity for re-representation.

This is a good point to step back and consider the issue of identicality in matching predicates. The structural alignment process we propose has a tiered identicality requirement. Three distinct patterns occur vis à vis identity. First, some conceptual identities are necessary to begin the alignment process. Second, these initial identities license other nonidentical correspondences: In particular, identical relations license matches between nonidentical functions and identities. Third, given a partial alignment, the process of re-representation finds new identities. From this perspective, identical relations are as much an *output* constraint as an input requirement in the structure-mapping model. Processes of re-representation can be modeled by adopting processes like those discussed earlier, capable of initiating and constraining re-representational activity. However, re-representation is still relatively unexplored in computational models.

Restructuring

Knowledge change can occur not only at the level of individual concepts, but also at the level of systems. The notion that analogy and metaphor prompt change at the system level has a long history in cognitive science (Black, 1979; Gentner, 1982; Gentner et al., 1997; Indurkhya, 1991; Nersessian, 1992; Schön, 1979; Verbrugge & McCarrell, 1977). In the following sections, we discuss how this sort of change may occur. To do so, we explore in some depth a historical example of restructuring by use of metaphor.

At the turn of the century, physicists were trying to conceptualize the atom (Wilson, 1983). One early proposal, made by J. J. Thomson, was that the atom was a sphere of positive electricity in which negatively charged electrons were stuck like plums in a pudding. The plum-pudding model was supported by a number of major scientists. Soon, however, the model was challenged by Thomson's own student, Earnest Rutherford. Rutherford's atom was a serious departure from Thomson's. How did this departure come about?

In the first decade of the century, Rutherford was working on determining the characteristics of alpha particles (helium atoms stripped of their electrons: i.e., two protons and two neutrons). Rutherford and Hans Geiger (of the Geiger counter), then Rutherford's lab manager, used a technique of aiming a beam of alpha particles at a thin piece of metal foil. By placing a zinc-sulfide screen behind the foil, the researchers could count the number of alpha particles passing through at various angles. Rutherford recognized that these experiments could bear on the structure of the atom. A discovery occurred in 1909 when Rutherford looked for scattering of alpha particles at large angles of deflection. The results were extraordinary. A small but significant number of screen marks indicated that a few of the alpha particles had been scattered backward, deflected through an angle of more than 90 degrees. As Rutherford noted: "It was as though you had fired a fifteen-inch shell at a piece of tissue paper and it had bounced back and hit you." The number of large-angle deflections was far greater than predicted from the simple accumulation of many small deflections.

As Wilson (1983) described it, for the next 18 months Rutherford pondered the finding. Thomson's atom was clearly incapable of predicting the severe change in direction. Sometime in December 1910, Rutherford appears to have settled on a model of the atom (Wilson, 1983) as containing a massively charged, very minute center (later to be called the nucleus) that was surrounded, at relatively great distances, by even smaller electrons distributed throughout a sphere. What led Rutherford to this model? A visit made by the Japanese physicist Nagaoka, which preceded Rutherford's announcement by several weeks, may have had an effect on his thinking. In 1904, Nagaoka had proposed a "Saturnian" model of the atom, a disk-shaped atom with a large, heavy center surrounded by rings of electrons. Another intriguing possibility is hinted at by the fact that in Rutherford's copy of Newton's *Principia* there are copious notes in the margins next to the sections on the inverse square law and the hyperbola (Wilson, 1983). These are sections where Newton described the mathematics underlying the paths of comets and other heavenly bodies around central masses like the sun. As has been suggested by several biographers (e.g., Andrade, 1964; Eve, 1939; Wilson, 1983), Rutherford probably recognized that the path of the alpha particle, being hyperbolic, was like the path of a comet. Strengthening the connection was the observation made by Rutherford and his colleagues that the severity of the alpha particles deflection increased with the atomic weight (and therefore the charge) of the metal in the foil, as would be expected by analogy with a comet being defected from a star of varying mass. Rutherford may have pursued the parallels and asked whether the path of an alpha particle is caused by a central force, like the comet's path with respect to the sun. The central force idea then perhaps invited a further mapping of electrons to planets. Electrons in the atom might be distributed around some central attractor like planets around the sun. The relative distances between the nucleus and the electrons might be large, like the relative distances between the sun and the planets. In a letter to Nagaoka in 1911, Rutherford wrote that the alpha particles must be considered as passing "right through the atomic system" and noted the similarities (and some differences) between his spherical atom and Nagaoka's Saturnian disk.

The shift from Thomson's plum-pudding model to Rutherford's solar system model is an example of restructuring via a series of implicit or explicit analogical comparisons. It resulted in a fundamental rearrangement of already known elements. According to Thomson, negatively charged electrons were surrounded by a sphere of positively charged electricity. According to Rutherford, it was nearly the other way around; the element being surrounded was now a positively charged nucleus. Besides this spatial reversal, there was also a fundamental change in how the positive charge was carried; rather than being distributed throughout a sphere, the positive charge was now carried by an entity with mass, a nucleus (see Fig. 11.5).

On the analysis, Rutherford's restructuring began with an alignment between the path of an alpha particle and the path of a comet. Given this local alignment, the structural mapping begins to take on a life of its own. Predicate structures not present in the target but connected to the matching predicate structure in the base are projected as a candidate inference (here, a comet's relation to the sun). New elements (a central object) or relations (attraction or repulsion between the central object and orbiting elements) may be postulated to exist in the target system. These new structures may be incompatible with pre-existing structures; a charged central object cannot be added to the plum-pudding atom without disturbing the balance between the sphere of positive electricity and the negative electrons. Conflicts may be resolved in a number of ways; structures may be eliminated from the target system (e.g., the distribution of

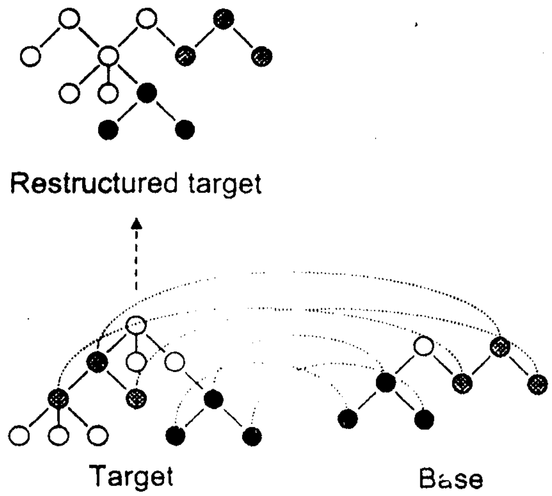

The plum-pudding atom is like the solar-system atom.

FIG. 11.5. Restructuring.

positive charge), or new structures may be induced from other sources of knowledge (e.g., incorporation of orbits as in the Bohr atom). Either way, the overall effect is a rearrangement of the system's basic elements. Rutherford's discovery (as constructed by Wilson) is a classic example of restructuring. It goes beyond candidate inferences and local re-representation and involves reorganization and revision of the previous representational structure.

A Little Restructuring Is a Dangerous Thing. Apart from their role in scientific creativity, we may ask whether metaphoric comparisons influence ordinary learning and reasoning about science. Gentner and Gentner (1983) carried out empirical studies to test whether people's conceptual inferences in a problem domain follow predictably from their metaphors for the domain. The domain they considered was electricity. One common metaphor for this domain compares electricity to water flow: electricity is water, wires are pipes, batteries are reservoirs, current is flowing water, voltage is pressure, and resistors are narrow constrictions in pipes. Another common metaphor involves comparing electricity to a crowd moving through a long hall. Here the correspondences are that electricity is a crowd, wires are paths, current is the number of entities that pass a point per unit time, voltage is how forcefully they push each other along, and resistors are narrow gates. Each of these models has its strengths and weaknesses. The "electricity is water" metaphor captures well how batteries combine to affect voltage. This is because the difference between serial and parallel reservoirs can be understood in terms of height of fluid, a relatively accessible distinction that can then map into electricity. However, this metaphor does not capture resistance well; people do not seem to reason fluently about how serial and parallel narrow constrictions affect flow rate in a water system. The moving crowd metaphor captures resistance much more naturally: In the moving crowd metaphor, resistance corresponds to gates, and people find it easy to simulate the movement of a crowd through various gate configurations.

Gentner and Gentner asked subjects which metaphor they used for electricity and then tested subjects for their ability to solve various circuit problems. As predicted, subjects using the water metaphor performed better on battery problems than on resistor problems, and subjects using the moving-crowd metaphor showed the reverse ordering.

Encouraged by these results, Gentner and Gentner asked whether teaching a new model could help people remedy deficiencies in their existing mental models (unpublished analysis). Subjects were first assessed as to their initial mental models of electricity. Then they were taught to use a particular metaphor, with half receiving the metaphor consistent with their view and half the other (inconsistent) metaphor. Thus, for subjects whose initial model was to view electricity as a teeming crowd, half were told more about these correspondences, and half were taught the metaphor that views electricity as water. The prediction was that this new metaphor would help subjects to solve problems that were hard to solve using their existing model. For example, we expected former *crowd* modelers to improve on battery problems when given the *water* model. However, subjects instead performed *worse* when asked to switch models than did the subjects in the consistent conditions. One interpretation is that the new model led to partial restructuring, leaving the learner with no consistent framework. Possibly matters would have improved had retraining continued over a long period; in this case, it would have been interesting to know whether subjects who learned the new models well showed a decline in performance on the problems supported by their original models. Learning in complex domains may be especially vulnerable to these sorts of transition costs (Spiro, Feltovich, Coulson, & Anderson, 1989; Wiser, 1986).

PRODUCTS OF ALIGNMENT AND MAPPING

In the previous sections, we have reviewed four kinds of change that can result from structural alignment and mapping: knowledge selection, projection, re-representation, and restructuring. Sometimes these changes in knowledge representation are temporary. They serve to get us through the moment—to afford a fleetingly interesting perspective. Sometimes the changes are lasting. Metaphor use may result in a new category or in a conventionalized extended mapping between two conceptual systems.

Metaphors Can Create New Categories

A classic example of metaphoric category creation occurred in the formulation of the scientific notion of *wave* (Hesse, 1966; Oppenheimer, 1956). Initially, the concept was based on the regular, rhythmic movements of water. More generally, it was recognized that these movements had regular relations: When two waves collide, they can either reinforce each other (constructive interference), or cancel each other out (destructive interference). Waves can be dispersed by going through an orifice (diffraction). These properties were found to be true with sound as for water waves. In extending beyond liquids to sound, however, wave phenomena were extended to encompass air as a possible medium through which waves could travel. Once abstracted this far, the concept of a wave was available for further extensions. Light was found to possess the abstract commonalities of constructive and destructive interference and diffraction as well as a major difference: Propagation was possible (indeed, more efficient) in a vacuum. (Like subatomic particles, to which the wave notion was also extended, light is not completely subsumed under the wave rubric; some of its properties are best explained in terms of quanta.) As the extension of the *wave* category broadened, its intension became more abstract.

Glucksberg and Keysar (1990) suggested that the connection between metaphors and categories is fundamental to the nature of metaphor. They suggested that categories are invoked or created in the comprehension of metaphor. Kennedy (1990) and Shen (1992) made similar arguments. As Glucksberg and Keysar pointed out, this view is consonant with the fact that metaphors can be phrased as class-inclusion statements—"Encyclopedias are gold mines"—but literal comparisons cannot—(*) "Encyclopedias are dictionaries." It also offers an explanation for why many metaphors, like class-inclusion statements, are highly asymmetric. Just as the class-inclusion statement "Leeches are parasites" cannot be turned around to make "Parasites are leeches," metaphors like "Suburbs are parasites" resist being turned around to form "Parasites are suburbs."

The claim that metaphors and categories are intimately related can be taken in different ways. The first is that metaphors may be used to *create categories*. The second is that metaphors are *processed by applying categories as opposed to through comparison*. We agree with the first claim, but not the second.

The first view, that metaphor can create categories, has been persuasively argued by Glucksberg and Keysar (1990, 1993) and Shen (1992), among others. Examples such as "He's a real Caligula" demonstrate that even an individual's description can give rise to a category via metaphor. Glucksberg and Keysar noted that metaphorical categories can become conventionalized into ordinary categories. For example, in American Sign Language the concept *furniture* is conveyed by a set of specific instances: bed, chair, and so on. As with the concept of *wave*, the new concept is often an abstraction or extension of the normal meaning of the base term of the metaphor. For example, one sense of the term *sanctuary* refers to a religious edifice; the other, to any location of safety. This second sense may be an extension of "a holy place where one is safe from persecution" to "any place of safety." Similarly, Miller (1993) suggested that "leg of a table" was once understood as a metaphorical comparison between the support of a table and the leg of an animal, but that "leg" has since acquired a secondary meaning. Category creation may also come about through use of a compound noun (Levi, 1978): For instance, "soldier ant" formed from the comparison of certain ants to soldiers. (See Wisniewski [1996] and Wisniewski & Gentner [1991] for evidence of structural alignment in noun–noun compounding.) Such concept creation from metaphors may extend to terms that do not normally strike us as metaphorically derived: for example, "antidote," "bait," "cannibal," "home," "parasite," "scavenger," "shield," and "trap." At one time, such uses as "Her sarcasm was a shield" might have seemed metaphoric. Now, however, these relational terms can apply widely across ontological boundaries. To the extent that metaphors provide mechanisms for concept extension, they offer a means of explaining not only the formation of metaphorical categories but also some aspects of polysemy in word meaning (Bowdle & Gentner, 1999; Gentner & Wolff, 1997; Lehrer, 1990; Miller, 1993).

We next turn to the second possible claim, that metaphors are understood not through comparison but through accessing (or creating) a category associated with the base. Comprehension through categories instead of comparison is an intriguing idea, but we suspect its application is limited to cases where a metaphorical category already exists. For example, if we hear "My boss is a pig," it seems likely that comprehension can proceed by inheritance from this conventional category; we do not have to infer the metaphor's meaning by comparing *pigs* and *bosses* anew (Bowdle & Gentner, 1995, 1999a, b; Gentner & Wolff, 1997; Glucksberg, Gildea, & Bookin, 1982; Wolff & Gentner, 1990, in press). When the metaphor to be understood is novel, however, category models face a critical selection problem: How is the correct category created? According to Glucksberg and Keysar (1990), an ad hoc category is formed of which the base is the prototypical member, and this category is applied to the target. Such a category must reflect both terms of the metaphor (Medin, Goldstone, & Gentner, 1993). We clearly do not derive the same interpretation (or apply the same category) for "My surgeon is a butcher" and "Ghenghis Khan is a butcher." Glucksberg, McGlone, and Manfredi (1997) noted this point and acknowledged that there must be some influence of the target as well as the base. We suggest that the easiest and most natural way to model the joint influence of the two concepts is by assuming a process of comparison via alignment and mapping (see Fig. 11.6). In other words, we suggest that for novel metaphors, the common structure—which may eventually become a category—*arises from* the comparison.

We have carried out several studies aimed at revealing the processing mechanisms for metaphor. The results suggest that novel metaphors are processed through structural alignment rather than through accessing a category associated with the base (Bowdle & Gentner, 1995, 1999a, b; Gentner & Wolff, 1997; Wolff & Gentner, 1992, in press). Our method was to prime metaphors with either the base or the target and then to ask which most facilitated metaphor comprehension. We reasoned as follows. If metaphors are interpreted in terms of base-derived categories rather than by comparison, then processing should begin with the base. For instance, given a base term like *jail*, a category like "situations that oppress" may be derived. Once a category is established, an interpretation of the metaphor would proceed by the application of this category to the target, *job*. If this is the underlying process, then we should see a base advantage. People should be faster to interpret a metaphor when it is preceded by the base rather than by the target.

In contrast, according to the structural alignment model, processing begins with alignment; the directional projection of features from the base term to the target occurs later in processing. Because this initial matching process requires simultaneous access to both the terms, there should be no base advantage. Thus, the structure-mapping model predicts that preceding a metaphor with its base should be no more facilitative than preceding a metaphor with its target.

Across a series of experiments we failed to find evidence for a base advantage. Seeing the base term in advance was no more advantageous than seeing the target in advance. However, seeing both terms consistently

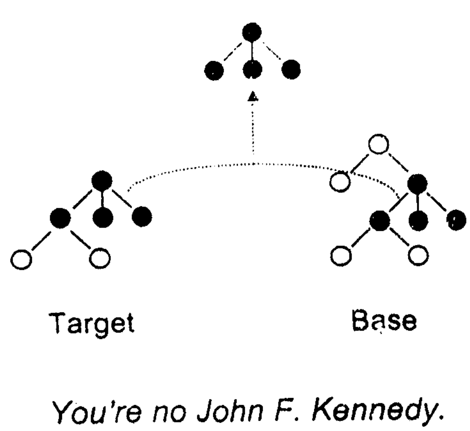

FIG. 11.6. Category creation.

resulted in faster metaphor comprehension than did seeing either of the terms alone.[10] Additional support for the structural alignment model comes from the fact that metaphors rated high in relational similarity were processed faster than metaphors rated low in relational similarity. These findings suggest that metaphor comprehension occurs by means of comparison and alignment.

However, there was one interesting exception: We found some evidence for a base advantage when highly conventional metaphoric bases were used (provided that the relational similarity between the terms was low). Further studies have confirmed this shift from novel to conventional metaphors (Bowdle & Gentner, 1999; in press; Wolff & Gentner, 1999). This suggests that the normal interpretation process for metaphors, as for analogies, is alignment and mapping. But when a metaphoric meaning becomes highly conventional, interpretation may proceed by directly accessing this stored meaning.

Conventionalizing Metaphoric Meanings. When two representations are aligned, common relational structure can emerge. When the same category is derived repeatedly from a given metaphoric base, it may come to be stored along with the base term. In this way, a conventionalized or "stock" metaphor may develop through schema abstraction (Brown, Kane, & Echols, 1986; Elio & Anderson, 1981; Forbus & Gentner, 1986; Gick & Holyoak, 1980, 1983; Hayes-Roth & McDermott, 1978; Medin & Ross, 1989). Processing can then take place by accessing the category abstraction rather than by comparing the normal full meanings of the terms.

These considerations lead to the suggestion of a "career of metaphor" (Bowdle & Gentner, 1995, 1999; in press; Gentner & Wolff, 1997; Wolff & Gentner, 1999): If conventionalization results in stored schema, then processing may change as the metaphor base becomes more associated with a conventionalized concept. As mentioned earlier, Wolff and Gentner found that processing could begin with the base if the base was associated with a conventionalized meaning. That is, the process of understanding a metaphor appeared to be one of structural alignment between the two literal meanings except when base conventionality was high, in which case the category associated with the base was accessed early in the process.

Other research is consistent with the claim that conventionalization results in a shift in processing from online active interpretation to retrieval of stored meanings (Cacciari, 1993; Cacciari & Tabossi, 1988; Clark &

[10]This result is strongly predicted by the structure-mapping model, but not by an initial pure category-projection model, in which any early advantage should reside in accessing the base term's category. However, Glucksberg et al.'s (1997) extension of the class-inclusion model, the *attributive category* model, postulates initial (differential) processing of *both* base and target, and thus can predict the obtained priming results.

Lucy, 1975; Gentner & Wolff, 1997; Gibbs, 1979, 1980; Hoffman & Kemper, 1987; Wolff & Gentner, 1992). For example, Blank (1988) found evidence that conventional metaphors are processed faster than those for which the structural alignments must be made online. He found that metaphorical targets from highly familiar conceptual families (*Time is money*) were responded to as quickly as literal controls. When the metaphorical mapping was less familiar (e.g., *Love is a sickness*), responses were slower to metaphorical targets than to literal controls. Likewise, Blasko and Connine (1993) found that subjects responded to metaphorically related targets from familiar metaphors as quickly as they responded to literally related targets. However, when the metaphors were low in familiarity, subjects responded more slowly to metaphorically related targets than to literally related targets.

The effect of conventionality on processing has been studied with other kinds of tropes, with similar results (Cacciari, 1993; Hoffman & Kemper, 1987). Studies by Clark and Lucy (1975) and Gibbs (1979) found that people were faster to verify pictures and paragraphs consistent with the nonliteral meaning of an indirect request than with the literal meaning. Studying idioms, Gibbs (1980) showed subjects paragraphs that could induce an idiom's literal or nonliteral meaning. The task was to say whether a paraphrase immediately following an idiom was consistent with the meaning of the idiom. Subjects were much faster to confirm conventional nonliteral uses of the idioms than literal ones. These studies suggest that conventional meanings are accessed and comprehended quickly[11] (see also Cacciari & Tabossi, 1988; Swinney & Cutler, 1979). These results are consistent with the claim that metaphor comprehension leads to the creation of conventionalized interpretations.

Metaphors: The Creation of Domains. Some writers (e.g., Boyd, 1979) have argued that when disciplines (e.g., physics) mature, their reliance on metaphors decreases. Others have argued that metaphors remain important (e.g. Campbell, 1920; Kuhn, 1979; Oppenheimer, 1956). Gentner and Grudin (1985) found support for the latter possibility. They examined metaphors for the mind that appeared across nearly a century (1891–1981) in the journal *Psychological Review*. Two interesting findings emerged. In terms of frequency, the number of metaphors used across the decades

[11]However, alignment processes may occur even in this case. In Gentner and Wolff's studies, a base advantage was found only when (a) the conventionality the base's metaphorical meaning was high and (b) the relational similarity between the target and base was low. When the terms were of high relational similarity, processing seemed to proceed in terms of comparison, even if base conventionality was high. Gentner and Wolff suggested a race model between direct comparison of the literal meaning and access to the stock meaning of the base. Wolff and Gentner (1999) also found evidence for initial alignment even for conventional metaphors.

was U shaped, with the bottom of the U occurring roughly when behaviorism was dominant. The use of metaphors for mental processes (like all discussion of mental processes) dwindled to a trickle. The other finding was that whereas the number of metaphors was approximately the same at the beginning and ending periods, the kinds of metaphors used were different. In the early decades, *animate* and *spatial* metaphors were by far the most usual. In the later decades, spatial metaphors continued to be frequent, but system metaphors (especially computer metaphors) became the most frequent class, reflecting the greater attention to mechanisms in current theories.

A related question is how metaphors fare across the development of expertise in an individual. Cooke and Bartha (1992) found that metaphor use actually increased with expertise. Participants in their study were asked to explain the results of several hypothetical experiments. The ratio of psychological metaphors to the total number of ideas was higher for subjects experienced in psychology than for inexperienced subjects. When compared with respect to everyday metaphors, the two groups were the same. Both these results suggest that even though any particular metaphor may become "bleached" and conventionalized, overall metaphor use may not diminish over the development of knowledge.

An interesting instance of metaphorically created categories occurs in our own field. Many technical terms in psychology have at least partial metaphorical status. Gentner (1982) noted that we have "reverberating circuits," "mental distance," "perceptual defense," "memory capacity," "mental image," and "depth of processing." In their historical survey, Gentner and Grudin (1985) noted the usage of such terms as "associative force" (Woodworth), "goal gradient" (Dennis), and "ego defenses" (Minard). As these examples suggest, the elements of a metaphoric mapping are often interconnected in ways that reveal domain relations. As we noted before, analogical mapping is implicitly oriented toward connected systems rather than toward isolated matches. The systematicity bias is a means by which structure-mapping processes bridge the gap between concepts and theories. If we think of theories as the domain structures in which the local concepts are embedded, then analogy and metaphor on this account are ideally suited to transfer such theories.

Consistent with this emphasis, we often find metaphorical categories forming an interrelated system. For example, one common conceptualization of the mind is that it is a physical space (Roediger, 1980). Two important correspondences in this metaphor are that memories are objects in this space and that recall involves spatial search. Metaphors like this derive their force not from a local mapping between the base concept of physical objects to the target concept of memories but rather from mapping the *system* of spatial relations in which these objects are embedded. As cognitive theories evolved (and as technological advances created a greater set of potential bases), Gentner and Grudin's historical trace shows a shift from such general spatial metaphors to more complex systems metaphors that yield more specific inferences about mental processes: telephone switchboard metaphors, circuitry metaphors, and most prominent of all, computer metaphors. Boyd (1979) identified a number of terms formed from the "mind is a computer" metaphor, including "information processing," "encoding," "decoding," "indexing," "feedback," and "memory stores." We now turn to systems of metaphoric mappings.

Global Domain Mappings

So far we have focused on metaphor as creating new categories. A more profound way that analogies and metaphors can lead to change in knowledge is through extended systems of mappings. Lakoff and his colleagues have suggested that a rich set of stored metaphoric mappings pervades our language (Gibbs, 1992; Lakoff, 1988, 1990; Lakoff & Johnson, 1980; Lakoff & Turner, 1989; Turner, 1988). They argue that many everyday expressions imply metaphorical parallels between abstract structures and structures grounded in our experience with the physical world. By correspondence with spatial orientation, for example, we can explain some of the meaning behind expressions like "She fell into disgrace" and "He rose in prominence" (Nagy, 1974). Other metaphorical mappings appeal to culturally grounded systems such as objects in motion, possession, and growth of plants, such as "He stole my idea" (*ideas are objects*), "I can't swallow your proposal" (*ideas are food*), or "The seeds of her great ideas were planted in her youth" (*ideas are plants*). The striking thing about these metaphors is that they involve large, coherent systems of related mappings. For example, Lakoff and Johnson (1980, pp. 90, 91) list examples of the *An argument is a journey* metaphor:

- *We have set out to prove that bats are birds.*
- *So far, we have seen that no current theories will work.*
- *We will proceed in a step-by-step fashion.*
- *This observation points the way to an elegant solution.*
- *He strayed from the line of argument.*
- *Do you follow my argument?*
- *I am lost.*
- *We have covered a lot of ground in our argument.*
- *You are getting off the subject.*

Lakoff and his colleagues have made a persuasive case for the importance of these conventional metaphors. However, two key issues are as yet unresolved: First, what is their psychological status, and second, how do they arise? We begin with the second question and then turn to the first.

How do metaphoric systems come into being? Clearly, no process that operates simply by finding local commonalities or common categories between pairs of concepts compared is adequate, for the essence of these mappings is in large-scale mappings between entire systems of concepts. We propose that these metaphors are processed initially as structural alignments (Murphy, 1996). The structural alignment model explains large-scale domain mappings in terms of correspondences between structured systems. The mechanisms of highlighting, projection, re-representation, and restructuring can be applied to large-scale metaphors.

We now come to the issue of the psychological status of metaphoric systems. Lakoff claimed that conceptual metaphors—including highly conventional metaphors—are psychologically real, enduring mappings that are used in thinking and argued for the invariance hypothesis: that schemas from the base domain are imported into the target. Both claims have been sharply challenged in recent times. Glucksberg, Brown, and McGlone (1993) offered evidence that people do not access the *anger is heat* metaphor when processing conventional metaphoric phrases like "lose one's cool." Murphy (1996) argued that conceptual metaphors are better accounted for as structural alignments between semantically parallel domains than as projective mappings from a base domain, which create meaning in the target domain.

Keysar and Bly (1995) convincingly showed that the seeming semantic transparency of a metaphoric or idiomatic system may be illusory. They asked English speakers to interpret "dead" English idioms such as "The goose hangs high" and found that interpretations were strongly driven by contextual information, rather than by the idiom itself. For example, subjects thought "The goose hangs high" conveyed good news in a happy story and bad news in a sad story. Worse, subjects were convinced that the idioms themselves were transparent. They maintained that they would have arrived at the same interpretation had they seen the idiom in isolation and furthermore, that any other English speaker would arrive at the same interpretation. Keysar and Bly's results, and the points raised by Murphy and by Glucksberg et al., make it clear that we cannot assume that metaphoric language necessarily implies a psychologically real domain mapping.

Gentner and Boronat (1992, 1999) conducted an empirical test of the domain-mapping hypothesis (Boronat, 1990; Gentner, et al., in press). If metaphors are processed by structural alignment, then extended metaphorical mappings should be processed fluently as long as they preserve the domain mapping. If an alignment and mapping process is actively guiding comprehension of "live" metaphors in discourse, people should be more fluent at reading sentences that consistently extend the existing structural alignment than at reading sentences based on a different mapping. For example, after reading Sentence 1, people should read Sentence 2a faster than Sentence 2b:

1. Her anger had been simmering all afternoon.
2a. When Harry got home, she was boiling over.
2b. When Harry got home, she was glacially cool.

Gentner and Boronat (1992, 1999) gave people passages containing extended metaphors such as *anger is a beast*. In four experiments, subjects were timed as they read passages containing these extended metaphors. The last sentence of each passage was always a metaphorical comparison. In some cases, this last metaphor was consistent with the metaphoric mapping underlying the passage. In other cases, the last sentence was based on a different metaphor from the mapping that informed the passage. As predicted, subjects read the last sentence significantly faster when it extended the existing mapping than when it drew on a new metaphoric mapping. This finding supports the basic tenet that metaphors are processed by alignment and mapping and suggests that people find it natural to incrementally extend such mappings. (See Forbus, Ferguson, & Gentner, 1994, for a computational extension of SME that performs such incremental mappings.) These findings are crucial to the claim that large-scale domain metaphors are psychologically real.

However, if conventionalization gradually results in metaphorical meanings coming to be stored with the base term, then we should see a very different pattern for highly conventional metaphors. For extremely conventional metaphors, the abstract metaphorical interpretation may simply be an alternative word sense. In this case, there should be no particular cost for switching metaphors, because comprehension in any case simply involves finding the appropriate word sense. On this account, in the extended metaphor task described earlier, we would predict no advantage in reading time for last sentences whose metaphors are consistent with their passages over those that are inconsistent. Indeed, Gentner and Boronat found that when the passages used highly conventional metaphors (Experiments 1 and 2), subjects were not significantly slowed by a shift in the metaphor.

This finding is convergent with the research on metaphors and idioms (discussed previously) that suggests that conventionalization results in a shift in processing from online active interpretation to retrieval of stored meanings (Blank, 1988; Bowdle & Gentner, 1995, 1999; Cacciari, 1993; Cacciari & Tabossi, 1988; Clark & Lucy , 1975; Gentner & Wolff, 1997; Gibbs, 1979, 1980; Hoffman & Kemper, 1987; Wolff & Gentner, 1992, in press). There is also supporting evidence for the claim that some metaphors are processed as large-scale conceptual systems. Allbritton, McKoon, and Gerrig (1995) found that large-scale conceptual metaphor schemas facilitate recognition judgments for schema-related sentences in text (see also Gibbs, 1990, 1994; Nayak & Gibbs, 1990; but see Glucksberg, Brown, & McGlone, 1993, for contradictory evidence). Further evidence comes from studies of metaphors from space to time (Gentner & Imai, 1992; Gentner, Imai, & Boroditsky, 1999; McGlone & Harding, 1998). This research capitalized on the existence of two English metaphoric *space* → *time* systems: the *ego-moving* metaphor, wherein the observer's context progresses along the timeline toward the future, and the *time-moving* metaphor, wherein time is conceived of as a river or conveyor belt on which events are moving from the future to the past. For example, Gentner and Imai (1992) asked subjects to process statements about time, stated in terms of spatial metaphors, such as "Joe's birthday is approaching" (*time-moving*) or "We are approaching the holidays" (*ego-moving*). As in the Gentner and Boronat studies, people's processing of the metaphors was slowed by a shift from one space–time metaphor to the other.

Metaphoric Mappings in the Creation of a Domain. We noted earlier that analogical mappings can be formative in the development of a new domain, as in the Rutherford atom example. The metaphors used in cognitive psychology vary in their systematicity. Pylyshyn (1989) contrasted coherent system-mapping metaphors, such as the computer metaphor for cognition, with other metaphors that lack systematicity, offering as examples of the latter Freud's hydraulic metaphor for the unconscious and the "mind's eye" metaphor of visual imagery. Such metaphors, he argued, are deceptive; they provide a spurious sense of comfort but not a clear theory. Gentner (1982) made a related point about Freud's theory of anal eroticism, with its claims that feces, money, gift, baby, and penis are often unconsciously treated as equivalent in dreams. This kind of many-to-many mapping places few constraints on the interpretations possible for a given dream.

The trouble with such metaphoric collections in theory building is that, lacking a systematic base and a structurally consistent mapping, they cannot support clear candidate inferences. Figure 11.7, adopted from Gentner (1982) and Markman and Gentner (1999), shows the inferential indeterminacy that arises from $n - 1$ correspondences (a particularly clear case of structural inconsistency).

It must be noted that although we have distinguished local metaphors from systematic metaphorical mappings, there is a continuum between them. Some cases seem purely local: "He's a real pig" seems relatively unconnected to any larger system (indeed, it is not even self-consistent, in that this same metaphor can mean that someone is a male chauvinist, that

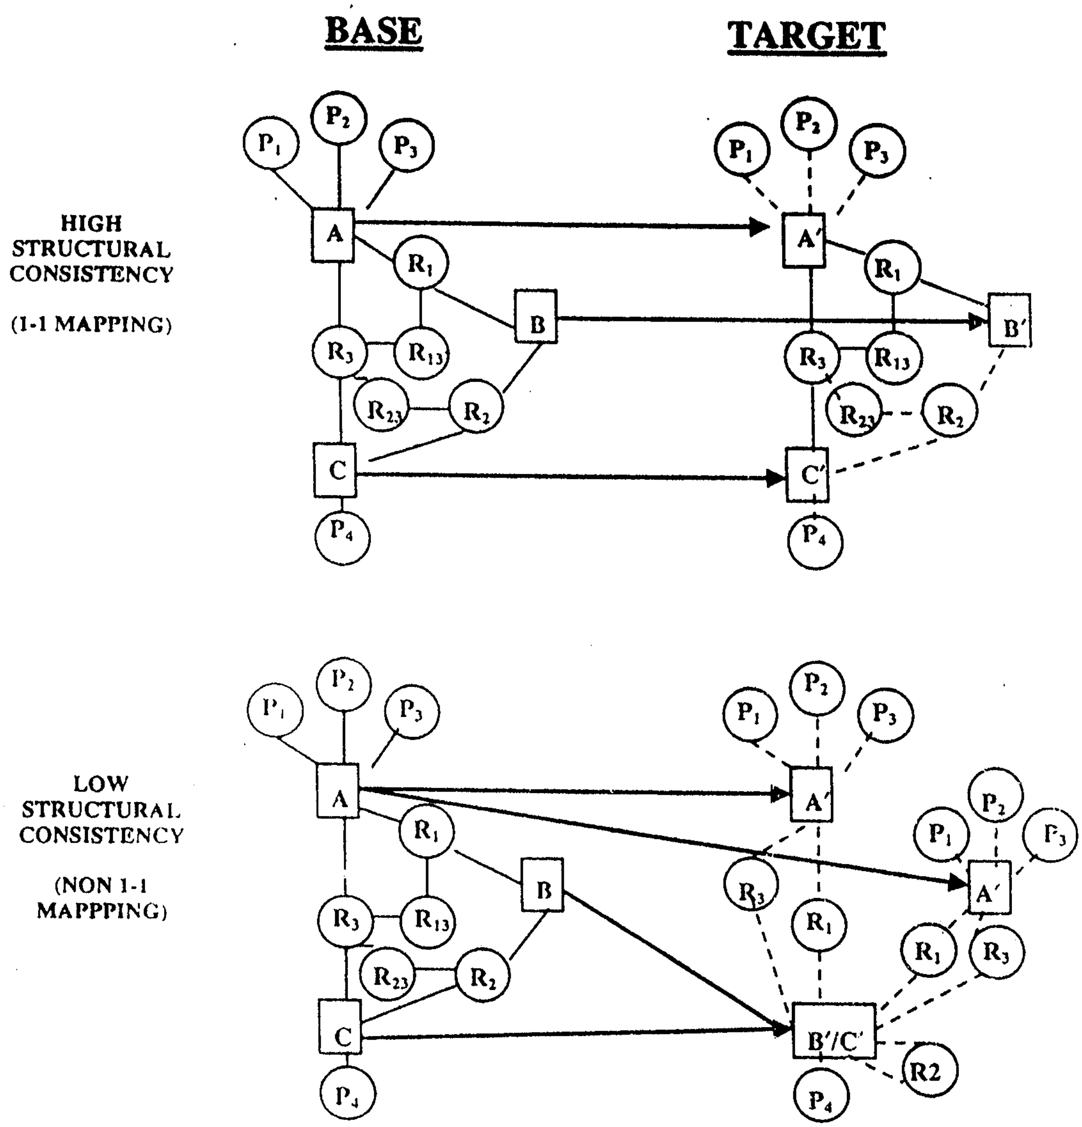

FIG. 11.7. Structural consistency and inference projection.

he eats too much, or that he is a policeman). Other cases seem to be global mappings, such as mapping from vertical dimension to emotions like sadness ("His spirits sank"; "He had never felt so low") or the computer metaphor for cognitive processing just discussed.

Still other metaphors seem intermediate between local categories and global system mappings. For example, Turner (1987) noted in his discussion of kinship metaphors that *mother* is often used to convey that one thing is the source or cause of another, as in "Necessity is the mother of invention." This entails a fairly local mapping—two base elements (mother and child) and some relations between them. Turner (p. 156) also cites examples that draw on a larger portion of the base: "Aristotle sayeth that the erthe is moder and the sonne fader of trees." Here the base elements include *mother, father,* and *child,* together with a relational system that is sufficiently elaborated to invite different father–child relations from mother–child relations and even perhaps a father–mother relation. Here *mother* participates in a system metaphor.

METAPHORIC CREATIVITY

Metaphors have always been associated with creativity. Recently, this aspect of metaphor has received increased attention. In a study by Tourangeau and Rips (1991), subjects wrote down interpretations for a series of metaphors and listed properties for each of the terms on which the metaphor was based. When interpretations were broken down by type, a majority of properties listed for the interpretations were not listed in either the target or base object descriptions. Many of the features listed for the interpretations were thus "emergent": that is, features that were not apparent in either the base or target representations. Similar results were obtained by Gentner and Clement (1988). The emergent nature of interpretations has implications for metaphor processing. To the extent that the properties of an interpretation are constructed or discovered online, then even for perfectly acceptable metaphors, there is no necessary pre-existing relation between the two words. One term of a metaphor need not prime the other. This prediction was tested by Camac and Glucksberg (1984). Subjects were shown a pair of letter strings and asked to decide whether one or both the strings were words. The word pairs were either preassociated, metaphorically associated, or nonassociated. Word pairs that were preassociated were responded to quickly. However, metaphorical pairs were responded to no faster than nonassociated pairs. The degree of preassociation for metaphorically related word pairs can therefore be quite low, supporting the constructive or emergent nature of metaphor interpretation.

Additional evidence for the existence of emergent features comes from a study by Blasko and Connine (1993). Subjects were instructed to make a lexical decision about a visually present word after hearing either the base or the target of a metaphorical sentence. Half the words to be verified were related to the meaning of the metaphor formed by the base and target. The other half were unrelated to the metaphor's meaning. The main result was that test words related to the meaning of the metaphor were verified no faster than unrelated words. In contrast, metaphorically related test words were facilitated when subjects heard both the target and base together. (See also Verbrugge & McCarrell, 1977.) These results suggest that the meaning emerges from combining the two forms and does not simply pre-exist in one of the individual terms. In sum, the results from these studies indicate that metaphor is most deeply and fruitfully understood when both terms are simultaneously present. These findings fit with the evidence offered earlier that a novel metaphor is first and foremost an alignment of representations (Gentner & Wolff, 1997; Wolff & Gentner, 1992, in press). Juxtaposing two terms invites a process of structural alignment and mapping that result in highlighting, candidate inferences, and re-representation. We suggest that comparison is crucial to creativity in metaphor.

These mechanisms and products are not unique to metaphor. On this account, metaphor is highly related to literal similarity—that is, to comparisons with substantial relational and object-attributional commonalties. To put it another way, we have argued here that novel metaphors are like analogies and elsewhere that similarity is like analogy (Gentner & Markman, 1995, 1997); to complete the pattern, we now suggest that metaphor is like similarity. The claim that metaphoric processing is like literal processing has a long history in the field (Gentner, 1983, 1989; Gibbs, 1984; Gibbs & Gerrig, 1989; Ortony, Vondruska, Foss, & Jones, 1985; Rumelhart, 1979).

We suggest that the processes of knowledge change discussed here apply to mundane literal similarity as well as to analogy and metaphor. Why then are the results so much more noticeable with analogy and metaphor? At least two factors enter into whether or not one concept is capable of changing another. One is the possibility of a good alignment. The other is the presence of some differences—in effect, some reason to change. Literal similarity has the advantage on the first factor; metaphor, on the second. In literal similarity, the concepts are already so close that the resulting adjustments to the representations are small. With metaphors, however, there are sufficient differences to permit substantial change.

The notion that differences are critical to metaphor comprehension may seem to fly in the face of common sense. According to Glucksberg and Keysar (1990), differences play little or no role in metaphor comprehension; if metaphors did focus on differences, their illuminatory effect would surely be overwhelmed by the vast number of differences. In the same vein, Ortony (1979a) suggested that both the difference weights in Tversky's (1977) contrast model equation for similarity might best be set to zero, to reflect the lack of relevance of differences in metaphor. Although this intuition is appealing, it leads to a conundrum: If differences between base and target are ignored, then how can the base transform the target? We suggest that *some* differences are important. Research by Markman and Gentner suggests that in carrying out comparisons people focus on *alignable differences*—differences connected to the common structure in the same way, such as the fact that cars have four wheels while motorcycles have only two. People discount *nonalignable differences*—differences unconnected to the common structure, or connected differently, such as the fact that motorcycles are popular with gangs and cars have steering wheels (Gentner & Markman, 1994, 1995; Markman & Gentner, 1993a). However, alignable differences are not only noticed, but are a prime source of new inferences. Bowdle and Gentner (1997) gave people pairs of stories varying in their alignability and relative systematicity and asked them to make inferences in either direction they chose. People strongly preferred to make inferences from the more systematic to the less systematic story, if the scenarios were alignable. (Nonalignable pairs yielded no directional preference, as predicted.) Furthermore, the inferences were typically based on alignable differences between the scenarios.

CONCLUSIONS

We have discussed the ways in which metaphors can change knowledge. Metaphors can highlight, project, re-represent, and, occasionally perhaps, restructure. These processes can lead to metaphorical categories and stored mappings. We have argued that these mechanisms are ways by which metaphors can lead to change of knowledge.

How have we fared with respect to Fodor's challenge that one can learn only what is already present? On the minus side, our proposals do not solve the problem of where the original predicates come from. Still, there are some kinds of knowledge change that we can account for:

- Candidate inference projection: A predicate *P*—previously expressed in the assertion $P\ (b_1, b_2)$—is mapped from base to target, so that the combination $P\ (t_1, t_2)$ is expressed for the first time.
- Re-representation: A predicate *K* is aligned with a predicate *L*, resulting in a re-representation that creates a slightly new predicate, *M*—for instance, *trail* (*b*1, *b*2) and *chase* (*t*1, *t*2) may result in *pursue* (*x*1, *x*2).
- Schema abstraction: A system of assertions common to both base and target is abstracted and stored as a schema, resulting in a new predicate—sometimes with the invitation of a relational label: for example, learning the term *symmetric*.

We suggest that these changes are psychologically significant. They offer a means of attaining a conceptual system richer than the initial system. Candidate inference projection, although its effects are modest on each application, is a generative mechanism. For every new application of a given analog, new inferences can occur, resulting in an indefinite number of potential new combinations of predicates and arguments. Schema abstraction—abstracting and storing a higher-order pattern—can facilitate noticing the same pattern in the future. After sufficient in-domain experience or after learning a word for *symmetric*, Kotovsky and Gentner's 4-year-olds could see cross-dimensional similarity patterns they had previously missed. If this is not learning, it is a good facsimile. Eventually, these processes can reveal ideas more general and powerful than their original instantiations suggested. As Bertrand Russell put it, "It must have required many ages to discover that a brace of pheasants and a couple of days were both instances of the number two."

ACKNOWLEDGMENTS

This work was supported by National Science Foundation grant SBR-95-11757 and ONR grant N00014-92-J-1098 to the first author. We thank Arthur Markman, Ron Ferguson, Ken Forbus, Sarah Brem, Brian Bowdle, and the rest of the Similarity and Analogy group for discussions during the evolution of this work.

REFERENCES

Allbritton, D. W., McKoon, G., & Gerrig, R. (1995). Metaphor-based schemas and text representations: Making connections through conceptual metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1–14.

Andrade, E. N. Da C. (1964). *Rutherford and the nature of the atom*. New York: Doubleday.

Asch, S. E. (1955). On the use of metaphor in the description of persons. In H. Werner (Ed.), *On expressive language* (pp. 29–38). Worcester, MA: Clark University Press.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94*, 115–147.

Black, M. (1962). *Models and metaphors*. Ithaca: Cornell University Press.

Black, M. (1979). More about metaphor. In A. Ortony (Ed.), *Metaphor and thought* (pp. 19–43). Cambridge: Cambridge University Press.

Blank, G. D. (1988). Metaphors in the lexicon. *Metaphor and Symbolic Activity, 3*, 21–36.

Blasko, D. G., & Connine, C. M. (1993). Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 205–308.

Boronat, C. B. (1990). *Effects of base shift and frequency in extended metaphor processing*. Unpublished master's thesis, University of Illinois at Urbana-Champaign.

Bowdle, B., & Gentner, D. (1995, November). *The career of metaphor*. Paper presented at the meeting of the Psychonomics Society, Los Angeles, CA.

Bowdle, B., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology, 34*(3), 244–286.

Bowdle, B., & Gentner, D. (1999a). Metaphor comprehension: From comparison to categorization. *Proceedings of the twenty-first annual conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Bowdle, B., & Gentner, D. (1999b). *The career of metaphor*. Manuscript in preparation.

Boyd, R. (1979). Metaphor and theory change: What is "metaphor" a metaphor for? In A. Ortony (Ed.), *Metaphor and thought* (pp. 186–201). Cambridge, England: Cambridge University Press.

Brown, A. L., Kane, M. J., & Echols, C. H. (1986). Young children's mental models determine analogical transfer across problems with a common goal structure. *Cognitive Development, 1,* 103–121.

Burstein, M. H. (1986). Concept formation by incremental analogical reasoning and debugging. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 351–369). Los Altos, CA: Morgan Kaufmann.

Cacciari, C. (1993). The place of idioms in a literal and metaphorical world. In C. Cacciari & P. Tabossi (Eds.), *Idioms: Processing, structure, and interpretation* (pp. 27–55). Hillsdale, NJ: Lawrence Erlbaum Associates.

Cacciari, C., & Tabossi, P. (1988). The comprehension of idioms. *Journal of Memory and Language, 27,* 668–683.

Camac, M. K., & Glucksberg, S. (1984). Metaphors do not use associations between concepts, they are used to create them. *Journal of Psycholinguistic Research, 13,* 443–455.

Campbell, N. R. (1920). *Physics, the elements.* Cambridge, England: Cambridge University Press.

Carbonell, J. G. (1981). Invariance hierarchies in metaphor interpretation. *Proceedings of the third annual conference of the Cognitive Science Society* (pp. 292–295). Hillsdale, NJ: Lawrence Erlbaum Associates.

Carbonell, J. G. (1982). Metaphor: An inescapable phenomenon in natural language comprehension. In W. G. Lehnert & M. H. Ringle (Eds.), *Strategies for natural language processing* (pp. 415–435). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chen, Z., & Daehler, M. W. (1989). Positive and negative transfer in analogical problem solving by 6-year-old children. *Cognitive Development, 4,* 327–344.

Clark, H., & Lucy, P. (1975). Understanding what is meant from what is said: A study in conversationally conveyed requests. *Journal of Verbal Learning and Verbal Behavior, 14,* 56–72.

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

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

Cooke, N. J., & Bartha, M. C. (1992). An empirical investigation of psychological metaphor. *Metaphor and Symbolic Activity, 7,* 215–235.

Dreistadt, R. (1968). An analysis of the use of analogies and metaphors in science. *Journal of Psychology, 68,* 97–116.

Elio, R., & Anderson, J. R. (1981). The effect of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 397–417.

Eve, A. S. (1939). *Rutherford.* Cambridge, England: Cambridge University Press.

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* (pp. 157–196). Los Altos, CA: Morgan Kaufmann.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of the fifth national conference on Artificial Intelligence* (pp. 272–277). Los Altos, CA: Morgan Kaufmann.

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

Fauconnier, G. (1990). Domains and connections. *Cognitive Linguistics, 1*(1), 151–174.

Fodor, J. A. (1975). *The language of thought.* New York: Random House.

Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure-mapping. *Proceedings of the 16th annual conference of the Cognitive Science Society* (pp. 313–318). Hillsdale, NJ: Lawrence Erlbaum Associates.

Forbus, K. D., & Gentner, D. (1986). Learning physical domains: Toward a theoretical framework. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 311–348). Los Altos, CA: Morgan Kaufmann.

Forbus, K. D., & Gentner, D. (1989). Structural evaluation of analogies: What counts? *Proceedings of the 11th annual conference of the Cognitive Science Society* (pp. 341–348). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the 12th annual conference of the Cognitive Science Society* (pp. 61–68). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gentner, D. (1977). Children's performance on a spatial analogies task. *Child Development, 48*, 1034–1039.

Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory, 4*(2), 161–178.

Gentner, D. (1982). Are scientific analogies metaphors? In D. S. Miall (Ed.), *Metaphor: Problems and perspectives* (pp. 106–132). Brighton, England: Harvester.

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

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

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

Gentner, D., & Boronat, C. B. (1992, May). *Metaphor as mapping*. Paper presented at the Workshop on Metaphor, Tel Aviv, Israel.

Gentner, D., & Boronat, C. (in preparation). *Novel metaphors are processed as generative domain-mappings*.

Gentner, D., & Bowdle, B. (1994). The coherence imbalance hypothesis: A functional approach to asymmetry in comparison. In *Proceedings of the 16th annual meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (in press). Metaphor is like analogy. In D. Gentner, K. J. Holyoak, & B. Kokinov (Eds.), *Analogy: Theory and phenomena*. Cambridge, MA: MIT Press.

Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., & Forbus, K. D. (1997). Analogical reasoning and conceptual change: A case study of Johannes Kepler. *Journal of the Learning Sciences, 6*(1), 3–40.

Gentner, D., & Clement, C. A. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower, (Ed.), *The psychology of learning and motivation, advances in research and theory* (Vol. 22, pp. 307–358). New York: Academic Press.

Gentner, D., Falkenhainer, B., & Skorstad, J. (1988). Viewing metaphor as analogy. In D. H. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence, cognitive science, and philosophy* (pp. 171–177). Dordrecht, The Netherlands: Kluwer.

Gentner, D., & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 99–129). Hillsdale, NJ: Lawrence Erlbaum Associates. (Reprinted in *Cognitive functions: Classic readings in representation and reasoning*, M. J. Brosnan, Ed., Eltham, London: Greenwich University Press)

Gentner, D., & Grudin, J. (1985). The evolution of mental metaphors in psychology: A 90-year retrospective. *American Psychologist, 40*(2), 181–192.

Gentner, D., & Imai, M. (1992). Is the future always ahead? Evidence for system-mappings in understanding space-time metaphors. *Proceedings of the 14th annual meeting of the Cognitive Science Society* (pp. 510–515).

Gentner, D., Imai, M., & Boroditsky, L. (1999). *As time goes by: Evidence for two systems in processing space–time metaphors.* Manuscript submitted for publication.

Gentner, D., & Jeziorski, M. (1993). The shift from metaphor to analogy in western science. In A. Ortony (Ed.), *Metaphor and thought* (2nd ed., pp. 447–480). Cambridge, England: Cambridge University Press.

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* (Vol. 5, pp. 855–862). San Mateo, CA: Morgan Kauffman.

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

Gentner, D., & Markman, A. B. (1995). Similarity is like analogy: Structural alignment in comparison. In C. Cacciari (Ed.), *Similarity in language, thought and perception* (pp. 111–147). Brussels, Belgium: BREPOLS.

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

Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition, 65,* 263–297.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes, (Eds.), *Perspectives on language and thought interrelations in development* (pp. 225–277). Cambridge, England: Cambridge University Press.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability and 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 T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 263–313). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gentner, D., & Stevens, A. L. (Eds.). (1983). *Mental models.* Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

Gibbs, R. W., Jr. (1979). Contextual effects in understanding indirect requests. *Discourse Processes, 2,* 149–156.

Gibbs, R. W., Jr. (1980). Spilling the beans on understanding and memory for idioms in conversations. *Memory and Cognition, 8,* 449–456.

Gibbs, R. W., Jr. (1984). Literal meaning and psychological theory. *Cognitive Science, 8,* 275–304.

Gibbs, R. W., Jr. (1990). Psycholinguistic studies on the conceptual biases of idiomaticity. *Cognitive Linguistics, 1,* 417–451.

Gibbs, R. W., Jr. (1992). Categorization and metaphor understanding. *Psychological Review,* 99(3), 572–577.

Gibbs, R. W., Jr. (1994). *Poetics of mind: Figurative thought: Language and understanding.* Cambridge: Cambridge University Press.

Gibbs, R. W., Jr., & Gerrig, R. J. (1989). How context makes metaphor comprehension seem "special." *Metaphor and Symbolic Activity, 4,* 145–158.

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

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

Glucksberg, S., Brown, M., & McGlone, M. S. (1993). Conceptual analogies are not automatically accessed during idiom comprehension. *Memory and Cognition, 21*, 711–719.

Glucksberg, S., Gildea, P., & Bookin, H. (1982). On understanding non-literal speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior, 1*, 85–96.

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review, 97*, 3–18.

Glucksberg, S., & Keysar, B. (1993). How metaphors work. In A. Ortony (Ed.), *Metaphor and thought* (2nd ed., pp. 401–424). New York: Cambridge University Press.

Glucksberg, S., McGlone, M. S., & Manfredi, D. (1997). Property attribution in metaphor comprehension. *Journal of Memory and Language, 36*(1), 50–67.

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

Goldstone, R. L., Gentner, D., & Medin, D. L. (1989). Relations relating relations. *Proceedings of the 11th annual conference of the Cognitive Science Society* (pp. 131–138). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects* (pp. 437–447). New York: Bobbs-Merrill. (Reprinted from *Experience and theory*, L. Foster & J. W. Swanson, Eds., 1970, University of Massachusetts Press)

Graesser, A., & Hopkinson, P. (1987). Differences in interconcept organization between nouns and verbs. *Journal of Memory and Language, 26*, 242–253.

Greenberg, J. H. (1966). *Language universals*. The Hague, The Netherlands: Mouton.

Hayes-Roth, F., & McDermott, J. (1978). An interference matching technique for inducing abstractions. *Communications of the ACM, 21*(5), 401–411.

Hesse, M. B. (1966). *Models and analogies in science*. Notre Dame, IN: University of Notre Dame Press.

Hoffman, R. R., & Kemper, S. (1987). What could reaction-time studies be telling us about metaphor comprehension? *Metaphor and Symbolic Activity, 2*, 149–186.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 19, pp. 59–87). New York: Academic Press.

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

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

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

Huttenlocher, J., & Lui, F. (1979). The semantic organization of some simple nouns and verbs. *Journal of Verbal Learning and Verbal Behavior, 18*, 141–162.

Indurkhya, B. (1991). Models in metaphor. *Metaphor and Symbolic Activity, 6*, 1–27.

Johnson-Laird, P. N. (198[illegible]). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.

Karmiloff Smith, A. (1991). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: The MIT Press.

Keane, M. T. (1988). Analogical mechanisms. *Artificial Intelligence Review, 2*, 229–250.

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

Kelly, M. H., & Keil, F. C. (1987). Metaphor comprehension and knowledge of semantic domains. *Metaphor and Symbolic Activity, 2*, 35–51.

Kennedy, J. M. (1990). Metaphor—Its intellectual basis. *Metaphor and Symbolic Activity, 5*, 115–123.

Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language, 28*, 375–385.

Keysar, B., & Bly, B. (1995). Intuitions of the transparency of idioms: Can one keep a secret by spilling the beans? *Journal of Memory and Language, 34*(1), 89–109.

Kittay, E. F., & Lehrer, A. (1981). Semantic fields and the structure of metaphor. *Studies in Language, 5*, 31–63.

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

Kuhn, T. S. (1979). Metaphor in science. In A. Ortony (Ed.), *Metaphor and thought* (pp. 186–201). Cambridge, England: Cambridge University Press.

Lakoff, G. (1988). Conventional metaphor as the basis for analogical reasoning. In D. H. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence.* Boston: D. Reidel.

Lakoff, G. (1990). The invariance hypothesis: Is abstract reason based on image-schemas? *Cognitive Linguistics, 1*(1), 39–74.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago: University of Chicago Press.

Lakoff, G., & Turner, M. (1989). *More than cool reason: A field guide to poetic metaphor.* Chicago: University of Chicago Press.

Lassaline, M. E., & Murphy, G. L. (1996). Induction and category coherence. *Psychonomic Bulletin and Review, 3*(1), 95–99.

Lehrer, A., (1990). Polysemy, conventionality, and the structure of the lexicon. *Cognitive Linguistics, 1–2*, 207–246.

Levi, J. N. (1978). *The syntax and semantics of complex nominals.* New York: Academic Press.

Lockhead, G. R., & King, M. C. (1977). Classifying integral stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 3*(3) 436–443.

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

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

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

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

Markman, A. B., & Gentner, D. (in press). Structure-mapping in analogy and similarity: The importance of being connected. *American Journal of Psychology.*

McGlone, M. S., & Harding, J. (1998). Back (or forward?) to the future: The role of perspective in temporal language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*(5), 1211–1223.

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

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

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving, and induction. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 5, pp. 189–223). Hillsdale, NJ: Lawrence Erlbaum Associates.

Miller, G. A. (1979). Images and models, similes and metaphors. In A. Ortony (Ed.), *Metaphor and thought* (pp. 202–250). Cambridge, England: Cambridge University Press.

Miller, G. A. (1993). Images and models, similes and metaphors. In A. Ortony (Ed.), *Metaphor and thought* (2nd ed., pp. 357–400). Cambridge, England: Cambridge University Press.

Murphy, G. L. (1996). On metaphoric representation. *Cognition, 60*(2), 173–204.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289–316.

Nagy, W. (1974). *Figurative patterns and the redundancy in lexicon.* Unpublished doctoral dissertation, University of California, San Diego.

Nayak, N. P., & Gibbs, R. W. (1990). Conceptual knowledge in the interpretation of idioms. *Journal of Experimental Psychology, 119*(3), 315–330.

Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. N. Giere, & H. Feigl (Eds.), *Minnesota studies in the philosophy of science* (pp. 3–44). Minneapolis: University of Minnesota Press.

Norman, D. A., Rumelhart, D. E., & the LNR Research Group. (1975). *Explorations in cognition.* San Francisco: Freeman.

Oppenheimer, R. (1956) . Analogy in science. *American Psychologist, 11*, 127–135.

Ortony, A. (1979a). Beyond literal similarity. *Psychological Review, 86*, 161–180.

Ortony, A. (1979b). The role of similarity in similes and metaphors. In A. Ortony (Ed.), *Metaphor and thought* (pp. 1–16). Cambridge, England: Cambridge University Press.

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.

Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review, 56*, 132–143.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology, 9*, 441–474.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259–303). Hillsdale, NJ: Lawrence Erlbaum Associates.

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance, 3*, 422–435.

Pylyshyn, Z. W. (1979). Metaphorical imprecision and the "top-down" research strategy. In A. Ortony (Ed.), *Metaphor and thought* (pp. 186–201). Cambridge, England: Cambridge University Press.

Pylyshyn, Z. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition, 32*(1), 65–97.

Rattermann, M. J., Gentner, D., & DeLoache, J. (1989). *Effects of competing surface similarity on children's performance in an analogical task.* Poster presented at the biennial meeting of the Society for Research in Child Development, Kansas City, MO.

Rattermann, M. J., Gentner, D., & DeLoache, J. (1999). *Effects of relational and object similarity on children's performance in a mapping task.* Manuscript in preparation.

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

Roediger, H. L., III. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition, 8*, 231–246.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(4) 629–639.

Ross, B. H. (1989). Remindings in learning and instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 438–469). New York: Cambridge University Press.

Rumelhart, D. (1979). Some problems with the notion of literal meanings. In A. Ortony (Ed.), *Metaphor and thought* (pp. 186–201). Cambridge, England: Cambridge University Press.

Rumelhart, D. E., & Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology, 5,* 1–28.

Rumelhart, D. E., & Ortony, A. (1977). Representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. C., & Abelson, R. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Schön, D. A. (1979). Generative metaphor: A perspective on problem-setting in social policy. In A. Ortony (Ed.), *Metaphor and thought* (pp. 186–201). Cambridge, England: Cambridge University Press.

Shen, Y. (1992). Metaphors and categories. *Poetics Today, 13,* 771–794.

Smith, L. B., & Sera, M. D. (1992). A developmental analysis of the polar structure of dimensions. *Cognitive Psychology, 24,* 99–142.

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

Spiro, R. J., Feltovich, P. J., Coulson, R. L., & Anderson, D. K. (1989). Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 498–531). New York: Cambridge University Press.

Steen, G. J. (1991). The empirical study of literary reading: Methods of data collection. *Poetics Today, 13.*

Swinney, D., & Cutler, A. (1979). The access and processing of idiomatic expressions. *Journal of Verbal Learning and Verbal Behavior, 18,* 523–534.

Thagard, P., & Holyoak, K. (1985). Discovering the wave theory of sound: Inductive inference in the context of problem solving. *Proceedings of the ninth international joint conference on Artificial Intelligence* (pp. 610–612). Los Altos, CA: Morgan Kaufmann.

Tourangeau, R., & Rips, L. (1991). Interpreting and evaluating metaphors. *Journal of Memory and Language, 30,* 452–472.

Tourangeau, R., & Sternberg, R. J. (1981). Aptness in metaphor. *Cognitive Psychology, 13,* 27–55.

Tourangeau, R., & Sternberg, R. J. (1982). Understanding and appreciating metaphors. *Cognition, 11,* 203–244.

Turner, M. (1987). *Death is the mother of beauty: Mind, metaphor, and criticism.* Chicago: University of Chicago Press.

Turner, M. (1988). Categories and analogies. In D. H. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence, cognitive science, and philosophy* (pp. 3–24). Dordrecht, The Netherlands: Kluwer.

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

Verbrugge, R. R., & McCarrell, N. S. (1977). *Metaphoric comprehension: Studies in reminding and resembling.*

Vosniadou, S., Ortony, A., Reynolds, R. E., & Wilson, P. T. (1984). Sources of difficulty in the young child's understanding of metaphorical language. *Child Development, 55,* 1588–1606.

Way, E. C. (1991). *Knowledge representation and metaphor.* Dordrecht, The Netherlands: Kluwer.

Wheelwright, P. (1962). *Metaphor and reality,* Bloomington: Indiana University Press.

Wilson, D. (1983). *Rutherford: Simple genius.* Cambridge, MA: MIT Press.

Winner, E. (1988). *The point of words: Children's understanding of metaphor and irony.* Cambridge, MA: Harvard University Press.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM, 23*(12), 689–703.

Wiser, M. (1986). *Learning about heat and temperature: A context-based historically inspired approach to a novice-expert shift*. Paper presented at the National Science Foundation Conference on the Psychology of Physics Problem Solving: Theory and Practice, New York.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language, 35*, 434–453.

Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 241–284). Amsterdam: Elsevier.

Wolff, P., & Gentner, D. (1992). The time course of metaphor comprehension. *Proceedings of the 14th annual conference of the Cognitive Science Society* (pp. 504–509). Bloomington.

Wolff, P., & Gentner, D. (1999). From symmetric to asymmetric processing: Two stages in the comprehension of metaphors.

Wolff, P., & Gentner, D. (in press). Evidence for role-neutral initial processing of metaphors. *Journal of Experimental Psychology: Learning, Memory and Cognition*.

Yukawa, H. (1982). *Tabibito* (L. Brown & R. Yoshida, Trans.). Singapore: National Printers.