In Developing Cognitive Competence: New approaches to process modeling, G. Halford & T. Simon, eds. Hillsdale: Erlbaum, 1995

6

Two Forces in the Development of Relational Similarity

Dedre Gentner
Northwestern University

Mary Jo Rattermann
Hampshire College

Arthur Markman
Columbia University

Laura Kotovsky
University of California at Los Angeles

Analogy commands the attention of developmental psychologists, first, because like grammar and mathematics, analogy is a supremely elegant form of thought; and, second, because of its importance in cognitive development. An appreciation of relational similarity is fundamental to learning beyond the basic level—to grasping theory-laden concepts that must be defined relationally, such as *predator* in biology and *limiting case* in mathematics. In this chapter, we explore two forces that promote the development of relational similarity. Our goal is to illuminate both the nature of the similarity mechanism and its role in experiential learning.

Despite the attention given to how and when children acquire the ability to process relational similarity—to carry out analogies—a number of important issues remain unresolved. Most researchers agree that there is a relational shift from early reliance on either holistic or object-level similarities to the possibility of purely relational similarities (Gentner, 1988; Gentner & Rattermann, 1991). However, there is disagreement as to the nature of the shift. Is it governed by cognitive stage or by degree of domain knowledge; is it maturational or the product of learning; is it an all-or-none shift from object similarity to relational similarity; and finally, when does

it occur, with nominations ranging from *during infancy* through *at formal operations*?

A precise model of the comparison process is needed to clarify the question of what changes with development. Our account is based on modeling the comparison process as one of alignment and mapping of representations. We make three chief points. First, we argue that analogical development is primarily a matter of changes in knowledge, rather than changes in global competence or processing capacity. Our second claim is that language learning—specifically, the acquisition of relational terms—is crucial in the development of relational comparison. Our third claim is that *the process of similarity comparison* is instrumental in the development of analogy. We present evidence that the comparison process itself invites attention to relational structure and that it can promote the acquisition of portable relational knowledge. We conjecture that the acquisition of relational language and the process of relational comparison provide mutual bootstrapping that drives representational change.

We begin with a review of the research and theoretical issues. Then, we outline the structure-mapping approach and present the Structure-mapping Engine (SME), our simulation of similarity and analogy comparison, along with some adult findings that illustrate the computational principles. Next, we present two studies of children's acquisition of relational comparisons and model these developmental studies using SME. Our results suggest that the basic similarity process is the same for children as it is for adults. What varies is the knowledge representations this process acts on.

THE DEVELOPMENT OF RELATIONAL SIMILARITY

There is considerable evidence for a relational shift in development of similarity (Gentner, 1988). After reviewing a large number of studies, Gentner and Rattermann (1991) proposed the following account of the *career of similarity*. Young infants tend to respond to overall (literal) similarity and identity between scenes, such as the similarity between a red ball rolling and another red ball rolling. The earliest partial matches are based on *object similarity*: direct resemblances between objects, such as the similarity between a round red ball and a round red apple. With increasing knowledge, children come to make single-attribute matches such as the similarity between a red ball and a red car, and finally, *relational similarity* matches, such as the similarity between a ball *rolling on* a table and a toy car *rolling on* the floor. For example, when asked to interpret metaphors like *A tape recorder is like a camera*, 6-year-olds produced object-based interpretations such as *Both are metal and black*, whereas 9-year-olds and adults focused chiefly on common relational structure (e.g., *Both can record* *something for later*; Gentner, 1988). Billow (1975) reported that metaphors based on object similarity could be correctly interpreted by children of about 5 or 6 years of age, but that relational metaphors were not correctly interpreted until around 10 to 13 years of age. Gentner and Toupin (1986) contrasted the effects of object similarities and relational structure. Children were shown a simple story acted out by toy characters, and then asked to re-enact the story with a new group of characters. For some stories, the new characters were similar to their corresponding original characters; in others, the characters were different; and in the worst case, the *cross-mapped analogy* condition, similar characters played different roles across the two stories, so that object similarities interfered with the best plot mapping. Both 6- and 9-year-olds were highly sensitive to object similarity: They were most accurate at re-enacting the story in the high-similarity condition and least accurate in the cross-mapped condition. For half the children, the relational structure was augmented by adding an explicit causal or moral summary statement. Under these conditions, 9-year-olds were nearly perfect even on the cross-mapped condition. However, 6-year-olds showed no improvement. Thus, both groups benefited from object commonalities that supported the structural mapping, but, in accord with the relational shift hypothesis, only the older group benefited from the presence of higher- order relational structure.

Global Competence Views of Relational Development

Explanations for the relational shift fall into two main classes: those that posit a global change in cognitive competence, and those that posit domain specific shifts driven by change of knowledge.

Piaget linked the development of higher order relational similarity to global shifts in competence—specifically, the acquisition of formal operations at around 11 or 12 years of age (Inhelder & Piaget, 1958; Piaget, Montangero, & Billeter, 1977). Piaget, Montangero, and Billeter presented children with pictures forming A:B::C:? analogies such as *Bicycle:Handlebars::Ship:?* and asked them to choose the best completing picture. Although older children (9-year-olds) were able to choose the correct relational response (here, a *rudder*), young children (5-year-olds) often responded with thematic associates (e.g., *seagull*). This shift from thematic to relational responding was also found by Sternberg and his colleagues using materials such as A:B::C:D analogies constructed from pictures of human forms (Sternberg & Downing, 1982; Sternberg & Nigro, 1980; Sternberg & Rifkin, 1979).

The idea that perception of relational similarity must await formal operations has few current adherents, because of the many studies demonstrating earlier abilities. However, it is still possible to maintain that there is a global cognitive shift that permits relational mapping, but that it occurs well before the formal operations stage. Halford and his colleagues have explored this position. They provide an elegant and precise account of the relational shift based largely on developmental increases in processing capacity (Halford, 1987, 1992, chap. 3, this volume). On their view, young children's cognitive capacity is insufficient to process the information needed to find relational similarities. Over the course of development, processing capacity increases (Halford, Mayberry, & Bain, 1986) permitting the child to perform increasingly more complex analogical mappings (Halford, 1992). Halford proposed that children go through four stages of analogical ability: (a) *element mappings*, in which element correspondences are based on matching common attributes; (b) *relation mappings* (at around 2 years), in which simple binary relations are placed in correspondence, along with the elements associated with these relations; (c) *system mappings* (at around 4 or 5 years), in which systems of two relations are placed in correspondence; and (d) *multiple-system mappings* (at around 11 years), in which systems of three or more relations are placed in correspondence. Although Halford (chap. 3, this volume) also allows for knowledge effects, such as chunking of information, his account of the relational shift emphasizes global change and a maturational increase in processing capacity.

Knowledge-Change Views of Relational Development

Knowledge-based accounts assume that the relational shift results from changes in knowledge, not from global and/or maturational changes (Brown, 1989; Brown & DeLoache, 1978; Brown & Kane, 1988; Chen & Daehler, 1989; Crisafi & Brown, 1986; Gentner, 1977a, 1977b, 1988; Gentner & Rattermann, 1991; Ortony, Reynolds, & Arter, 1978; Vosniadou, 1987, 1989). On this view, the chief predictor of whether a child can carry out a comparison is his or her knowledge of the two domains.

One line of evidence for the knowledge-change view is that even very young children can show considerable analogical ability provided the domains are familiar. For example, Gentner (1977a, 1977b) demonstrated that preschool children can perform a spatial analogy between the familiar base domain of the human body and simple pictured objects, such as trees and mountains. When asked, "If the tree had a knee, where would it be?," 4-year-olds (as well as 6- and 8-year-olds) were as accurate as adults in performing the mapping of the human body to the tree, even when the orientation of the tree was changed or when confusing surface attributes were added to the pictures. Ann Brown has conducted a number of ingenious experiments demonstrating early transfer abilities. In one study, she found that after one experience using a tool to pull a desirable toy closer, 20- to 30-month-old children would transfer the notion of *pulling tool* to a new situation and choose an appropriate tool from a new set of transfer objects (Brown, 1989; 1990). In other studies, Brown and Kane (1988) taught 4-year-olds about biological mechanisms such as mimicry and camouflage and found that children could map from a base scenario of *using ladybugs to control aphids* to propose the solution *use purple martins to control mosquitoes*. It is probably true that in both these tasks the relational mapping was supported by some object-level similarity. For example, both the pulling tools had to be long with hooklike protuberances, resulting in some common shape. In the second study, the similarity between aphids and mosquitoes may have contributed to children's performance. It is notoriously difficult to control object similarity in designing naturalistic analogical materials, partly because real causal scenarios often involve similar objects in similar relations (See Gentner, Rattermann, & Campbell, 1994). However, although these results do not provide evidence for purely relational transfer, they do support the claim that greater knowledge permits greater degrees of structure-sensitive mapping.

A second line of evidence for the claim that the relational shift is a knowledge-driven shift rather than a global cognitive shift is that it occurs at different ages for different domains and tasks. Gentner and Rattermann (1991) came to this conclusion on the basis of a survey of developmental findings. For example, in the story-mapping task discussed in the preceding section, Gentner and Toupin (1986) found a shift between 6 years and 9 years: Only the older children benefited from the presence of higher-order relational structure. Yet 4-year-olds show some ability to perform relational mappings in Brown and Kane's biological mechanisms task and on Gentner's body parts mapping task. There is even evidence that a relational shift can occur during infancy. Kolstad and Baillargeon (1991) showed babies repeated events of salt being poured into and out of containers, first showing the infants the container, including its sides and bottom. After the babies were familiarized to this event of filling and pouring out salt, they were shown a transfer event: the same pouring event with one of two new objects. One of the transfer objects was perceptually different from the training objects, but was otherwise a perfectly normal container. The other object was perceptually similar to the original, but appeared to have no bottom. Kolstad and Baillargeon found that 5.5-month-old infants looked reliably longer (indicating surprise) at the perceptually dissimilar event, but not at the "bottomless" container. In contrast, 10.5-month-old infants looked reliably longer at the causally different event, the "bottomless" container. Their encoding of the first event apparently included the notion that the bottom of the container had supported the salt. Their surprise reaction occurred in response to causally relevant relational commonalities and not to overall similarity. This suggests that during their first year of life, babies may undergo a small but distinctive relational shift within the highly familiar arena of *containment*.

The evidence summarized suggests that the shift occurs at different times for different domains because of differential knowledge: the deeper the child's domain knowledge, the earlier the relational shift. This provides support for a relational shift based on change in knowledge. But by itself it cannot be conclusive, for a global change of competence could also account for this decallage if we assume that the relations vary in complexity across domains. In this chapter, we pursue two lines of inquiry that can provide firmer support for the knowledge-change account. First, we show that children who initially fail to perform a relational mapping can show a pronounced gain in performance when given more knowledge. Second, we simulate the process of similarity matching, varying the degrees of knowledge assumed. We show that, using the same process, the model shifts from performing nonrelationally to performing relationally when given more knowledge.

STRUCTURE-MAPPING AND STRUCTURAL ALIGNMENT

To sort out these issues requires a theory of how analogy and similarity are processed. We propose that both children and adults compare mental representations via a structure-mapping process of alignment of conceptual representations (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1982, 1983, 1989). According to this view, the commonalities and differences between two situations are found by determining the maximal structurally consistent match between the representations of the two situations (Gentner, 1983, 1989; Gentner & Markman, 1994; Goldstone, 1994; Goldstone & Medin, 1994; Markman & Gentner, 1990, 1993; Medin, Goldstone, & Gentner, 1993). A *structurally consistent* match conforms to the *one-to-one mapping* constraint (i.e., an element in one representation corresponds to at most one element in the other representation) and to the *parallel connectivity* constraint (i.e., if elements correspond across the two representations, then the elements that are linked to them must correspond as well). When more than one structurally consistent match exists between two representations, contextual relevance and the relative systematicity of the competing interpretations are used. All else being equal, the match with the richest and deepest relational match is preferred (the *systematicity* principle).

Arriving at a maximally deep structural alignment might seem to require an implausibly discerning process, or even advance knowledge of the point of the comparison. But in fact, as the SME simulation makes clear, structural alignment can be realized with a process that moves from a rather blind stage of forming local (often inconsistent) matches to a stage in which deep, structurally consistent alignments are formed by capitalizing on connections between predicates (Falkenhainer, Forbus, & Gentner, 1986, 1989) . We describe SME in detail later; for now a sketch of the general process is sufficient.

SME takes as input propositional representations that consist of entities, attributes, functions and relations.[1] *Entities* are simply the objects in the domain and *attributes* are used to describe these objects. *Functions* are used to describe dimensional properties such as size and position. *Relations* are predicates that represent links between entities, attributes, functions, or other relations. Given two representations, SME operates in a local-to-global manner to find one or a few structurally consistent matches.[2] In the first stage, SME proposes matches between all identical predicates at any level (attribute, relation, higher-order relation, etc.) in the two representations. At this stage, there may be many mutually inconsistent matches. In the next stage, these local correspondences are coalesced into larger mappings by enforcing structural consistency: *one-to-one mapping* (each object in one representation is constrained to match to at most one object in the other representation) and *parallel connectivity* (matching predicates are constrained to have matching arguments). Although relations must match identically, SME allows correspondences between nonidentical functions or objects if they are arguments of matching relations. This allows a structural match to be made across nonidentical objects. It also permits a match across nonidentical dimensions (because dimensions are represented as functions). This is a way of capturing a psychological claim that dimensionalized attributes—that is, attributes such as size, weight, and brightness—that have been extracted as dimensions of a domain are particularly likely to participate in cross-dimensional mappings. (We return to the issue of dimensionalization in development later.)

SME then gathers these structurally consistent clusters into one or two global interpretations (called *GMAPS*). It then makes candidate inferences in the target. It does this by adding to the target representation any predicates that currently belong to the common structure in the base that are not yet present in the target. These function as possible new inferences imported from the base representation to the target representation. Finally,

---

[1] Relations and attributes are predicates taking truth-values, whereas functions instead map objects or sets of objects onto objects or values. However, for brevity we sometimes use the term *predicate* to refer to all three categories—relations, attributes, and functions.

[2] In this chapter, we describe SME in its Literal Similarity mode that allows attributes, functions, and relations to determine the best match between two representations. In order to model other tasks, SME can also be run in Analogy mode (relations only) and Mere Appearance (attributes only) modes. We think that the all-purpose Literal-similarity matcher is the better model of the normal similarity process.

the interpretations are given a structural evaluation. This evaluation is calculated by giving each matching predicate an evidence score. Predicates then pass a portion of their evidence to their arguments in a cascadelike fashion that implements a preference for systematicity (Forbus & Gentner, 1989). All else being equal, SME prefers a deep system of matches to a collection of isolated matches (even given the same number of correspondences). Thus, the process begins with local matches, but the final interpretation of a comparison is a global match that preserves large-scale structures.

Two immediate predictions arise from this model. First, carrying out a comparison process should promote a structural alignment. Second, the interpretation of a comparison emerges from the interaction of relational and object-based matches. A study by Markman and Gentner (1990, 1993) with adults illustrates these phenomena. We presented subjects with pairs of scenes like those displaying the *monotonic decrease* relation, as shown in Fig. 6.1. (Similar materials are used in the developmental studies presented later in this chapter.) The pairs were designed to contain cross-mapped objects—perceptually similar objects that play different roles in the relational structure of the two scenes (Gentner & Toupin, 1986). The task was a simple *one-shot mapping task*. We pointed to the cross-mapped object in one scene (e.g., Circle B on the left side of Fig. 6.1) and asked subjects to select the object in the other scene that best went with that object. They generally selected the perceptually similar object (Circle AA in the right hand configuration, which is the same size as Circle B). In the key experimental manipulation, other subjects were asked to rate the similarity of all the pairs of pictures prior to performing the mapping task. These subjects were significantly more likely to place the cross-mapped object in correspondence with the object playing the same relational role (e.g., Circle BB, the middle circle in the right-hand configuration) than the subjects who only performed a one-shot mapping. It seems that the simple act of carrying out a similarity comparison promotes a structural alignment.

The second prediction, that the winning alignment is a function of both

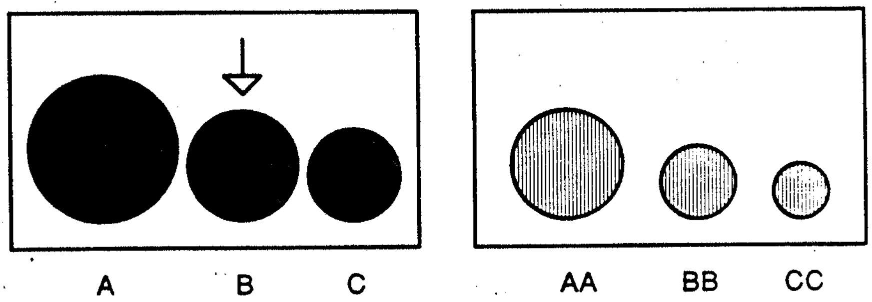

FIG. 6.1. Sample stimuli displaying the monotonic change in size relation like those used by Markman and Gentner (1990, 1993).

object similarity and relational similarity, was tested in a second study. This study made use of the fact that in a cross-mapping, the object-based interpretation and the relational interpretation are incompatible. Thus, increasing the similarity of the objects should increase the likelihood that the preferred match will be based on object similarities. Similarly, increasing the depth and strength of the relational match should increase the likelihood that the relational match will be the winner.

Subjects were again given pairs of pictures in patterns such as monotonic decrease (Fig. 6.1). For half the subjects, the pairs contained simple objects like circles and squares. The results were as given: Those subjects who rated similarity prior to making a mapping made more relational responses than subjects who simply did the mapping task. The other half of the subjects received richer, more complex objects, such as globes, scales, and faces, for which the pairwise object similarity was high. These subjects tended to make object-based mappings, even when they rated similarity prior to making their mapping. Although carrying out a comparison tends to induce a structural alignment, when object similarities are high relative to the degree of relational overlap, then object matches still may dominate the interpretation. This pattern is evidence that, even for adults, the process of interpreting a comparison involves both object matches and relational matches.

Simulating the Comparison Process

When we simulated these comparisons using SME, the results were consistent with subjects' performance. Before discussing the results, we describe SME's representational format. Although the specific details of a given representation are not crucial, the general assumptions are psychologically important. As stated previously, the model takes as input two propositional representations constructed from entities, attributes, functions, and relations. Within SME, these representational elements have strict definitions:

- *Entities* correspond to the objects in a domain.
- *Attributes* are unary predicates and are used primarily to describe independent descriptive properties of objects such as *thin*(Mary) or *short*(John).
- *Functions* (which unlike attributes and relations, do not take truth values, but rather, map objects onto objects or values) are used primarily to state dimensional properties like *size*(Mary) or *weight*(John).
- *Relations* are multi-place predicates that represent links between two or more entities, attributes, functions, or relations. We distinguish between *first-order relations*, which link objects (e.g., above[circle, triangle]) or their attributes (e.g., greater(*height*[Mary], *height* [John]) from *higher-order relations*, which link other relations (e.g., cause(buy[Mary, ball], possess[Mary, ball]).

Two points must be made about these representations. First, they are assumed to be *conceptual structures*, not verbal formulas, although for convenience words are used to label the nodes and predicates. For example, synonymous words typically have nearly identical representations. Second, these representations are intended as *psychological construals*. We are not aiming to capture the best, most complete, or most logical description of a situation, even if such a thing were possible. On the contrary, we assume that the same person may have alternate construals on different occasions, and that which comparisons are made and how they are interpreted depends on the person's current representation. The idea is to capture the processes of comparison between two situations as currently represented.

To simulate the Markman and Gentner results, we gave SME structural representations of two same-dimension monotonic-decrease stimuli (Fig. 6.2). These representations embodied four additional assumptions. Of these, the final two are of special interest, because later in this chapter we make different representational assumptions for children.

1. Similar objects are represented by nodes with similar sets of attributes. Thus, the cross-mapped objects in these scenes had identical size and shape attributes.
2. Complex objects have more attributes than simple objects. Thus, the object match between cross-mapped objects is stronger the richer and more complex the matching objects (Tversky, 1977).
3. We assume that adults conceive of certain attribute types such as size, darkness, and spatial position as dimensions, and we represent these as functions. (We suggest later that the extraction of dimensions is an important representational change during development.)
4. Finally, we assume that the adult notion of *monotonic decrease* can be represented as a higher-order relation connecting first-order pairwise comparisons (represented as greater-than relations) along dimensions. The higher-order relation expresses a coordination between monotonic change along two dimensions: for example, left–right position and size.

Representations for these stimuli were submitted to SME, run in general similarity mode as already described. All possible interpretations (GMAPS) were generated.[3] When the stimuli contained sparse objects, the winning GMAP—that is, the interpretation that received the highest evaluation of all those generated for this pair (evaluation score = 16.50)—corresponded to the relational match between scenes, as shown in Fig. 6.3a. In this

---

[3] In SME's normal (and most plausible) simulation mode, it produces only one or two best matches for a comparison (Forbus & Oblinger, 1990). However, SME can be run in exhaustive mode when, as here, one wishes to see all possible interpretations of a comparison and their evaluations.

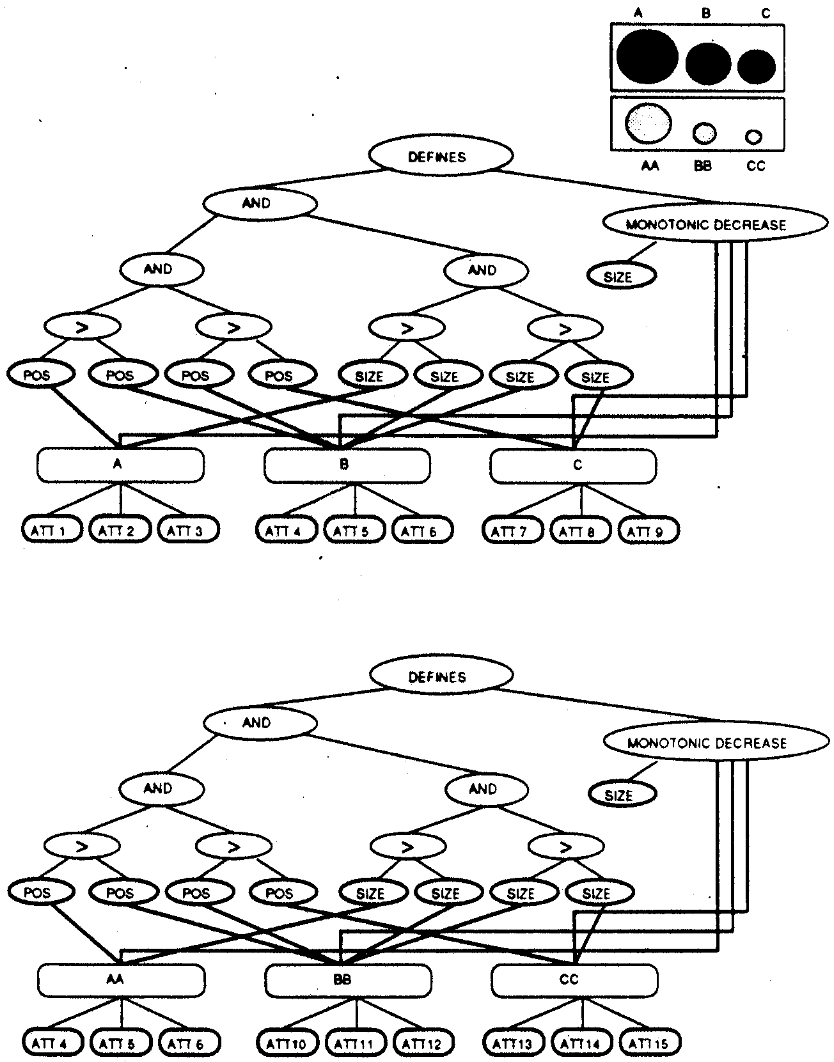

FIG. 6.2. Sample propositional representations from the simulation of the study by Markman and Gentner (1990, 1993). The representations shown here depict the adult representation of monotonic decrease.

interpretation the matching relational structure has been preserved: The objects have been placed in correspondence based on the similarity in their relational roles. Figure 6.3b shows the best ranked object interpretation for this pair. This interpretation aligns the matching attributes of the cross-mapped objects (Circle B with Circle AA), but shows very little matching relational structure. In the sparse-object simulation, the evaluation score for this object-based interpretation (evaluation score = 11.50), lagged

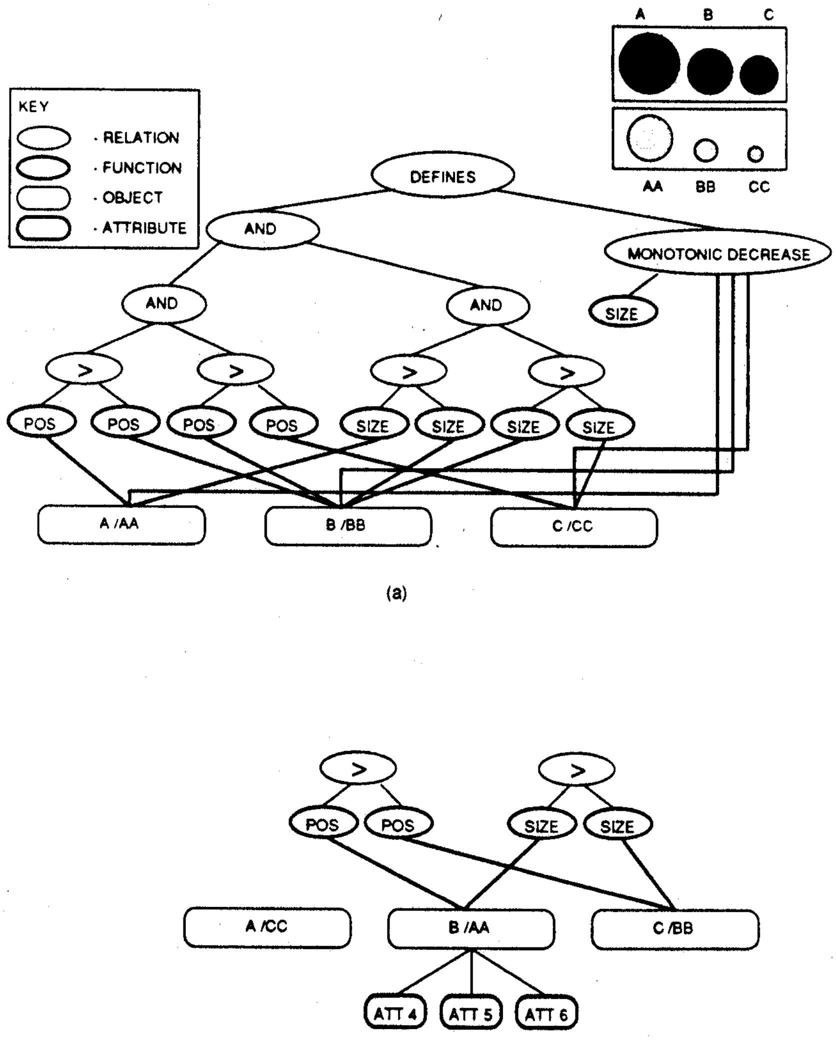

(a)

(b)

FIG. 6.3. Relation-based (a) and object-based (b) interpretations generated by SME when presented with the representations in Fig. 6.2. The relational interpretation received a higher evaluation score than the object interpretation.

behind that of the relational alignment shown in Fig. 6.3a. However, when the simulation was given richer object representations with greater object-matching potential, the cross-mapped-object-based interpretation overcame the relational match. This fits with the psychological findings concerning cross-mapped pairs: First, the comparison process induced a relational alignment in our subjects, which prevailed when the objects were sparse;

and, second, as object richness increased, so did the probability that the object match would win over the relational alignment. If the cross-mapping is removed, so that we have a simple analogy with dissimilar objects in the base and target, then SME, like adult subjects, readily makes the structural alignment.

STRUCTURAL ALIGNMENT AND THE RELATIONAL SHIFT

Now let us apply these ideas to development. We suggest that for children as well as for adults, similarity and analogy involve aligning two representations. When the object matches and the relational alignment are correlated, as in literal similarity, they are mutually supporting and there is one dominant interpretation. But when there is conflict, as in a cross-mapping, then whether the relational match or the object match will prevail depends on several factors. As with adults, the richer the object match, the more likely it is to prevail, and the larger and deeper the relational match, the more likely it is to prevail. Developmentally, these assumptions interact with considerations of change of knowledge. When children's domain theories are weak, their representations of the objects are likely to be much richer than their representations of the relations. As their knowledge of domain relations increases, children's relational representations become richer and deeper, increasing the likelihood that their comparisons will focus on matching relations. Thus, there occurs a relational shift: Children become able to carry out primarily relational matches. But, as evidenced in the adult results, the relational shift does not imply the disappearance of object similarity as a psychological factor. Rather, it refers to the *possibility* of making purely relational matches.

We have suggested that changes in knowledge drive the relational shift. In the following section, we consider two kinds of evidence for this claim. Our first line of evidence is a set of experiments that ask whether children who initially fail a relational mapping task are better able make a relational mapping after learning about the domain relations. If giving children knowledge of the relevant relational structure improves their performanceon a relational mapping task, then this constitutes an in principle demonstration that changes in domain knowledge, rather than maturation, are sufficient to account for the normally observed improvements.

A second way to gain insight into this phenomenon is to make use of a computer simulation of similarity processing. If we can simulate both younger and older children using the same processes and changing only the knowledge representations, this will support the change-of-knowledge account of the relational shift. The simulation allows us to ask what we never could about a real human: What would happen if we changed only the knowledge and not the processing?

We present two studies of mapping in children. Both use simple analogies based on perceptual relations such as monotonicity. This was done so that the relational structure could be grasped without presupposing extensive conceptual knowledge. Both studies follow the same logic. We first devise a task sufficiently difficult that young children normally fail to make a relational mapping. Then we try to create a change in their knowledge and ask whether they are then able to see the relational mapping. We consider two kinds of knowledge change: *structural enrichment*—adding higher order relational knowledge to initially shallow representations; and *re-representation* of initially holistic representations in a way that permits relational commonalities to be extracted. The first is a special case of *enrichment*, the second, of *restructuring* (Carey, 1985, 1991; Chi, 1981; Karmiloff-Smith, 1991, 1992; Norman & Bobrow, 1979). The studies we now describe provide evidence for both possibilities.

Finally, in addition to testing whether change of knowledge occurs, we wished to investigate how it occurs. Our studies further suggest two mechanisms or promoters of knowledge change: structural alignment and acquisition of relational language. We defer discussion of these until after we describe the results. For now, we list five predictions:

1. There will be a relational shift: Older children will perform more relationally than younger children.
2. Because the shift arises from knowledge, not age *per se*, giving children additional domain knowledge will lead to earlier relational mappings.
3. (Mechanism 1) The use of relational language can promote relational mapping.
4. (Mechanism 2) Because structural alignment induces attention to relational structure, the comparison process itself can promote relational mapping.
5. Because the comparison process is multiply constrained, both object similarity and relational similarity will affect performance.

If these five empirical predictions are borne out, this will support the claim that the relational shift is governed by changes in knowledge rather than changes in processing.

INVESTIGATING CHILDREN'S ACQUISITION OF RELATIONAL COMPARISON

Investigation 1: Solving Cross-Mappings

In these studies, we assessed children's ability to make relational mappings that varied in their predicted difficulty (Rattermann & Gentner, 1990;

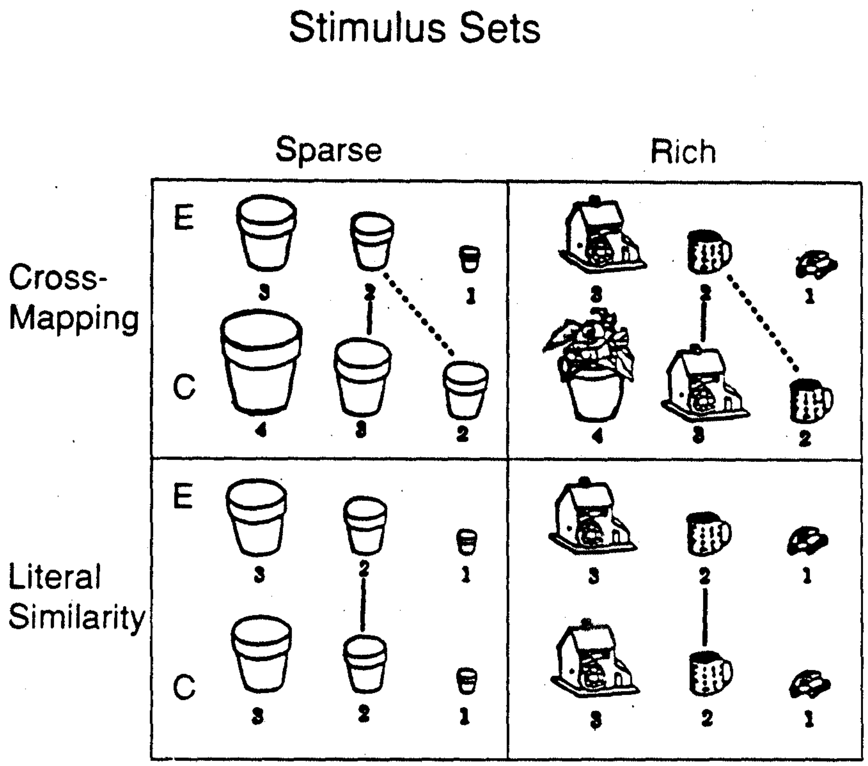

FIG. 6.4. Sample stimuli displaying the monotonic change relation from the mapping studies (Gentner & Rattermann, 1991; Rattermann & Gentner, 1990; Rattermann et al., 1994). Two of the objects are cross-mapped for both the (a) sparse and (b) rich pairs.

Rattermann, Gentner, & DeLoache, 1994). The easy pairs were literal similarity pairs, in which the object similarities supported the required relational alignment; the difficult pairs were cross-mapped pairs, in which the object similarities opposed the relational alignment (Gentner & Toupin, 1986). As in the study by Markman and Gentner described earlier, we varied the richness of the objects so that the interplay between object similarities and relational similarities could be examined.

Children were presented with two configurations of objects, each arranged according to the *monotonic increase (or decrease) in size* relation, operationalized as three objects in a row, increasing in size from left to right or right to left (see Fig. 6.4). One set of objects was designated as the child's (C) set, the other as the experimenter's (E). The child was asked to close his eyes while the experimenter hid stickers under one object in each set.[4] Then

---

[4] For clarity of presentation, we assume a male subject and a female experimenter, although in fact, roughly equal numbers of boys and girls participated.

the child opened his eyes and watched as the experimenter lifted one of her objects to reveal where the sticker was located in her set. The child was told that if he watched carefully, he could figure out where the sticker was hidden in his set. The rule was always the same: The child's sticker was hidden under the object that had the same relative size and position as the chosen object in the experimenter's set. To succeed, the child had to focus on common relational roles (e.g., the largest and leftmost object in both sets). If the child found the sticker on the first attempt, he was allowed to keep it. If not, he was shown where it was but was not allowed to keep it.

Within this basic task, two variables were manipulated: the richness and complexity of the objects and the mapping type (cross-mapping vs. literal similarity). Richness was manipulated by using either sparse objects, such as clay pots and blue plastic boxes, or rich objects, such as a pot of brightly colored silk flowers, a toy house, a colorful mug, and a toy car (see Fig. 6.4 for examples of rich and sparse stimuli.). The second variable was mapping type. Half of the subjects received literal similarity mappings, in which the object similarities suggested the same correspondences as the relational mapping. The other half received cross-mappings, in which the object similarities suggested different correspondences than did the relational mapping. Using a between-subjects design, we tested twelve 3-year-olds and twelve 4-year-olds in each of the four possible conditions formed by crossing stimulus richness (sparse vs. rich) with task type (literal similarity vs. cross-mapping). Additionally, twelve 5-year-olds were tested in each of the two richness conditions of the cross-mapping task.[5]

The predictions for the basic task were as follows. First, children in the literal similarity condition should perform better (i.e., should make more relational responses) than those in the cross-mapping condition, because in literal similarity, the object matches draw the children toward the correct relational alignment. We also expected an interaction between object richness and mapping type. In the literal similarity task, for which the object similarities supported the relational interpretation, we expected children to perform better on the rich stimuli than on the sparse stimuli. The reverse pattern should be found in the cross-mapping task, for which object similarity was in conflict with the relational interpretation. Because the rich object match should provide a more tempting competitor to the relational mapping than the sparse object match, children should perform worse on this relational task with the rich stimuli than with the sparse stimuli.

Figure 6.5 presents the proportion of correct responses across age in the literal similarity and cross-mapped conditions. Not surprisingly, the 4-year-olds were nearly perfect on literal similarity mappings (top graph), whether

---

[5] Because of the overall high performance of the 3- and 4-year-olds in the literal similarity task, the 5-year-olds were tested only in the two cross-mapping conditions.

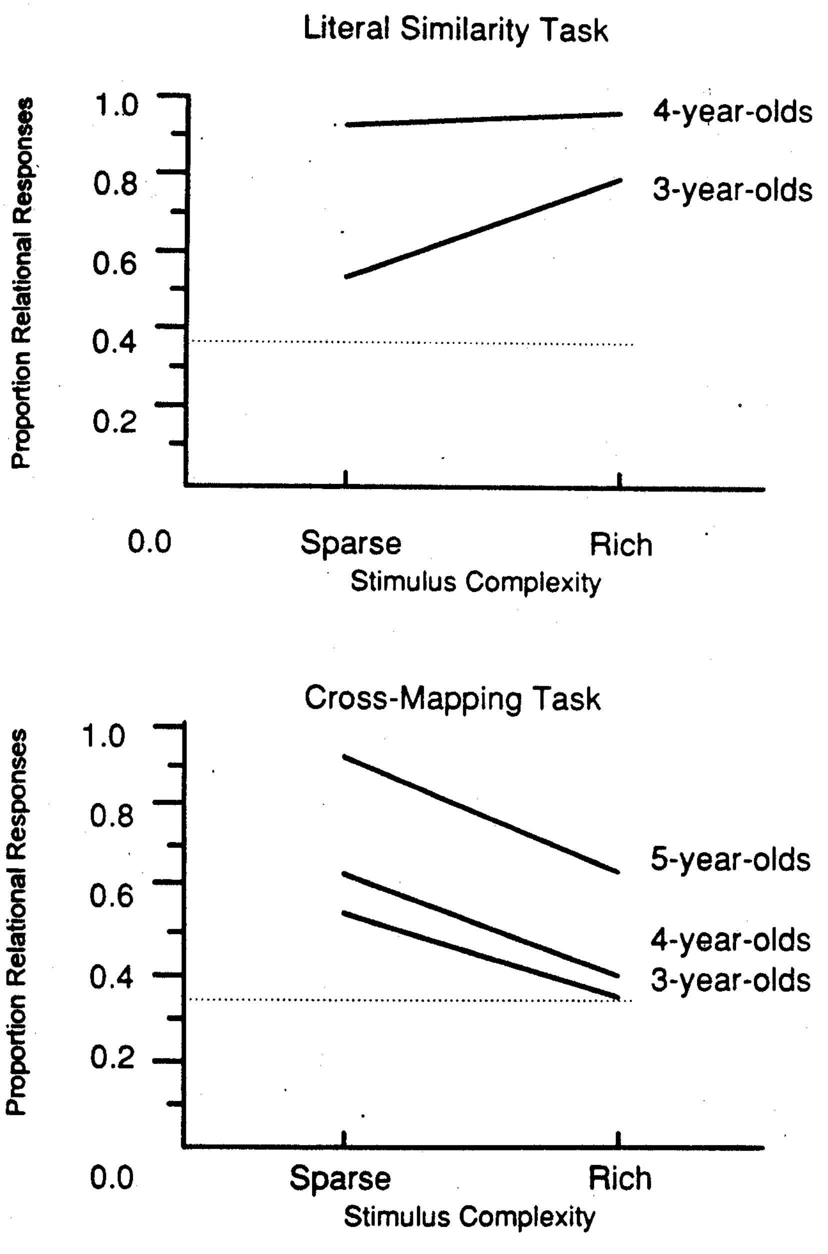

FIG. 6.5. Proportion of relational responses for different age groups in each condition of the mapping study (Gentner & Rattermann, 1991; Rattermann & Gentner, 1990; Rattermann et al., 1990).

they had rich stimuli (89% correct) or sparse stimuli (94% correct). However, the 3-year-olds benefited from object similarity, even in this "easy" condition. They made significantly more relational responses to the rich stimuli (85% correct) than to the sparse stimuli (55% correct). This pattern is striking, because even 3-year-olds might have been expected to find the literal similarity task easy even with sparse objects.

As predicted, the cross-mapped condition was harder (bottom graph). Also as predicted, the effects of object similarity in cross-mapped pairs were opposite to those in literal similarity. Both 3- and 4-year-old children had difficulty producing relational responses and the difficulty was greater for rich stimuli than for sparse stimuli. The results bear out the prediction that children should have trouble focusing on the matching relational structure in the face of competing object similarities. This conclusion is reinforced by the finding that children selected the identical object (rather than the object playing the same relational role) more often for the rich stimuli (42% for the 3-year-olds, 32% for the 4-year-olds) than for the sparse stimuli (23% for the 3-year-olds, 23% for the 4-year-olds). Consistent with the relational shift predictions, 5-year-olds were considerably better at maintaining a relational mapping despite cross-mapped objects, although they showed the same pattern of greater difficulty with rich objects (68% correct for rich as compared to 95% correct for sparse stimuli).

These orderly results underscore the joint contribution of object similarities and relational commonalities to the comparison process. Rich object similarities make relational responding easy for children when they suggest the same correspondences as the relational mapping (as occurs in the literal similarity condition), and hard for children when they suggest different correspondences from the relational mapping (as occurs in the cross-mapped condition). Sparse object similarities have weaker effects, giving mild support in literal similarity mappings and having mild negative effects on relational mappings for cross-mappings.

Parenthetically, the object richness effect here is a particularly clear instance of Tversky's (1977) self-similarity effect, whereby rich objects are more similar to themselves than are sparse objects. Self-similarity is theoretically important because it is counterevidence to a mental distance account of similarity, in which all identical pairs are at distance = 0. Because much of the prior evidence for self-similarity comes from perceptual confusion patterns in speeded tasks, the present evidence constitutes a significant broadening of the evidential basis for self-similarity effects, and hence, for a componential account of similarity over a mental distance account.

These results are consistent with the predictions of the structural alignment framework. Alignment and mapping involve interplay between local object similarities and global relational similarity. The more compelling the local similarities, the more they influence the overall relational mapping for better or worse. Furthermore, consistent with the relational shift hypothesis, with increasing age and experience, children become better able to focus on relational commonalities across a range of object similarities.

Can It Be Taught? The Effect Of Relational Language. The next study has two purposes. The first is to shed light on whether the improvement in relational performance across age is due to increases in knowledge, as we have suggested, or to maturational increases in processing capacity, as Halford's framework would predict. The second is to explore a more specific hypothesis, namely, our third prediction, that the acquisition of relational language contributes to the relational shift.

In this study, we again gave children the cross-mapping task, but this time we gave them labels for the higher-order relational pattern of *monotonic decrease*. In the previous study, some children had spontaneously applied the labels *Daddy, Mommy, Baby* to the objects (see also Smith, 1989). Because these terms seemed to apply to the monotonic decrease pattern, in this next study we taught twenty-four 3-year-olds to use these labels. We gave them "families" in which the largest object was labeled *Daddy*, the middle, *Mommy* and the smallest, *Baby*. The children received explicit training trials with labeled families of penguins and bears. For example, the experimenter pointed and said, "This is my Daddy, this is my Mommy, and this is my Baby. This is your Daddy, this is your Mommy, and this is your Baby. If my sticker is under my Daddy, then your sticker is under your Daddy." Then the children were tested on the same stimul – boxes and baskets in the sparse condition and houses, cars, and so on, in the rich condition – as in the first experiment. We had previously ascertained that children understood all the first-order relations between objects (the pairwise size comparisons); the question was whether the family labels would increase children's ability to appreciate the higher-order pattern by inviting them to import a familiar relational schema. If so, this would increase the level of relational responding.

The results of the labeling manipulation were dramatic: The 3-year-olds' performance in the cross-mapping task improved on both the sparse (89% relational responding) and rich (79% relational responding) stimuli. This is a substantial gain over their performance in Experiment 1 (54% and 32% correct, respectively). In fact, the 3-year-olds in this study performed at a level comparable to that of 5-year-olds, as though the children had gained 2 years of insight (see Fig. 6.5).

We might worry that this impressive performance depended on maintaining an artificially high level of explicit labeling. However, in a subsequent study we found that 3-year-olds given the label training could successfully maintain their performance without continued use of labels by the experimenters. In this test, after the first transfer trial, we gave children new stimuli and told them to arrange them "as they should be" so that we could play the game. Then we conducted the transfer task as in the original experiments, with no mention of the labels by the experimenter. We found that children who had received the label training did well. Of the 3-year-olds who had received training, 81% reached criterion with the sparse stimuli and 50% with the rich stimuli, as compared with the control (no training) group, of which 50% reached criterion with sparse stimuli and 12% with rich stimuli.[6] Interestingly, sparse objects remained easier than rich objects even under training. That the adverse effects of object richness persisted is consistent with our view that the relational shift is not all or none.

A second objection might be that *Daddy, Mommy,* and *Baby* do not name relational roles, but instead serve as a set of three object names. This interpretation contrasts with our suggestion that the task improvement stemmed from the fact that *Daddy-Mommy-Baby* labeled a higher-order relational schema for monotonic decrease. This possibility was tested in another study in which we taught children nonsense labels (such as *jiggy, gimli,* and *fantan*) that could only serve as pure object names. Their performance was not improved over the nonlabeled condition; in fact there was a tendency toward *worse* performance (Rattermann & Gentner, 1990; Rattermann, Gentner, & DeLoache, 1994). Thus, we suggest that the use of relational labels invited attention to the relation of monotonic change, making it more likely that young children would notice the matching relational structure. These results suggest the importance of possessing compact labels for relational patterns.

Finally, we note that we did not provide children with entirely new knowledge in this study. Rather, we invited an analogy that suggested that the stimuli could be viewed in terms of a relational structure the children already knew. We suspect this kind of cross-domain analogizing is a powerful force in development, and that it is often promoted by the use of common labels.

Overall, these findings are consistent with the view that representational change is the underlying mechanism of the relational shift. These data do not fit well with maturational stage theories. The children succeeded in a higher-order mapping at a mean age of 3 years, 6 months, slightly below Halford's hypothesized transition age of 4 years for multiple relation mappings, and far before Piaget would have granted the capacity for higher order relational mapping. More fundamentally, it is hard to see how any maturational theory could account for the radically different performance that occurred between two groups of the same age as a function of training.

---

[6] The dependent measure for this study was the number of children (out of 12 per group) who reached a criterion of four consecutive correct trials within the training set of 10 trials.

Yet the relational shift is not all or none. Rather, there is an interplay between the richness of the object matches and the depth of the relational matches. In cross-mapped trials, young children performed worse with rich object matches than with sparse object matches, and the reverse was true for literal similarity trials (for the same reason – namely, that the rich object matches were more alluring than the sparse object matches). Just as increasing object richness increased the strength of object-based interpretations, so increasing the amount and depth of the matching relational structure (by using relational labels) increased the likelihood of a relational interpretation. By the age of 5 years, the necessary relational knowledge was firmly in place, and this structure could be used as the basis of a mapping, even in the presence of a competing rich object match. However, even 5-year-olds performed better with sparse objects than with rich objects in cross-mapping trials. This pattern argues against the suggestion that "children would only solve analogies on the basis of object similarities when they were ignorant of the relations on which the analogy was based" (Goswami, 1992, p. 92). Rich object matches are perennially attractive.

These studies exemplify *structural augmentation* or *enrichment* of knowledge. Via the analogy with family relations, a higher-order pattern was added to the children's representation. We do not rule out the possibility that more radical restructuring may have taken place for some children; but this assumption is not necessary to account for the shift. We return to these findings later in this chapter, when we present a simulation of this comparison process.

Overall, these studies are consistent with the hypothesis that increases in children's relational knowledge play a significant role in their developing ability to match on the basis of relations. However, if our assumptions are correct, these results bear chiefly on augmentation or enrichment, albeit structural augmentation. We turn now to a set of studies that addresses the issue of *re-representation of knowledge*.

Investigation 2: Development of Cross-Dimensional Similarity

The previous section examined how adding higher-order relations to children's representations can increase their ability to make relational mappings within a domain. We now describe a series of studies that examine how re-representing the components of a domain can ease the determination of cross-dimensional similarities (Kotovsky & Gentner, 1990, 1994). In these studies, the difficulty for children lay in seeing patterns across different dimensions. We first describe the basic task and demonstrate an age shift in performance. Then we describe two manipulations that increase young children's ability to make these mappings.

We showed 4-, 6-, and 8-year-old children (12 children at each age) triads of figures and asked them to say which of two alternatives was most similar to the standard. The figures were sets of squares or circles differing in size and darkness. The standard was always constructed to fit one of two higher order perceptual relations: either monotonic change or symmetry. ***Monotonic change*** was operationalized as three objects in a line, identical except for the dimension of interest—either size or darkness—that increased (or decreased) steadily across the three objects. *Symmetry* was operationalized as three objects in a line with a central object flanked on either side by objects that were identical to each other. The middle object and the outer objects differed only along the dimension of interest—either size or darkness (see Fig. 6.6).

Although the child could select either response—there was no feedback—one of the two choices was always clearly more similar to the standard from the adult point of view. This alternative, the *relational choice,* depicted the same relational structure as the standard, but contained different objects. The other comparison figure (the *foil* or *nonrelational choice*) used the same objects as the relational choice, but these objects were haphazardly arranged so that there was no good higher-order relational structure. Thus, both alternatives matched the standard equally well at the object level, and the relational alternative matched better at the relational level, making the relational match preferable to anyone who recognized the matching rela-

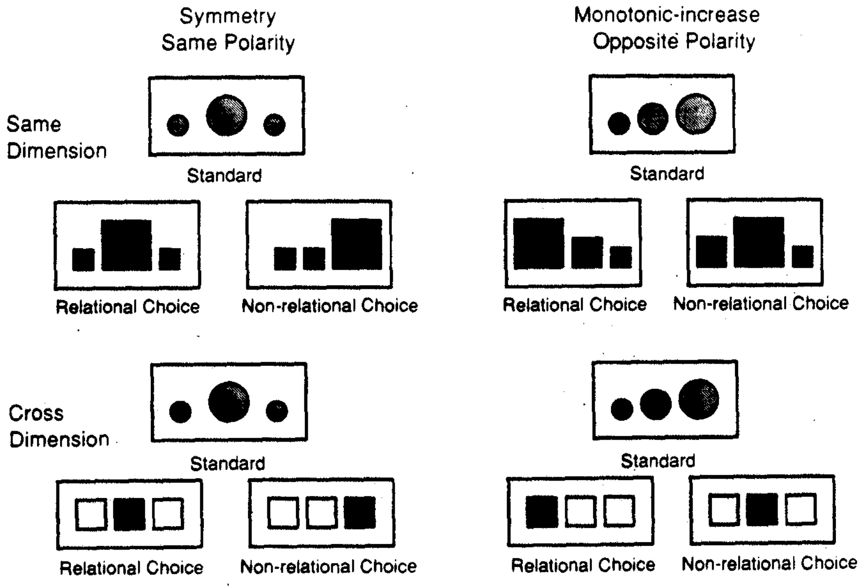

FIG. 6.6. Sample stimuli from the similarity triads task (Kotovsky & Gentner, 1990).

tional structure. (Adult subjects showed virtually 100% relational responding.) Thus, we can refer to the relational alternative as the *correct choice*, even though children were allowed to make whichever choice they liked.

The key manipulation was to vary the transparency of the mapping. The standard and the relational choice could match either within (e.g., size to size) or across (e.g., size to darkness) dimensions, and could either have the same or opposite direction of increase (polarity). That is, the design varied whether the polarity and dimension of variation matched between the standard and comparison figures. In *same-polarity trials*, first-order increases in the standard figure mapped to first-order increases in the correct figure (e.g., xXx and oOo). In *opposite-polarity* trials, first-order increases in the standard figure mapped to first-order decreases in the correct figure (e.g., xXx and OoO). In *same-dimension* trials, the comparison figures and standard both varied along the same dimension (e.g., size). In *cross-dimension* trials, the comparison figures varied along a different dimension than did the standard (e.g., darkness vs. size).

This design gave rise to four different types of trials: same polarity/same dimension, different polarity/same dimension, same polarity/different dimension and different polarity/different dimension (see Fig. 6.6 for a sample of each type). Let us consider the choice difficulty from the child's point of view. First, object matches were never of any use to the child, for in all the triads (even on the same-dimension trials) the correct response and the foil were equally similar (or dissimilar) to the standard at the level of objects. However, the four trial types varied in whether they could only be solved at a highly abstract level, or whether they could also be solved (that is, a choice could be made) on the basis of shared lower-order relations. The easiest trial type was the same polarity/same dimension type, for it could be solved with no recourse to higher-order structure. The child only needed to recognize the matching first-order relations to find the relational choice. For example, in the stimulus in Fig. 6.6, the second object is bigger than the first and the third objects for both the standard and the correct choice. Because these *bigger-than* relations are scrambled in the foil, the relational choice could be selected simply by matching the specific first-order relations. Knowledge of the higher-order symmetry relation was not needed.

The task was more difficult for opposite-polarity and cross-dimensional trials, where these low level relational matches were no longer available. For opposite-polarity trials, the child had to notice that the same overall pattern held in both items, even though increases in size or darkness in one array corresponded to decreases in size or darkness in the other array. Similarly, to make the relational response on cross-dimensional trials, children had to recognize that the overarching symmetry or monotonic-change relation dictated placing *bigger-than* relations in correspondence with *darker-than* relations. Because a deeper understanding of the relational systems was

TABLE 6. 1
Mean Proportion of Relational Responses by Age and Condition

	Same-Polarity		*Opposite-Polarity*	
Age	*Same-Dimension*	*Cross-Dimension*	*Same-Dimension*	*Cross-Dimension*
4	0.68	0.49	0.49	0.48
6	0.90	0.75	0.77	0.72
8	0.96	0.90	0.93	0.80

required to solve the opposite-polarity and cross-dimensional trials, we expected that children would make fewer relational choices on these stimuli than on same-polarity, same-dimension items.

In the first experiment, polarity match was a between-subjects factor and dimension match was a within-subjects factor (Kotovsky & Gentner, 1994, Experiment 1). Each child received 16 same-polarity trials or 16 opposite-polarity trials.[7] Table 6.1 shows the proportion of times children of different ages selected the (correct) relational alternative. The results support the predictions just described. The 4-year-olds seemed to choose randomly in all but the most concrete condition: They were reliably above chance only for the same-polarity, same-dimension stimuli. The 6-year-olds were above chance in all conditions, but, like the 4-year-olds, selected the relational choice more often on same-polarity, same-dimension trials than on any other type of trial. Finally, the 8-year-olds performed well in all conditions, although they too tended to make the fewest correct responses in the condition predicted to be the most difficult: the opposite-polarity, cross-dimension condition. These results are consistent with Chipman (1977) and Chipman and Mendelson's (1979) findings that perception of higher-order visual structure increases developmentally. These findings provide evidence for the predicted relational shift.

The Effect of Relational Labels. So far, we have evidence for a relational shift with age. As before, the central question is whether changes in knowledge and experience underlie the apparent age shift. In subsequent studies, we used only 4-year-olds and investigated whether training on the higher-order perceptual structure of these stimuli would improve their cross-dimension performance. We used only same-polarity triads. (Recall that 4-year-olds had performed at chance on even the same-polarity

---

[7] In a subsequent session children were shifted to the opposite-polarity condition (Kotovsky & Gentner, 1994, Experiment 1b). Children performed better on the second day when the task order was same-polarity to opposite-polarity than in the reverse order, consistent with our thesis here.

cross-dimensional triads.)[8] The first training task involved teaching labels for the relations (*more-and-more* for monotonic change, and *even* for symmetry). During the training task, children learned (with feedback) to classify the stimuli, one at a time, as to whether they were *more-and-more* or *even*.

After training, the children were given the eight cross-dimensional (same-polarity) triads, with the same similarity choice task as in Experiment 1. The 5 (out of 12) 4-year-olds who scored above criterion in the labeling and sorting task (75% correct categorizations and/or four productions of the labels) were well above chance on cross-dimensional trials (72% relational responding).[9] As in the Rattermann, Gentner and DeLoache (1990, 1994) studies, the use of relational labels increased children's attention to common relational structure. But whereas the children in the Rattermann et al. studies were given *Daddy, Mommy,* and *Baby* labels that could tap their existing schemas, the kind of training provided in the Kotovsky and Gentner (1990, 1994) task allowed children to build up the higher-order relational patterns for *more and more* and *even* over the course of the experiment. These newly reified relational patterns could then be more readily noticed and used.

In the Rattermann et al. task, the alignment is one of matching identical relational structure. The difficulty, of course, lay in ignoring the misleading cross-mapped object. The Kotovsky and Gentner task posed a different difficulty: namely, that of perceiving cross-dimensional commonality. To see the cross-dimensional similarity, children must align representations containing different first-order relations. For example, they must match monotonic change across *darker-than* relations with monotonic change across *bigger-than* relations. From the previous study, we know that common labels are one impetus to such a creative alignment. In the next study we investigated a different mechanism of change: repeated alignment itself. We have evidence from the Markman and Gentner (1993) studies that similarity comparisons promote structural alignment. We asked now

---

[8] In this and all subsequent studies, only the 16 same-polarity triads were used. Half were same-dimension (four size, four darkness) and half cross-dimension. All subsequent studies use 4-year-olds only (see Kotovsky & Gentner, 1994 Experiments 3-5). To avoid identity between the standard and the relational alternative, half of the darkness-change stimuli were blue circles, and the other half pink squares. For size-change, half were black-and-white patterned circles and half black squares.

[9] The results across all twelve 4-year-olds were weaker: 59% correct, only marginally significantly different from chance. However, we suspect that the children would have done better with more experience in the labeling and sorting task. (They were given only one pass through the cards.) A later training study with a more extended training regime produced strong results (Kotovsky & Gentner, 1994, Experiment 5).

whether prior experience in aligning same-dimension comparisons (which we know are accessible to 4-year-olds) would help them to see the relational structure necessary to align cross-dimensional comparisons.

Progressive Alignment. In this experiment, a simple change was made. The same set of triads was presented as in the first study, but the order was changed. Trials were blocked so that same-dimension trials were seen before cross-dimensional trials. Thus, children saw the easier trials before the more difficult ones. This blocking improved 4-year-olds' performance to 60% correct on the cross-dimension trials, marginally above chance. In comparison, 4-year-olds had achieved only 49% matching choices (chance performance) on cross-dimension trials when the trials were randomly mixed in Experiment 1. The difference becomes striking if we consider only the children who understood the same-dimension trials. When the same-dimension and cross-dimension trials were mixed in Experiment 1, even children who performed above the 75% criterion on the same-dimension trials were correct on only 48% of the cross-dimension trials (chance performance). In contrast, when the same-dimension trials were blocked initially, children who were at least 75% correct on the same-dimension trials ($n = 5$) went on to choose correctly on 80% of the cross-dimensional trials (significantly above chance).

This finding suggests that there is transfer from the easier same-dimension trials to the cross-dimensional trials when the same-dimensional trials are massed together. From other studies we have evidence that this is not a mere effect of task practice, for doubling the number of trials with size alone is not enough. Children must have concentrated experience in alignment within each of the two dimensions, size and darkness. Such repeated within-dimension alignments apparently potentiate subsequent cross-dimensional alignment.

How might this happen? The fact that young children initially fail to see the similarity among cross-dimensional comparisons suggests that such nonmatching relations are hard for them to align. We assume that the relations here are initially representing domain-specific manner (e.g., *darker than* and *bigger than*). Given two different domain-specific relations, some kind of re-representation is required in order to see these patterns as alike. The notion of re-representation to improve alignment is important in theories of analogy and case-based reasoning (Burstein, 1983; Falkenhainer, 1988; Gentner & Rattermann, 1991; Kass, 1989; Keane & Brayshaw, 1988) as well as in theories of conceptual development (Karmiloff-Smith, 1991).

What kind of re-representation might apply here? We speculate that children may initially view dimensional relations in a holistic manner. Specifically, we suggest that their representation of a difference in magnitude is typically conflated with the dimension of difference: for example, *darker*(a, b). Later, they re-represent these differences in a manner that separates the comparison and the dimension: for example, *greater*[*darkness*(a), *darkness*(b)]. Such a re-representation would make it possible to notice that there is some commonality between change in size and change in darkness. The idea is that extracting the specific dimensions from the relation of change along a dimension permits flexible cross-dimensional alignment. We examine this claim in more detail in the simulations to follow (see Fig. 6.8).

This proposal is in the spirit of the research of Smith, Kemler and their colleagues, who have demonstrated that the acquisition of adult dimensional structures is a lengthy process involving a shift from holistic to analytic processing (Smith, 1989; Smith & Kemler, 1977). Our claim is that repeated alignments, sometimes abetted by relational labels, help the child extract common structure. On this account, similarity comparisons contribute to the child's gradual disembedding (or decontextualizing, or desituating) of initially fused knowledge into separable representations (see Nunes, Schliemann, & Carraher, 1993). Also in accord with our thesis, Smith and Sera (1992) provided persuasive evidence that language learning contributes to children's learning to dimensionalize the world.

Children must not only separate perceptual knowledge into dimensions, but must come to see them as dimensions, as possessing a unified (often ordinal) structure. Once dimensions are extracted and represented, it becomes possible to grasp analogous structure across different dimensions. It is this kind of dimensionalization and alignment that permits humans to deal fluently with cross-domain metaphoric systems such as *up/down → good/bad* and the others discussed by Lakoff and his colleagues (Gibbs & O'Brien, 1990; Lakoff & Johnson, 1980; Turner, 1987, 1991).

All this suggests that children's early perceptual representations are conservative and context-specific and that they gradually develop dimensionally separated representations. The process of disembedding or desituating dimensions is promoted when children receive repeated opportunities to align the embedded dimensional structure; it is also promoted by learning common language that invites the extraction of dimensional commonalities. As the child gradually extracts the dimensions that apply within and across domains, cross-dimensional alignments become increasingly available. The child can see consistent mappings between structures across different dimensions.

This process of extracting dimensions is not smooth, as the work of Smith, Kemler, and Shepp showed convincingly (Shepp, 1978; Smith, 1984, 1989; Smith & Kemler, 1977). Smith (1989) showed that 2-year-olds do not possess anything like the adult notion of uniform dimensions within and across domains. For example, they fail to group according to like dimensions, and they fail to attend to dimensional identity in classification (Smith & Kemler, 1977). When asked to order three objects, a 2-year-old may be as pleased with *little, medium-sized, dark* as with *little, medium-sized, big*. Smith theorized that to the 2-year-old, *bigger* and *darker* both count as *more*. Overall, the picture that emerges from these studies, as well as from our own, is that between the ages of 2 and 5 years of age children come to see perceptual similarity in terms of what adults consider *like* dimensions. Clearly, children's ability to extract and attend to the dimension of variation in a given event is a necessary aspect of learning to group on the basis of like dimensions.

MODELING THE EFFECTS OF KNOWLEDGE CHANGE USING SME

The data from the developmental studies described suggest that knowledge change plays a crucial role in the development of relational sensitivity. In the studies by Rattermann, Gentner, and DeLoache (1990, 1994), the knowledge change seems to be an augmentation of children's domain knowledge, whereas in the studies by Kotovsky and Gentner (1990, 1994), the knowledge change seems to be a re-representation of existing knowledge structures. In this section we present simulations of the effects of both kinds of changes. The idea is to use SME to keep the process of comparison fixed and then vary the knowledge representation on which it operates. If the postulated changes in knowledge representation produce the observed changes in children's similarity performance, then we have evidence that change of knowledge could provide a sufficient explanation for the observed effects. Note that we are not simulating the re-representation process itself here, although such re-representation during analogy is an important aspect of our ongoing research (e.g., Falkenhainer, 1988). As earlier, we use the Structure-mapping Engine (SME; Falkenhainer, Forbus, & Gentner, 1986, 1989) in literal similarity mode to simulate the process of structural alignment and mapping. (The section surrounding Figs. 6.1–6.3 describes the simulation.) Our goals are, first, to achieve greater specificity in our discussion of representational change, and second, to discover whether change of knowledge is sufficient to produce the observed changes in children's performance.

Simulating Investigation 1: Structural Augmentation

In the Rattermann, Gentner, and DeLoache (1990, 1994) study, children saw two sets of objects and were asked to map from one set to the other.

There was a marked shift toward more relational mappings with age (and, in the later studies, with training).

Representational Assumptions. We made the assumption that the 3- and 5-year-olds differed primarily in whether their representations included the higher-order relational structure of monotonic change in size. The 3-year-olds were assumed to lack this higher-order relational pattern unless given training with relational labels. Thus, the following assumptions were made in the representations given to SME:

1. 1.For 5-year-olds, we assumed that the representations included higher-order relations of *monotonic change*: for example, *monotonic decrease in size*, as shown in Fig. 6.7.
2. For 3-year-olds prior to training, we assumed knowledge of only the first-order relations between objects, as shown in the dotted box in Fig. 6.7.

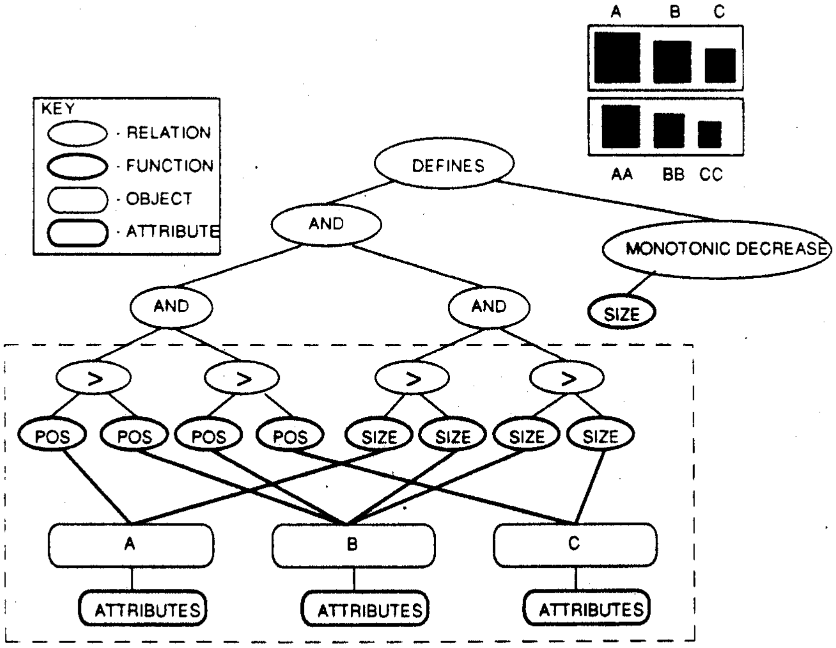

Knowledge representation used in simulation of developmental results. (Area within the dashed box is the novice representation.)

FIG. 6.7. Representations assumed for the older children (full representation) and the younger children (boxed section of representation) in the mapping study (Gentner & Rattermann, 1991; Rattermann & Gentner, 1990; Rattermann et al., 1994).

3. We assumed that 3-year-olds after training with relational labels have augmented their representations to include monotonic decrease, so that their representations are equivalent to those of older children.

Further Representational Assumptions. As in our simulations of the Markman and Gentner adult studies, we modeled the sparse and rich stimuli by varying the number of attributes each object possessed (its intrinsic richness) as well as the number of its attributes not shared by the other objects in its stimulus set (its distinctiveness).

4. Sparse objects were modeled with three attributes: two shared in common with the other objects in its set, and one that was distinctive.
5. Rich objects each possessed five attributes, none shared with any of the objects in its set (to capture the distinctiveness of the rich objects).
6. Literal similarity was modeled by placing object similarity and relational structure in synchrony in two representations. Objects playing the same relational role in the two representations were described by the same attributes.
7. The cross-mapping condition was modeled by placing object similarity and relational similarity in conflict. Thus the objects in matching relational roles were described by different attributes. The cross-mapped objects (which played different relational rules in the two representations) had identical attribute sets.

Simulation Results. The results of the computational simulation are comparable to the results found in the empirical work, as shown in Table 6.2. For *literal similarity* pairs, when the representations were local, a higher evaluation score was given to the relational mapping with the rich objects (15.5) than with the sparse objects (12.5). This result corresponds to the performance of the 3-year-olds, who made more relational responses to the literal similarity materials for rich stimuli than for sparse stimuli. When the representations included the higher-order relation of monotonic decrease, SME's evaluation scores for the relational mapping were higher than for the shallow representation, although the evaluation was still higher for the rich pairs (19.0) than for the sparse pairs (16.0). This result is consistent with the performance of the 4-year-olds, who responded relationally to the literal similarity stimuli, but still performed better with the rich items than with the sparse items.

For the *cross-mapping* task, the results of the simulation also paralleled the performance of the children in the experiments. For these pairs, SME

TABLE 6.2
Results of Simulation of Rattermann, Gentner and DeLoache Task

		Representation Type			
		Local Representation		*Higher-Order Relations*	
	Richness	*Object GMAP*[a]	*Relational GMAP*	*Object GMAP*	*Relational GMAP*
Literal Similarity	Sparse	12.50[b]	(same)	16.00	(same)
	Rich	15.50	(same)	19.00	(same)
Cross-mapping	Sparse	10.00	9.00	10.00	13.00
	Rich	11.50	8.00	11.50	11.50

[a] All values are GMAP evaluation scores generated by SME. [b] In Literal Similarity simulations, the object and relational GMAPs are the same.

generated both an object similarity interpretation and a relational similarity interpretation. When local relational structure was used (to simulate younger children), for sparse sets, the relational similarity mapping (evaluation = 13.0) received a higher evaluation score than the object similarity match (evaluation = 10.0). In contrast, for rich stimulus sets the object interpretation (evaluation = 16.5) received a higher evaluation score than the relational interpretation (evaluation = 8.0). This result corresponds to the empirical finding that 3-year-olds (and 4-year-olds) could make relational mappings for sparse stimuli, but not for rich stimuli.

When SME was given representations with higher-order relational structure for the sparse stimuli (simulating 5-year-olds and the 3-year-olds who were given familiar labels), the relational interpretation received a higher evaluation score than the object interpretation. For the rich stimuli, the relational interpretation and the object interpretation received the same evaluation score. This pattern is consistent with the behavior of 5-year-olds (and 3-year-olds given relational labels). On sparse object sets the deep relational mapping is clearly preferred, whereas on rich object sets there is a mixture of relational and object-based alignments.

The major conclusion from these simulations is that change of knowledge is sufficient to account for the relational shift. The same process model running on two different knowledge representations can simulate older children and younger children. When higher-order relational structure is included, SME's performance is like that of the older children, who readily master the relational mapping task. When the higher-order relational structure is removed, SME resembles the younger children, who fall prey to object matches and fail to master the task.

The simulations lend concrete support to our claim that the comparison process is an interaction of object matches and relational matches. In our simulations, as in the children's performance, object commonalities could either increase the likelihood of a relational response (in the literal similarity case, in which object commonalities supported the relational interpretation) or decrease the likelihood of a relational match (in cross-mapped pairs, for which object commonalities and relational commonalities were in competition). Only with deep relational representations did the relational interpretations prevail over rich cross-mapped objects. These simulations suggest that the relational shift need not involve a casting aside of objects. It results not from neglect of objects in the child's representations, but rather from an increase in the amount and depth of relational knowledge represented.

Investigation 2: Dimension-Specific and Dimension-General Representations

The results to be simulated from the Kotovsky and Gentner (1990, 1994) task are that (a) both younger and older children respond to within-dimension relational matches such as monotonic decrease in size, (b) older (but not younger) children spontaneously respond to cross-dimensional relational matches, and (c) younger children can be brought to notice the higher order cross-dimensional commonality if they are first given concentrated experience on both of the within-dimension comparisons.

Representational Assumptions. The ability to notice cross-dimensional commonality results from the recognition that dimensions like *bigger* and *darker* share some underlying similarity: They state that one value is somehow "more" than another along their respective dimensions. We speculate that the younger children represented magnitude difference in a dimension-specific way—roughly, *x is bigger than y*—whereas the older children represented magnitude difference in a dimension-general manner. Their encoding, we assumed, separates the specific dimension out of the magnitude comparison: *x's size is greater than y's size*. Note that the same information is encoded in both representations; the difference is in how analytical the representation is. The dimension-embedded encoding bigger(square1, square2) could be described as contextually embedded, conservative, or situated. The dimension-general encoding of the same relation greater[size(square1), size(square2)]. is more analytic than the first. It requires that the dimension of size has been extracted. This way, which dimension is affected can be separated from how it is affected (increase, decrease, and so on). We represent such extracted dimensions as functions: for example, *size*(x).[10]

SME treats these representations differently. Because SME matches

---

[10] In showing SME's representations, we represent relations in boldface and functions in boldface italics.

relations only if they are identical, it would see no correspondence between, say, bigger(x, y) and darker (a, b). In contrast, when given the second, more analytic encoding:

greater[*size*(a), *size*(b)]
greater[*darkness*(x), *darkness*(y)],

SME can readily perform a cross-dimensional alignment. Because SME allows nonidentical functions to correspond, two representations with identical relations over different dimensions can match. Thus, the second, more analytic encoding permits cross-dimensional alignment in SME, because the match between the two magnitude relations is then apparent.

Our representational assumptions in simulating the developmental change were as follows:

1. We assumed that 4-year-olds encoded the arrays in terms of dimension-specific relations (embedded relations) as shown in Fig. 6.8a, such as bigger(square1, square2) and darker(square1, square2).
2. We assumed that 4-year-olds encoded a dimensionally embedded change-in-size relation across the three objects (monotonic-decrease-in-size in Fig. 6.8a).
3. We assumed that 8-year-olds (and 4-year-olds after training) encoded the arrays in terms of dimension-general representations as shown in Fig. 6.8b, such as:

   greater[*size*(square1), *size*(square2)] and
   greater[*darkness*(square1), *darkness*(square2)],

   (where *size*(x) and *darkness*(x) are functions).
4. We assumed that 8-year-olds (and 4-year-olds after training) encoded a dimension-general relation of change (monotonic-decrease in Fig. 6.8b).
5. The objects were simulated in the same way as the sparse objects in the prior studies: namely, as having a few attributes shared with other objects in the figure, and one distinctive attribute (because only one dimension varied within each figure; see Fig. 6.8).

Simulation Results. When embedded relations were used, the relational interpretation received a much higher evaluation for the same-dimensional comparison (evaluation score = 11. 50) than for the cross-dimensional comparison (evaluation score = 7.50). This result is consistent with the finding that 4-year-olds performed significantly better on same-dimension trials than on cross-dimension trials. In contrast, when monotonic decrease

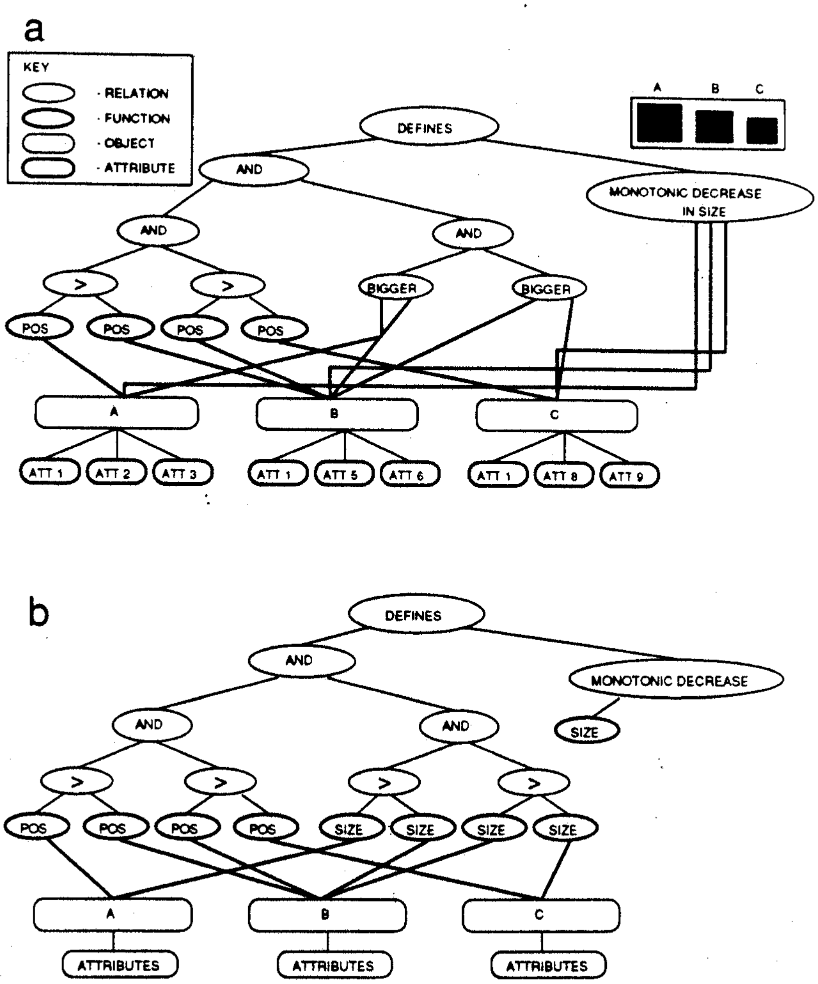

FIG. 6.8. Relational structures corresponding to the younger (a) and older (b) children in the triads study (Kotovsky & Gentner, 1990). The main difference between these representations lies in whether decrease along a dimension is represented holistically or analytically.

was represented using dimension-general encoding, the relational interpretation was preferred in both the same-dimensional comparison (evaluation score = 16.50) and the cross-dimensional comparison (evaluation score = 16.50). This result is consistent with the finding that 8-year-olds could spontaneously notice cross-dimensional comparisons.

The results of this simulation show that representational change alone is sufficient to account for the results presented here. The development of cross-dimensional comparisons can be modeled by assuming that children initially represent comparisons between values in a dimension-dependent fashion, but later re-represent these comparisons in a dimension-independent way. The simulation suggests that the dimensionalization of domains may be a key factor in acquiring the adult arsenal of reasoning tools. We return to this point later.

GENERAL DISCUSSION

The research summarized here suggests several conclusions. First, it supports the *career of similarity* thesis: Children begin with highly concrete similarity matches and gradually become able to appreciate partial matches. Second, among partial matches we see a relational shift from early focus on object-based matches to a later ability to perceive purely relational commonalities. Third, this development is driven by changes in domain knowledge. Fourth, we found support for several specific claims of the structure-mapping theory of the comparison process. On this account, the ability to make a purely relational alignment – an alignment not supported by, or even inconsistent with, the object similarities – requires that the structural alignment be rich enough and deep enough to prevail given the pressures from object similarity. We found that relational matches could be promoted in several ways predicted by the theory: (a) by diminishing the salience of competing object similarities (e.g., by using sparse objects in the Rattermann, Gentner, & DeLoache mapping studies); (b) by augmenting the salience of supporting object similarities (as in the rich-object literal similarity matches in the mapping studies); (c) by augmenting the depth of the matching relational structure (e.g., by adding higher-order relations, as in the mapping studies); (d) by highlighting the matching relational structure, as occurred in both studies; and (e) by re-representing two mismatching concrete relational structures so that they become instead two partially matching structures, as in the cross-dimensional task. We also considered two experiential forces that promote highlighting and re-representation: First, the acquisition of relational language and second, the progressive alignment of a series of cases so as to reveal common relations.

Language and Representational Change

These studies suggest that learning relational language fosters relational insight in several ways (see Gentner & Rattermann, 1991, for a review). One [illegible] allow children to transfer familiar relations to learning is that the use of a common label may lead children to search for relational similarities between two different situations, despite the presence of dissimilar perceptual attributes. That is, we suggest that alignment and re-representation can be invited through common language. Language may provide a kind of conceptual juxtaposition that is perhaps more powerful than the temporal juxtaposition that led to alignment in the within-to-cross-dimensional transfer studies.

Research on word learning has shown that words have the power to focus children's attention on commonalties among objects (Gelman & Markman, 1987; Gentner, 1978; Landau, Smith, & Jones, 1988; Markman, 1989; Markman & Hutchinson, 1984; Waxman, 1991; Waxman & Senghas, 1992). This research has demonstrated the powerful effect of labels on children's acquisition of concepts. For example, Markman and Hutchinson asked 2- and 3-year-old children to select another object that "goes with" a standard object. Given a choice of two objects, the children chose a thematically related choice (e.g., a web for a spider). However, if a novel word was used along with the task (e.g., "This is a dax, show me another dax") then children often selected the taxonomic choice (e.g., a spider for a fly). Thus, the label—even one whose meaning was not known in advance—oriented the child toward some notion of *like kind.* Although most of this research has focused on nouns referring to objects, we believe the same phenomena can occur with relations. That is, we believe that the use of a common term can serve as an invitation to find a common relation, provided the children have already learned names for the objects in the scene (Gentner, 1982; Markman, 1989). There is some preliminary evidence for this claim. Smith and Sera (1992) found evidence that children's acquisition of dimensional language influences their cognitive organization of the dimensions. In some pilot studies, Gentner and Wolff (1994) found that introducing a new relational term helps 5-year-olds solve relational analogies.

Progressive Alignment and Representational Change

Kotovsky and Gentner's results suggest that re-representation is a natural extension of the comparison process. We were able to increase children's performance on cross-dimensional trials simply by blocking the within-dimension trials before the harder cross-dimensional trials, and this despite the fact that children were given no feedback on their responses. We speculate that the within-dimension comparisons, being strong overall matches, are easy for children to perceive. Each time a pair of these dimensionally embedded relational structures is aligned, their common structure is highlighted. Thus, repeated experience on within-dimension pairs permits the child to extract deep common structures, which then form the basis for cross-dimensional alignment and re-representation. When these deep but dimensionally specific structures are juxtaposed in the cross-dimensional trials, the alignment process operates to promote re-representation of the comparison relations into a more domain-general format. We refer to this process as *progressive alignment*.

Our position that re-representation is central is clearly related to Karmiloff-Smith's (1991, 1992) theory of cognitive development, especially to her claim that initially implicit procedural knowledge becomes available explicitly through representational redescription. However, there are some differences between the accounts. First, the scale of the changes discussed here is more local than in Karmiloff-Smith's discussion. Whereas Karmiloff-Smith emphasized metalevel insights into one's own processes, we see a role for re-representation even at the simple content level. Second, our research focuses particularly on *mechanisms* of re-representation. We believe this level of explanation is crucial to understanding the phenomenon. Third, whereas Karmiloff-Smith proposes that redescription processes begin only after behavioral mastery is attained in a given domain, we assume that alignment and re-representation happen from the start. The reason that re-representation does not seem to be occurring in very young children is that their earliest representations are so richly embedded in concrete detail, and so lacking in higher-order abstractions, that only the most conservative similarity matches can be made. These early matches and their resulting inferences may go unnoticed; they are far too simple to be the kind of insights that parents proudly relate. Nonetheless, we suggest that they pave the way for the more dramatic comparisons to come.

The Early Conservativeness Of Similarity

We have stressed that early similarity matches tend to rely on massive overlap between the items, and that only with experience do children become able to appreciate partial similarity and analogy. Our current findings support this conclusion. In both studies, children could respond well to literal similarity before they could take advantage of purely relational commonalities. One implication of the "career of similarity" thesis is that children's earliest similarity matches should be highly conservative; that is, they should rely on extremely large overlap. Our survey of the development of similarity turned up considerable evidence for this claim (Gentner & Rattermann, 1991). For example, Baillargeon found a fascinating ability in infants to perform a rudimentary kind of inferential mapping, but only under conditions of near identity (Baillargeon, 1987, 1990, 1991; Baillargeon, Spelke, & Wasserman, 1985). Normally, 4½-month-old infants who have been habituated to a screen rotating back and forth through an 180° arc show no surprise when a solid box is placed behind the screen and in the path of its trajectory, and is (apparently) crushed into a tiny fraction of its former size. (Note that the apparent crushing of the box takes place behind the screen and out of the infant's line of sight.) However, if another box of the same size and shape is placed next to the to-be-crushed box, the babies show surprise at the crushing event, provided that this second box (which remains visible throughout the event) is identical or highly similar to the first "to-be-crushed" box. For example, given a visible box that was red with white dots, the 4½-month-olds could successfully make the mapping (and thus show surprise) if the "crushed" box was red with green dots, but not if it was yellow with green dots or, worse, yellow with a clown face.

We interpret this finding as suggesting that the babies are doing a kind of similarity-based mapping, using the box that is visible to infer (or remember) the size of the occluded box as it disappears behind the screen (Gentner & Rattermann, 1991). What is striking is the conservativeness of the process. The babies appear to require a strong overall similarity match before they can make the match. Results like these bring home the magnitude of the human achievement in acquiring the kind of flexible, purely relational similarity capability that adults take for granted. Thus, the development of similarity proceeds from the perception of overall similarity between two situations to the ability to perceive partial similarity matches, and among these partial matches, object-matches precede relational matches.

Relation to Other Views

The Primacy of Relations View. We have discussed the positions of Brown and Goswami under the general rubric of the knowledge-based account of the relational shift. However their position is somewhat more complex than this. Goswami and Brown stress the early availability of relational similarity, and at times seem to argue that there is no relational shift: Relational similarity is dominant from the start (Brown, 1989, 1990; Goswami, 1992; Goswami & Brown, 1990). Clearly, our results do not support such a position. The results of both studies show a strong relational shift with experience. In the Rattermann, Gentner, and DeLoache mapping studies, there was a shift from object matches to relational matches in the cross-mapping studies: Children showed increasing dominance of relational similarity with increasing knowledge of relations. Furthermore, children performed better when the object matches were consistent with the relational alignment, and worse when they were inconsistent. In the Kotovsky and Gentner studies, young children needed concrete same-dimension matches; only children with more domain experience (acquired over time or by training) could appreciate cross-dimensional matches sharing purely higher-order relational similarity. These findings, which demonstrate a relational shift in development, are hard to reconcile with the claim that relational mapping behavior is fully present from the start.

However, this apparent disagreement may be partly terminological. Goswami at times seems to adopt a very liberal criterion for the use of the term *relational similarity*. For example, as evidence that infants use relational similarity, Goswami (1992) cited Meltzoff's (1990) finding that 14-month-olds will watch an adult who is imitating their current behavior (e.g., shaking a toy when the infant shakes a similar toy) in preference to one who is imitating their past behavior. Meltzoff's intriguing evidence certainly suggests that babies are aligning their own actions with those of the adult. But to describe this match as relational similarity obscures the fact that the two events—*adult shaking toy* and *baby shaking [similar] toy*—match closely at the object level as well as at the relational level. Such a match fits the description of *overall* or *literal similarity*. Again, in her reanalysis of Baillargeon's occluded box example discussed earlier, Goswami (1992) referred to babies' use of a near-identical standard as a *relational comparison* and cited this study as evidence that even young infants can perceive relational similarity. We would term the likeness between a red box with white dots and a red box with green dots one of literal similarity, and not one of purely relational similarity.

If Goswami is using the term *relational similarity* to subsume both literal similarity and analogy, then there would be no disagreement as to its primacy. The claim that overall similarity can be perceived from the start coincides with the claims of the relational shift hypothesis. However, to equate relational similarity and overall similarity would seem to render discussions of the later development of purely relational similarity rather cumbersome. If, on the other hand, Goswami's claim is that an appreciation of purely relational similarity precedes an appreciation of object-based similarity, then this view is quite distinct from our own, and is countermanded by the results here.

A related but distinct position on early relational similarity is Bryant's (1974) thesis that relational similarity precedes absolute similarity. Bryant pointed out that when young children are given comparisons like "5 is greater than 3," they find it easier to match on the basis of the dimensional relation than on the basis of the absolute values. The results of the Rattermann et al. mapping studies point the way to a reconciliation between Bryant's findings and the relational shift claim. Absolutes are values along one dimension only, and hence are extremes in the direction of sparseness. Thus, for absolutes, as for the sparse objects in our experiment, the pull of object similarity is very low (i.e., self-similarity is very low). For when the objects differ only along one dimension, say size, then in a cross-mapping, say *5 2* and *2 1*, the best object match (*2* → *2*) is only one feature better than the object match (*2* → *1*) required to support the relational mapping.

Thus, even with a beginner's knowledge of relational structure, relational similarity may win out over object similarity when the objects are absolutes. In contrast, rich objects like the ones in our experiment have high self-similarity: This means that the $2 \rightarrow 2$ object match is much stronger than the $2 \rightarrow 1$ object match. Thus, the relational mapping becomes less likely as the richness of the objects (and therefore the degree of object similarity) increases. It becomes more likely either as the objects become less interesting or as the relational structure becomes deeper and more salient. Consistent with this account, when a focus on relations is desired—for example, when teaching mathematics—sparse objects such as x's and y's are often preferred (Uttal & DeLoache, 1994). According to this account, we can expect to see a relational shift in most domains. Relational similarity gains ascendancy with increasing knowledge.

The relational shift is not absolute: It does not represent a shift from using exclusively object similarity to using exclusively relational similarity. Adults can use purely relational matches, and tend to find them more sound and apt than object matches (Gentner & Clement, 1988; Gentner, Rattermann, & Forbus, 1993), but they use overall similarity matches whenever possible. Furthermore, when objects and relations are pitted against one another we find effects of both kinds of similarities in studies of comparison (Gentner & Markman, 1993; Goldstone, 1994; Goldstone & Medin, 1994; Markman & Gentner, 1990, 1993; Medin, Goldstone, & Gentner, 1993), in problem-solving research (Novick, 1988; Ross, 1984, 1989) and in device transfer (Schumacher & Gentner, 1988). The comparison process typically involves both object similarity and relational similarity throughout development.

The Global Change View. These results argue against theories that propose that relational similarity must await some advanced stage of cognitive development (Piaget, Montangero, & Billeter, 1977). They accord with the position that children should be seen as domain novices rather than as underdeveloped information processors (Brown & DeLoache, 1978; Carey, 1985; Chi, 1981).

Our findings are also problematic for Halford's account, according to which children's performance on relational analogies is governed largely by maturational increases in their processing capacity (Halford, 1987, 1992, 1993). Halford's account of the relational shift is that children's ability to perform analogical reasoning tasks depends on whether their processing capacity is equal to the structural complexity of the task. Tasks that require only unary predicates, such as object attributes, make fewer demands than those requiring the mapping of a binary relation. Tasks whose solution requires mapping systems of relations induce an even higher processing load. Halford and his colleagues have carried out many closely reasoned studies of topics such as transitivity and class inclusion that demonstrate the expected shift in performance with age. They have simulated their view in the STAR model, which represents relational structure in a distributed connectionist system using tensor products (Halford, Wilson, Guo, Wiles, & Stewart, in press; Smolensky, 1990).

Halford's emphasis on a maturational increase in processing capacity contrasts with our claim that the relational shift derives from the acquisition of higher-order structure. In part this difference derives from a difference in the ways the two theories represent complex relational structure (Gentner, 1992). Halford models structures such as monotonic increase as consisting of multiple relations between objects. The idea then is that greater processing capacity is required to find correspondences between representations containing multiple relations. In contrast, we represent complex relational patterns in terms of higher-order relations that take lower-order relations or propositions as arguments. Because the human comparison processor (like SME) favors connected systems, resolving a similarity comparison is easier when there are higher-order relations connecting lower-order relations that would otherwise be independent. Simulations using SME bear out this claim. SME is faster and more certain of the best match when given deep representations than when given flat representations. Thus, increases in domain knowledge can actually decrease the processing load associated with relational tasks in that domain. This argument is cousin to the demonstrations of Bower and Winzenz (1969), Bransford and Johnson (1972), Mandler and Mandler (1964) and others to the effect that adding more connecting material can make a task easier.

Our results that children can shift from object-based similarity to relational similarity in the space of a few minutes run contrary to the claim that the relational shift stems from a maturational increase in processing capacity. In particular, the fact that 3-year-olds who are taught higher-order relations can then use them to perceive relational similarity runs against Halford's suggestion that the ability to process relational similarity is acquired at about 3½ or 4 years of age. However, Halford's account also allows for the effects of learning; for example, chunking of relations into larger relations. Although our account stresses learning more and maturation less than his view, there is considerable agreement between the two views.

SME and Re-representation

The SME simulations serve a dual purpose. First, they demonstrate that the observed developmental changes in ability to make relational mappings can be mimicked by holding the similarity process constant and changing the representations it operates over (and, moreover, changing the representations in ways consistent with our task interventions). Second, the simulations suggest that the process of comparison might itself be part of the learning process, in two respects: (a) repeated alignment might focus a learner on large common structures, and (b) finding aligned mismatches provides candidates for re-representation. Structural alignment thus acts as a domain-general process that is sensitive to domain-specific information.

We have discussed re-representation (e.g., decomposing two similar predicates to reveal some identical subpredicates) as a key mechanism in progressive alignment. However, another possibility is suggested by the ACME simulation of Holyoak and Thagard (1989). Rather than requiring identical relations, it uses a similarity table: The more similar two predicates, the better their match.[11] In this approach, learning to make cross-dimensional mappings could be done simply by entering a high similarity for the two relations; the child would just learn that *darker* is similar to *bigger*, and so on. Although this approach seems less cumbersome than the re-representational approach proposed here, we believe it is psychologically incorrect, for two reasons. First, *darker* is not in fact similar to *bigger*. The child would be mistaken to suppose that a dark horse would resemble a big horse, for example. Second, entering a local similarity value for these two predicates misses the fact that they correspond by virtue of a larger domain mapping. If in a given context *darker* corresponds to *bigger*, then *lighter* must correspond to *smaller, pale* to *small*, *darkening* to *growing*, and so on. All these would have to be given high values in the similarity table. It is difficult to conceive of a mechanism that could accomplish this without losing the simplicity that made the similarity table attractive in the first place. Moreover, if the table could be changed to reflect these domain mappings, we would face an even worse version of the first problem. Our similarity metric would tell us that a big horse is like a dark horse *and* that a pale horse is like a small horse, and so on.

The problem is made more acute by the large number of system-mappings in common use: *dark/light* → *sad/happy, dark/light* → bad/good, dark/ligh → *confusing/clear*, and so on (Lakoff & Johnson, 1980; Turner, 1987, 1991; see also Kittay, 1987). To account for all these mappings within a similarity table would render the table incoherent and factually incorrect. SME's technique avoids these problems. It does not try to capture dimensional mappings by setting the pairwise similarity of individual predicates, but by representing them as system mappings in which nonidentical functions representing dimensions (like *darkness* and *size*) correspond by virtue of their roles in the larger matching structures. Once two dimensions are placed in correspondence, the mapping can readily be extended: *increase*

---

[11] In contrast, SME's decompositional approach captures degree of similarity through the number of overlapping representational components. Later versions of ACME have explored other similarity algorithms.

in darkness (*darkening*) corresponds to *increase in size* (*growing*), and so on. This generativity is a significant advantage of our dimension mapping approach over the use of similarity tables.

In our simulations we did not simulate the re-representation process itself. This is because our aim is to show that change of representation is sufficient to account for the observed developmental changes. However, simulating the re-representation process has been the focus of considerable recent research in analogy (Burstein, 1983; Falkenhainer, 1988; Gentner & Rattermann, 1991; Kass, 1989; Keane & Brayshaw, 1988). Falkenhainer's (1988) Phineas system, which uses contextual structure-mapping to model scientific discovery, re-represents predicates to improve the alignment under carefully specified conditions. For example, it tries to find a common superordinate for two nonmatching antecedents whose match is invited by the overall alignment, provided that their consequents match. Keane's Incremental Analogy Matcher (IAM) models the analogical process as one of iterative mapping, with later mappings incorporating more information (Keane & Brayshaw, 1988). Kass (1989) described a set of "tweaks" by which explanations of prior cases can be adapted to apply to a current situation. A complete model of the developmental changes described here will include the mechanisms of re-representation itself. Once two dimensions are placed in correspondence, the mapping can readily be extended. This extendability and generating is a significant advantage of a dimension-mapping approach over the use of similarity tables.

Although we have emphasized the importance of domain knowledge, other factors may also play a role in the development of comparison ability. As Klahr and Wallace (1976), Siegler (1984), and Sternberg (1984) have emphasized, it is unlikely that one explanation will cover all of cognitive development in an arena this size. One possibility that is consistent with structure-mapping theory is that children do not initially share the adult preference for structurally consistent and systematic mappings. That is, children may only gradually develop the preference for higher-order relational matching. The possibility that a preference for systematicity and structural consistency is culturally influenced is supported by an examination of the writings of medieval alchemists. Their aesthetic was different from the modern one; it encompassed rich, structurally inconsistent analogies and many-to-one mappings (Gentner & Jeziorski, 1989, 1993). Recent connectionist models of analogical mapping also suggest ways to capture a lack of structural consistency. Holyoak and Thagard's (1989) ACME and Goldstone and Medin's (1994) SIAM both use localist connectionist networks to determine the best match between scenes. In both systems, the one-to-one correspondence rule, for example, is a pressure rather than a firm constraint. Thus, another developmental shift worth examining is a potential shift in the firmness of the constraints on alignment and mapping.

Cross-Domain Mappings

We speculate that analogical mapping between domains may be a major mechanism of learning and discovery in the developing child (see Gentner & Rattermann, 1991; Halford, 1993; Siegler, 1989) as in the scientist (Gentner, 1982; Gentner & Jeziorski, 1993; Nersessian, 1992). We suggest that alignments, sometimes abetted by common relational labels, contribute to the child's gradual analysis (or disembedding, or desituating) of initially context-bound knowledge into separable representations and help the child see common structure across different dimensions (see Nunes, Schliemann, & Carraher, 1993). The gradual dimensionalization of the child's world brings with it the ability to align structure across different domains. The ramifications of this representational change are vast. As mentioned earlier, the mapping of structure across different domains underlies the rich set of cross-domain metaphoric systems that pervade our language (Carbonell, 1982; Gentner & Boronat, 1992; Gibbs & O'Brien, 1990; Kittay, 1987; Kittay & Lehrer, 1981; Lakoff & Johnson, 1980; Nagy, 1974; Turner, 1987, 1991). Such cross-domain analogies can occur seemingly unconsciously, as when Bowerman's (1981, 1982) preschool child asked "May I have some candy behind dinner?" (a possible time–space analogy). In other cases, the analogy is noticed, as when a child in our cross-dimensional task announced with delight, "It's exactly the same, but different."

The discarding of Piaget's global stage system in favor of a domain-knowledge view of cognitive development threatens to leave us with a piecemeal account, one that lacks any link between, for example, conservation of volume and conservation of weight. We speculate that analogy provides that link. The child who has caught on to conservation in one or two prior domains is more likely to learn the principle in the next domain. For example, in an intriguing study, Gelman (1969) taught 5-year-olds, who initially failed to conserve length, number, mass, and liquid, a discrimination learning task with length and number. Their subsequent conservation performance was near perfect on length and number; but more impressively, the children also improved substantially on conservation of the two nontrained quantities, mass and liquid amount. In another study, Gelman (1982) taught children conservation of small numbers and found that they subsequently improved their performance on tasks involving conservation of large numbers. Simon and Klahr's (chap. 7, this volume) simulation of this finding using their Q-Soar simulation further demonstrates how knowledge of conservation can transfer from small numbers to large numbers. Consistent with our transfer-of-knowledge account, Simon and Klahr suggest that an understanding of discrete numbers provides the basis for learning to reason about continuous quantities.

We suggest that there is a kind of mutual promotion cycle, whereby analogy and similarity act to increase representational uniformity (through re-representation to increase alignment), and are in turn promoted by uniform representations (because the more alignable the representations, the more likely it is that the likeness will be noticed and the comparison made). This positive feedback cycle contributes to what we have called the *gentrification of knowledge* (Gentner & Rattermann, 1991)—to the gradual replacement of the idiosyncratic perceptions of childhood by the sturdy, relatively uniform representations of the adult cultural world view.

SUMMARY AND CONCLUSIONS

We have focused on two ways in which children become fluent at higher-order relational comparison. First, as they gain information about higher-order relations in a domain, they become better able to make complex relational mappings, even in the face of cross-mappings that may distract them from the relational correspondences. Second, information they already possess may be re-represented to determine deeper similarities.

A number of forces drive these representational changes. One force is the comparison process itself. Our results suggest that simply carrying out similarity and analogy comparisons may play a fundamental role in the development of representations. Alignment of structure may focus the child on a limited number of areas where knowledge enrichment or re-representation is likely to be fruitful. Although similarity is often treated rather slightingly in current theories of cognitive development, these results suggest that similarity—even mundane within-dimension similarity—can act as a positive force in learning and development.

A second force is language. Language provides names for abstract relational structures, reifying complex information and making it easier to manipulate. By applying familiar labels in a new domain, children may learn to transfer relational structures learned in one situation to novel circumstances. Finally, language and comparison may act in concert: The conceptual juxtaposition and alignment invited by common language can lead children to form relational categories.

ACKNOWLEDGMENTS

This chapter grew out of a symposium presented by the authors at the third annual Midwestern Artificial Intelligence and Cognitive Science Society conference in Carbondale, IL, 1990. The developmental research was supported by NSF grant BNS-87-20301 and by The Center for the Study of Reading at the University of Illinois at Urbana-Champaign Grant 400-31-0031. The development of the Structure-Mapping Engine was supported by ONR contract N00014-89-J1272. We thank Ken Forbus, Doug Medin, Judy DeLoache, David Uttal, Phil Wolff, and Ron Ferguson for insightful discussions on these issues, and Tony Simon for helpful comments on the manuscript.

REFERENCES

Baillargeon, R. (1987). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology, 23,* 655–664.

Baillargeon, R. (1990). *The role of similarity in infants' use of visible objects as cues for hidden objects.* Unpublished manuscript.

Baillargeon, R. (1991). Reasoning about the height and location of a hidden object in 4.5- and 6.5-month-old infants. *Cognition, 38,* 13–42.

Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition, 20,* 191-208.

Billow, R. M. (1975). A cognitive developmental study of metaphor comprehension. *Developmental Psychology, 11,* 415–423.

Bower, G. H., Clark, M. C., Lesgold, A. M., & Winzenz, D. (1969). Hierarchical retrieval schemes in recall of categorized word lists. *Journal of Verbal Learning and Verbal Behavior, 8,* 323-343.

Bower, G. H., & Winzenz, D. (1969). Group structure, coding and memory for digit series. *Journal of Experimental Psychology Monograph, 80*(2, Pt. 2).

Bowerman, M. (1981). The child's expression of meaning: Expanding relationships among lexicon, syntax and morphology. In H. Winitz (Ed.), *Native language and foreign language acquisition* (Vol. 379, pp. 172-189). New York: New York Academy of Sciences.

Bowerman, M. (1982). Starting to talk worse: Clues to language acquisition from children's late speech errors. In S. Strauss (Ed.), *U-shaped behavioral growth* (pp. 101-145). New York: Academic Press.

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior, 11,* 717-726.

Brown, A. L. (1989). Analogical learning and transfer: What develops? In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 369–412). New York: Cambridge University Press.

Brown, A. L. (1990). Domain specific principles affect learning and transfer in children. *Cognitive Science, 14,* 107-134.

Brown, A. L., & DeLoache, J. S. (1978). Skills, plans, and self-regulation. In R. S. Siegler (Ed.), *Children's thinking: What develops?* (pp. 3-35). Hillsdale, NJ: Lawrence Erlbaum Associates.

Brown A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology, 20,* 493-523.

Bryant, P. E. (1974). *Perception and understanding in young children: An experimental approach.* New York: Basic Books.

Burstein, M. H. (1983). Concept formation by incremental analogical reasoning and debugging. *Proceedings of the International Machine Learning Workshop* (pp. 19-25). Urbana: University of Illinois.

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.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Carey, S. (1991). Knowledge acquisition: Enrichment or conceptual change? In S. Carey & R. Gelman (Eds.), *The epigenisis of mind: Essays on biology and cognition* (pp. 257–291). 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.

Chi, M. T. H. (1981). Knowledge development and memory performance. In M. Friedman, J. P. Das, & N. O'Conner (Eds.), *Intelligence and learning* (pp. 221–230). New York: Plenum.

Chipman, S. F. (1977). Complexity and structure in visual patterns. *Journal of Experimental Psychology: General, 106,* 269–301.

Chipman, S. F., & Mendelson, M. J. (1979). Influence of six types of visual structure on complexity judgments in children and adults. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 365–378.

Crisafi, M. A., & Brown, A. L. (1986). Analogical transfer in very young children: Combining two separately learned solutions to reach a goal. *Child Development, 57,* 953–968.

Falkenhainer, B. (1988). *Learning from physical analogies: A study of analogy and the explanation process* (Tech. Rep. No. UIUCDCS-R-88-1479). Urbana: University of Illinois, Department of Computer Science.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. In *Proceedings of the Meeting of the American Association for Artificial Intelligence* (pp. 272–277). Los Altos, CA: Morgan Kaufmann.

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

Forbus, K. D., & Gentner, D. (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? In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 341–348). Hillsdale, NJ: Lawrence Erlbaum Associates.

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 61–68). Hillsdale, NJ: Lawerence Erlbaum Associates.

Gelman, R. (1969). Conservation acquisition: A problem of learning to attend to relevant attributes. *Journal of Experimental Child Psychology, 7,* 167–187.

Gelman, R. (1982). Accessing one-to-one correspondence: Still another paper about conservation. *Journal of Psychology, 73,* 209–220.

Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development, 58,* 1532–1541.

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

Gentner, D. (1977b). If a tree had a knee, where would it be? Children's performance on simple spatial metaphors. *Papers and Reports on Child Language Development, 13,* 157–164.

Gentner, D. (1978). Testing the psychological reality of a representational model. *Proceedings of Theoretical Issues in Natural Language Processing* (Vol. 2, pp. 1–7). Urbana: University of Illinois, Association for Computing Machinery.

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. (1986). *Evidence for structure-mapping analogy and metaphor* (Tech. Rep. No. UIUCDCS-R-86-1316). Urbana: University of Illinois, Department of Computer Science.

Gentner, D. (1988) Analogical inference and analogical access. In A. Prieditis (Ed.), *Analogica* (pp. 63–88). Los Altos, CA: Morgan Kaufmann.

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

Gentner, D. (1990). *Metaphor as structure mapping: The relational shift.* (Tech. Rep. No. 488). Urbana: University of Illinois, Center for the Study of Reading.

Gentner, D. (1992). Commentary on Halford. *Human Development, 35,* 218–221.

Gentner, D., & Boronot, C. B. (1992). *Metaphors are (sometimes) processed as generative domain-mappings.* Unpublished manuscript.

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

Gentner, D., & Jeziorski, M. (1989). Historical shifts in the use of analogy in science. In B. Gholson, A. Houts, R. A. Neimeyer, & W. R. Shadish (Eds.), *The psychology of science and metascience* (pp. 296–325). New York: Cambridge University Press.

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. (1994). Similarity is like an analogy: Structural alignment in comparison. In C. Cacciari (Ed.), *Similarity.* Brussels: Brepols.

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). London: Cambridge University Press.

Gentner, D., Rattermann, M. J., & Campbell, R. (1994). *Evidence for a relational shift in the development of analogy.* Unpublished manuscript.

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

Gentner, D., & Wolff, P. (in press). Metaphor and knowledge change. In A. Kasher & Y. Shen (Eds.), *Cognitive aspects of metaphor, structure, comprehenson, and use.*

Gibbs, R. W., & O'Brien, J. E. (1990). Idioms and mental imagery: The metaphorical motivation for idiomatic meaning. *Cognition, 36,* 35–68.

Gibbs, R. W., & O'Brien, J. E. (in press). Idioms and mental imagery: The metaphorical motivation for idomatic meaning. *Cognition.*

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

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

Goswami, U. (1991). Analogical reasoning: What develops? A review of research and theory. *Child Development, 62,* 1–22.

Goswami, U. (1992). *Analogical reasoning in children.* Hove, UK: Lawrence Erlbaum Associates.

Goswami, U., & Brown, A. L. (1990). Higher-order structure and relational reasoning: Contrasting analogical and thematic relations. *Cognition, 36,* 207–226.

Halford, G. S. (1987). A structure-mapping approach to cognitive development. *International Journal of Psychology, 22,* 609–642.

Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development*

Halford, G. S. (1993). *Children's understanding: The development of mental models.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Halford, G. S., Maybery, M. T., & Bain, J. D. (1986). Capacity limitations in children's reasoning: A dual task approach. *Child Development, 57,* 616–627.

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

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

Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence.* New York: Basic Books.

Karmiloff-Smith, A. (1991). Beyond modularity: Innate constraints and developmental change. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and cognition* (pp. 171-197). Hillsdale, NJ: Lawrence Erlbaum Associates.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science.* Cambridge, MA: MIT Press.

Kass, A. (1989). Strategies for adapting explanations. In *Proceedings: Case-Based Reasoning Workshop* (pp. 119-123). San Mateo, CA: Morgan Kaufmann.

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: Morgan Kaufmann.

Kittay, E. (1987). *Metaphor: Its cognitive force and linguistic structure.* Oxford, England: Clarendon.

Kittay, E., & Lehrer, A. (1981). Semantic fields and the structure of metaphor. *Studies in Language, 5*(1), 31–63.

Klahr, D., & Wallace, J. G. (1976). *Cognitive development: An information processing view.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Kolstad, V., & Baillargeon, R. (1991). *Appearance and knowledge-based responses to containers in infants.* Unpublished manuscript.

Kotovsky, L. & Gentner, D. (1990). Pack light: You will go farther. In *Proceedings of the Second Midwest Artificial Intelligence and Cognitive Science Society Conference* (pp. 60–72). Hillsdale, NJ: Lawrence Erlbaum Associates.

Kotovsky, L., & Gentner, D. (1994). *Progressive alignment: A mechanism for the development of relational similarity.* Unpublished manuscript.

Lakoff, G., & Johnson, M. (1980). The metaphorical structure of the human conceptual system. *Cognitive Science, 4,* 195-208.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3,* 299-321.

Mandler, G. (1967). Organization and memory. In K. W. Spence & J. T. Spence (Eds.), *Psychology of learning and motivation* (Vol. 1, pp. 328-372). New York: Academic Press.

Mandler, J. M., & Mandler, G. (1964). *Thinking: From association to Gestalt.* New York: Wiley.

Markman, A. B., & Gentner, D. (1990). Analogical mapping during similarity judgments. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 38–44). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Markman, E. M. (1989). *Categorization in children: Problems of induction.* Cambridge, MA: MIT Press.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology, 16,* 1-27.

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

Meltzoff, A. N. (1990). Foundations for developing a concept of self: Role of imitation in relating self to other and the value of social mirroring, social modeling and self-practice in infancy. In D. Cicchetti & M. Beeghly (Eds.), *The self in transition: Infancy to childhood* (pp. 139-164). Chicago: University of Chicago Press.

Miller, G. A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81-93.

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

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., & Bobrow, D. G. (1979). Descriptions: An intermediate stage in memory retrieval. *Cognitive Psychology, 11,* 107-123.

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

Nunes, T., Schliemann, A. D., & Carraher, D. W. (1993). *Street mathematics and school mathematics.* New York: Cambridge University Press.

Ortony, A., Reynolds, R. E., & Arter, J. A. (1978). Metaphor: Theoretical and empirical research. *Psychological Bulletin, 85,* 919-943.

Piaget, J., Montangero, J., & Billeter, J. (1977). La formation des correlats [The formation of correlations]. In J. Piaget (Ed.), *L'Abstraction reflechissante* (pp. 115-129). Paris: Presses Universitaires de France.

Rattermann, M. J., & Gentner, D. (1990). The development of similarity use: It's what you know, not how you know it. In *Proceedings of the Second Midwest Artificial Intelligence and Cognitive Science Society Conference* (pp. 54-59). Hillsdale, NJ: Lawrence Erlbaum Associates.

Rattermann, M. J., Gentner, D., & DeLoache, J. (1990). Effects of labels on children's use of relational similarity. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 22-29). Hillsdale, NJ: Lawrence Erlbaum Associates.

Rattermann, M. J., Gentner, D., & DeLoache, J. (1994). *Effects of relational and object similarity on children's performance in a mapping task.* Unpublished manuscript.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology, 16,* 371-416.

Ross, B. H. (1989). Some psychological results on case-based reasoning. In *Proceedings: Case-Based Reasoning Workshop* (pp. 144-147). San Mateo, CA: Morgan Kaufmann.

Schumacher, R. M., & Gentner, D. (1988). Remembering causal systems: Effects of systematicity and surface similarity in delayed transfer. *Proceedings of the Human Factors Society 32nd Annual Meeting* (pp. 1271-1275). Santa Monica, CA: Human Factors Society.

Shepp, B. E. (1978). From perceived similarity to dimensional structure: A new hypothesis about perceptual development. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 135-167). Hillsdale, NJ: Lawrence Erlbaum Associates.

Siegler, R. S. (1984). Mechanisms of cognitive growth: Variation and selection. In R. J. Sternberg (Ed.), *Mechanisms of cognitive development* (pp. 141-162). Prospect Heights, IL: Waveland Press.

Siegler, R. S. (1989). Mechanisms of cognitive development. *Annual Review of Psychology, 40,* 353-379.

Smith, L. B. (1984). Young children's understanding of attributes and dimensions: A comparison of conceptual and linguistic measures. *Child Development, 55,* 363-380.

Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 146–178). New York: Cambridge University Press.

Smith, L. B., & Kemler, D. G. (1977). Developmental trends in free classification: Evidence for a new conceptualization of perceptual development. *Journal of Experimental Child Psychology, 24*, 279-298.

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

Smolensky, P. (1990). Neural and conceptual interpretation of PDP models. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing. Vol. 2: Psychological biological models* (pp. 390–431). Cambridge, MA: MIT Press.

Spinillo, A. G., & Bryant, P. (1991). Children's proportional judgments: The importance of "half." *Child Development, 62*, 427–440.

Sternberg, R. J. (1984). Mechanisms of cognitive development: A componential approach. In R. J. Sternberg (Ed.), *Mechanisms of cognitive development* (pp. 163-186). Prospect Heights, IL: Waveland Press.

Sternberg, R. J., & Downing, C. J. (1982). The development of higher-order reasoning in adolescence. *Child Development, 53*, 209-221.

Sternberg, R. J., & Nigro, G. (1980). Developmental patterns in the solution of verbal analogies. *Child Development, 51*, 27-38.

Sternberg, R. J., & Rifkin, B. (1979). The development of analogical reasoning processes. *Journal of Experimental Child Psychology, 27*, 195-232.

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

Turner, M. (1991). *Reading minds: The study of English in the age of cognitive science.* Princeton, NJ: Princeton University Press.

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

Uttal, D., & DeLoache, J. (1994). *Mapping and symbolization processes in young children.* Paper presented at the meeting of the Piaget Society, Chicago, IL.

Vosniadou, S. (1987). Children and metaphors. *Child Development, 58*, 870–885.

Vosniadou, S. (1989). Analogical reasoning as a mechanism in knowledge acquisition: A developmental perspective. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 413–437). New York: Cambridge University Press.

Waxman, S. (1991). Semantic and conceptual organization in preschoolers. In J. Byrnes & S. Gelman (Eds.), *Perspectives on language and thought* (pp. 107-145). Cambridge, England: Cambridge University Press.

Waxman, S. R., & Senghas, A. (1992). Relations among word meanings in early lexical development. *Developmental Psychology, 28*(5), 862–873.

Wolff, P., & Gentner, D. (1992). The time course of metaphor comprehension. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 504–509). Hillsdale, NJ: Lawrence Erlbaum Associates.