Chapter 10

The Development of Relational Category Knowledge

Dedre Gentner
Northwestern University

Fish gotta swim and birds gotta fly. Robins eat worms, dogs chase cats, roses love sunshine. These kinds of relations are as much a part of our understanding of the world as are the direct properties of the entities themselves. Relational knowledge is a prominent feature of human categories—indeed, of human cognition in general. Nowhere is this clearer than in our ability to learn relational categories like *gift, weapon, predator,* or *central force system.* By *relational categories,* I mean categories whose meanings consist either of (a) relations with other entities, as in *predator* or *gift,* or (b) internal relations among a set of components, as in *robbery* or *central force system.*

My purpose here is, first, to discuss relational categories and how they are learned, and second, to discuss how children learn relational information about object categories. Relational categories contrast with object categories[1] (e.g., *tiger* or *cow*), whose members share intrinsic features, often including perceptual commonalities. Of course, object categories typically contain not only property information but also relational information. For example, that tigers hunt and eat animals is part of our concept of a tiger, along with intrinsic attributes such as their stripes. I return to the role of relational information in ordinary object categories later. For now, I make a strong contrast between object categories and relational categories, to better reveal the dimensions of difference.

---

[1]Elsewhere I have referred to this dichotomy as the entity category versus relational category distinction (Gentner & Kurtz, in press).

As noted by Gentner and Kurtz (in press; Kurtz & Gentner, 2001; see also Asmuth & Gentner, 2004), relational categories abound in ordinary language: for example, *gift, father, accident, priority, benefit, setback, symmetry,* and so on. Some are fairly restricted: for example, in the relational category *father(x, y),* x and y must be animals of the same species (unless the term is metaphorically extended, an issue we consider later). But for many relational categories, the arguments can range widely. For example, a *goal* can be a physical hoop (for a basketball player), or a high mark on an exam, or a college degree, or a good marriage. Relational categories differ from object categories in that (as in the case of *goal*) the instances of a relational category do not need to have intrinsic properties in common. Gentner and Kurtz (in press), roughly following Markman and Stilwell (2001), divided relational categories into *relational role categories* (or *role categories*), and *relational schema categories* (or *schema categories*). Role categories, such as *thief,* are defined by *extrinsic relations*: Their members all play the same role in a relational schema. Schema categories, such as *robbery,* are defined by *internal relational structure.* Schema categories denote relational systems, and they generally take arguments. Role categories often serve as the arguments of implicit or explicit schema categories.

For example, *robbery* is a relational schema category with three conceptual arguments,[2] which are each relational role categories: *robbery* (*thief, goods-stolen, victim*). The three role categories are *thief* (the agent who steals), *goods* (the things stolen), and *victim* (the one stolen from). As this example illustrates, not all the relational roles have to be specified on any given occasion. For example, one can refer to a *bank robbery* without explicitly specifying the thief.

The relations that enter into relational categories may include common function (e.g., *both are edible*), mechanical causal relations (e.g., *both are strong so they can bend things*), biological causal relations (e.g., *both need water to grow*), role relations (e.g., *both grow on trees*), and progeneration (e.g., *both have babies*). Relational categories can also be based on perceptual relations such as *symmetric in form,* mathematical relations such as *prime,* or logical relations such as *deductively sound.* It is these relational systems that provide the theory-like aspects of concepts and categories.

Although relational categories occur frequently in adult language, their acquisition poses a challenge to simple accounts of learning. Indeed, the human facility at learning relational categories has led many theorists to conclude that they are largely built in, or that there is innate preparation in the form of nascent belief structures or skeletal theories. Frank Keil (1994) gave a particularly clear statement of the position that general learning

[2]Syntactically, *robbery*'s argument structure includes *thief* and *victim,* but not *goods stolen.* Thus we can refer to *the thief's robbery of the bank,* but not to **the jewels' robbery* nor to **the robbery of the jewels.*

processes are insufficient to explain children's acquisition of relational categories: ". . . the extraordinary ease with which all of us do learn about functional objects, such as tools, relative to other species that exhibit sophisticated learning in so many other areas also argues against reduction to general learning procedures" (p. 251). This conclusion would surely be correct if the only general learning processes were elementary association and perceptual generalization. But at this point there is sufficient evidence to warrant examining a more sophisticated learning mechanism—analogical learning. As I will describe, analogical mapping, or more broadly, structural alignment and mapping[3] is a learning mechanism capable of abstracting relational structure.

A central claim of this chapter is that relational categories can be learned via experiential learning, guided by language. To preview the argument, I suggest that the simple, ubiquitous act of comparing two things is often highly informative to human learners. I suggest that human comparison processes involve structure-mapping—a process of structural alignment and mapping that naturally highlights relational commonalities, and that promotes the learning of connected relational systems. I begin by reviewing relational categories and contrasting them with object categories. Next, I briefly summarize structure-mapping theory and consider its role in learning. Then I show evidence that comparison processes are central in children's learning of relational knowledge in categories. A recurring theme throughout the chapter is the crucial role of relational language in inviting and shaping relational categories.

RELATIONAL CATEGORIES

The study of nominal relational categories is largely uncharted. With a few exceptions (e.g., Barr & Caplan, 1987; Markman & Stilwell, 2001), relational categories have been largely ignored in studies of concepts and categories. This is surprising, given their frequency. Informal ratings of the 100 highest and 100 lowest frequency nouns in the British National Corpus by Asmuth and Gentner (in preparation) revealed that a third to a half are defined primarily by common relational structure rather than by common object properties. Relational categories such as *spouse, contradiction, deviation,* and *symmetry* (from the low-frequency list) are nearly as common in adult discourse as object categories like *suitcase, garlic,* and *pigeon* (from the same list).

[3]The term "analogy" is typically restricted to nonliteral comparisons. Our evidence suggests that the same mechanisms are used in literal comparison (e.g., Markman & Gentner, 1993; Medin, Goldstone, & Gentner, 1993).

Relational categories named by nouns (like the ones we have been considering) have some commonalities with those named by verbs and prepositions, in that their meanings include relations between other concepts. The semantic similarity between nominal relational categories and verb categories carries over into other commonalities as well, leading Gentner and Kurtz (in press) to advance a speculative analogy (see also Asmuth & Gentner, 2004):

Relational categories : Object categories : : Verb categories : Noun categories

This analogy is only partial. As Croft (2001) pointed out, nouns differ from verbs not only in their semantic properties, but in their pragmatic functions: The function of nouns is reference and that of verbs is prediction. However, one way to gain insight into the contrast between relational categories and object categories is to explore the contrast between nouns and verbs.

Contrasting Verbs and Nouns

Gentner (1981; Gentner & Boroditsky, 2001) described a large set of interrelated processing differences between verbs and nouns. For instance, verbs are more cross-linguistically variable (Talmy, 1975; Langacker, 1987); harder to translate, less likely to be borrowed in language contact (Sobin, 1982; Morimoto, 1999); acquired later[4] in both first and second languages (Caselli et al., 1995; Gentner & Boroditsky, 2001; Tardif, Gelman, & Xu, 1999); and distributed at higher word frequencies than nouns (Gentner, 1981). Verbs are less hierarchically structured and more likely to show multiple branching (Graesser & Hopkinson, 1987) or matrix structure (Huttenlocher & Lui, 1979). In addition, verbs are less likely to be accurately remembered or recalled than nouns (Kersten & Earles, in press; Earles & Kersten, 2000) and are more polysemous, more context-sensitive, and more semantically mutable than nouns (Gentner, 1981; Gentner & Boroditsky, 2002; Gentner & France, 1988).

If this analogy is correct, relational categories should behave (relative to object categories) as verbs do relative to nouns. Although only a few of these predictions have been tested so far, the results are encouraging. For example, consider the related phenomena of mutability and memory. Gentner (1981) suggested that verbs fare more poorly in recognition and recall than nouns in part because they are more likely to adapt their meanings to the current context than are nouns (Gentner & France, 1988). Building on this hypothesis, Kersten and Earles (in press) gave participants a list of simple intransitive noun-verb sentences and asked them to recognize either the nouns or the verbs from a later list of four kinds of intransitive sentences (old/new noun × old/new verb). Two findings are of interest: First, recognition for nouns was better overall than recognition of verbs; and second, verb recognition was substantially (and significantly) better when the verb was paired with the same noun at test as at encoding. This effect of context was much smaller for nouns.

Asmuth and Gentner (2004) tested whether this pattern would hold for relational nouns vis-à-vis object nouns; that is, whether relational nouns would adapt more to context and therefore be less well recognized than object nouns. Participants read noun-noun combinations consisting of a relational noun and an object noun (randomly combined from controlled-frequency lists)—for example, *a mountain limitation*—and rated their comprehensibility. Later, participants were given a recognition test with four kinds of pairs: old/new relational nouns and old/new object nouns. Participants were more likely to accept as old a combination with an old object noun and a new relational noun than the reverse. When participants saw an old object noun in a phrase, they were likely to judge the whole phrase as old (regardless of whether the relational noun was old or new); but when they saw an old relational noun in a phrase, their judgment of new or old depended to a large extent on the familiarity of the object noun. These results suggest that relational nouns are more context sensitive than object nouns, mirroring the behavior of verbs relative to nouns.

In addition to this evidence that relational nouns are more context sensitive and poorer in recognition than object nouns, there is also evidence that they are more difficult to generate. Kurtz and Gentner (2001) compared relational categories (e.g., *barrier* and *target*) to superordinate object categories (e.g., *furniture* and *vegetable*) using an exemplar generation task. Both the total number of subcategories generated and the rate of generation were higher for object categories than for relational categories. Further, not surprisingly, the items generated for the object categories were independently rated as much more similar to each other than those generated for the relational categories. (This fits with the greater overlap in intrinsic properties for object categories.) Interestingly, when the task was run in the other direction, the pattern reversed. Asked to generate the categories to which dog could belong, people generated many times more relational categories than taxonomic object categories.[5] This low-level "bushi-

[4] Although it has been claimed that verbs are acquired as rapidly or more rapidly than nouns in languages whose input favors verbs, such as Mandarin (Tardif, 1996) and Korean (Gopnik & Choi, 1995), this claim has so far proved false when the vocabularies are explored more completely (Au, Dapretto, & Song, 1994; Pae, 1993; Tardif, Gelman, & Xu, 1999; see Gentner & Boroditsky, 2001, for a review).

[5] In total, there were 6 object categories, and these could be hierarchically arranged: for example, *canine, mammal, animal, organism, living thing,* and *being.* There were 27 relational categories: *carnivore, pet, creature, guard, companion, friend, guide, hunter, racer, playmate, rescuer, fighter, showpiece, barrier, social parasite, threat, weapon, food, profit-maker, host for parasites, disease, carrier, cat chaser, swimmer, escapee, mess-maker,* and *transportation.*

ness" of relational categories—wherein a given item can have multiple upward branches—parallels the kinds of conceptual structures found for verbs (Graesser & Hopkinson, 1987; Pavlicic & Markman, 1997).

Acquiring Relational Categories

At some point in their first or second year, children catch on to the idea that words refer to things in the world. This first referential understanding centers around object and object categories. For these categories, the mapping from word → world is straightforward. The entities they refer to can readily be individuated on the basis of direct experience with the world (Gentner's [1982] Natural Partitions hypothesis). Often, the child has already individuated the referents prelinguistically, and has only to attach the word to the referent. I argued in my prior work that relational terms such as verbs and prepositions pose a greater challenge. Their referents are not simply "out there" in the experiential world; they are selected according to a semantic system. They have to be learned from the language as well as from the world. Indeed, as this line of theorizing predicts, names for entities—especially names for animate beings—are learned early in many languages (Gentner, 1982; Gentner & Boroditsky, 2001; Pae, 1993; Tardif et al., 1999). In this chapter I want to extend this argument to relational nouns. Just as concrete nouns predominate over relational-term verbs and prepositions in children's early vocabularies (because they refer to readily individuated things in the world), so object nouns should predominate over relational nouns. The same handicaps that make verbs hard to learn also make relational nouns hard to learn. Unlike object terms, they cannot be learned simply by correspondence to the world.

To see this, consider a highly simplified summary of the steps in initial word learning: (a) isolate part of the environment; (b) isolate part of the sound stream; (c) attach the sound segment to the environmental bit. For highly individuable entities and categories, step (a) is already done. To establish the word-to-world mapping, it only remains to find and attach the right speech segment. In contrast, for relational nouns, the parsing of the perceptual world into individual referents is not obvious. A child can tell which things in the room are *apples* by looking at them; but she cannot tell which things are *gifts*, or *goods-for-sale*, or *weapons*, without knowing something about their relations to other entities. The mapping from world-to-word is equally confusing: A carving knife is a *gift* if it is in a box under the Christmas tree, but a *kitchen implement* if it's used to carve a turkey, and a *weapon* if Mom uses it to scare away a burglar.

All this suggests that relational categories should be acquired later than object categories. One indication that this is the case comes from the MacArthur Communicative Developmental Inventory, which can be taken as a rough upper bound of the words children are likely to have in their vocabularies at a given age. For 8- to 16-month-olds, it lists 296 nouns, of which 93% are entity nouns (objects, animals, and people) and 7% are mixed entity-relational nouns, with no purely relational nouns. For 17- to 30-month-olds, the MCDI lists 411 nouns: 79% entity nouns, 13% mixed, and 8% relational nouns. Second, when children *do* learn relational terms, they often initially treat them as entity terms (Gentner & Rattermann, 1991). For example, a young child may describe a *brother* as a boy about 12 years old, and only later come to realize that it can be any male (however young or old) who is someone's sibling (E. Clark, 1993). For example, Hall and Waxman (1993) found that 3½-year-olds had difficulty learning novel relational nouns denoting concepts like *passenger*. Even when they were explicitly told (for example) "This one is a *blicket* BECAUSE IT IS RIDING IN A CAR," they tended to interpret the novel noun as referring to the object category. Likewise, Keil and Batterman (1984) found that 4-year-olds define a *taxi* as a yellow car. The relational knowledge that a taxi is a car that someone can hire to go where they want does not appear until a few years later. Eve Clark (1993) noted a similar pattern with kinship terms: Young children think of a *brother* as a young boy, and are surprised to see a middle-aged man described as someone's brother.

Having established the comparative difficulty of relational categories, we now turn to how they are acquired. I suggest that a major mechanism of relational learning is comparison—or more precisely, structure-mapping is a learning mechanism that can transmute experiential knowledge into abstract relational structures. I further suggest that abstract conceptual knowledge can be acquired by conservative initial learning followed by comparison processes that highlight common relational structure. Finally, I show that comparison can be invited by common language labels—*symbolic juxtaposition*—as well as by actual experiential juxtaposition. This gives humans immense flexibility in which relational structures are extracted and transmitted. Because a central issue in this work is mechanisms of learning, in the next section I review structure-mapping in comparison and discuss its role in children's learning.

The Role of Comparison in Learning Relational Categories. Comparison is a general learning process that can promote deep relational learning and the development of theory-level explanations (Forbus, 2001; Gentner, 1983, 2003). According to structure-mapping theory, comparison acts to highlight commonalities, particularly relational commonalities that may not have been noticed prior to comparison (Gentner, 1983, 2003). This is because the structural alignment process operates to promote common systems of interconnected relations (as described further later). Thus when two representations are aligned, common structure is preferentially highlighted. This can result in the extraction of common higher order relational structure that was not readily evident within either item alone.

It is this property of comparison processing—that it heightens the salience of common structure—that is the key to how relational knowledge can develop via experiential learning. For example, there is considerable evidence that children and adults who compare two analogous cases are likely to derive a relational schema and to succeed subsequently at relational transfer to a greater degree than those who have seen the same materials but not compared them (Gentner, Loewenstein, & Thompson, 2003; Gick & Holyoak, 1983; Loewenstein, Thompson, & Gentner, 1999). This effect of alignment in promoting relational abstraction is evident across a wide span of ages and cognitive tasks (Gentner & Namy, 1999; Kotovsky & Gentner, 1996; Loewenstein & Gentner, 2001).

Structure-Mapping as a Domain-General Learning Mechanism

The defining characteristic of analogy is that it involves an alignment of relational structure. But it is not necessary to begin with a clear relational description in order for this process to operate. The simulation that embodies the process model of structure-mapping, SME (the Structure-Mapping Engine), begins by forming all possible identity matches, and gradually coalesces these to reach a global alignment (Falkenhainer, Forbus, & Gentner, 1989; Forbus et al., 1995). In the initial, local-match stage, SME typically has a large number of mutually inconsistent matches. Then it uses connections between the matched elements to impose structural consistency and to propose further matches, eventually arriving at one or a few global alignments that constitute possible interpretations of the comparison (for details, see Forbus, Gentner, & Law, 1995; Gentner & Markman, 1997; Markman & Gentner, 2000). The alignment must be structurally consistent; that is, it must be based on a *one-to-one correspondence* between elements of the representations, and it must satisfy *parallel connectivity* (i.e., if two predicates correspond, then their arguments must also correspond).

The final characteristic of analogy is *systematicity*: what tends to win out in an analogical match is a *connected system of relations* rather than an isolated set of matching bits (Gentner, 1983; Gentner & Markman, 1997). It is as though we had an implicit aesthetic built into our comparison process that likes connected systems better than lists of separate matches. Our penchant for systematicity seems to betoken a tacit preference for coherence and causal predictive power in analogical processing. Systematicity operates in several ways in the structure-mapping process. First, during the mapping itself, connected matches get more activation than do isolated matches (because match evidence is passed from a predicate to its arguments—one reason that interconnected matches win out). Second, systematicity influences whether a rerepresentation processes occur, as described later. Third, it influences which inferences are drawn: new inferences are only proposed if (a) there is an aligned common system; and (b) there is a predicate that is connected to the system in the base and not yet present in the target. In that case the predicate is projected into the target structure as a candidate inference. Finally, systematicity also influences the selection of an interpretation when more than one structurally consistent match exists between two situations. Assuming contextual relevance is equivalent, the largest and deepest, that is, the most systematic, relational match is preferred.

An important feature of the structure-mapping process—particularly if one intends to model learning in children—is that achieving a deep structural alignment does not require that the common schema be known in advance. Common structure can emerge through a comparison process that begins blind and local. This makes it an interesting candidate for a developmental learning process. For example, a child who notices a chance similarity may end up noticing common causal patterns without explicitly seeking to do so.

This process model suggests that comparison can promote learning in several different ways, including (a) highlighting common relational systems, thereby promoting the disembedding of subtle and possibly important structure; (b) projection of candidate inferences—inviting the importing of new knowledge about one domain on the basis of the other; and (c) rerepresentation—altering one or both representations so as to improve the match, as amplified later. It is clear that candidate inferences can qualify as learning. But at first glance it may seem that highlighting common structure is out of place on this list. I believe that highlighting does count as a learning process, based on the following assumptions. First, I suggest that human representations tend to be rich and contextually situated. Second, this is especially true of early domain learning, which tends to be conservative and specific to the initial context. In this case, focusing on a sparse common relational structure can be highly informative to the learner.

A third way that comparison fosters learning, in addition to highlighting common structure and inviting inferences, is by rerepresentation. Once a partial alignment has been built up, if two potentially corresponding relations do not match, and if their match would substantially improve the analogy, they can sometimes be rerepresented to match better (Falkenhainer, 1990; Yan, Forbus, & Gentner, 2003). For example, a group of participants asked to compare "George divorced Martha" and "Megacorp divested itself of Regal Tires" wrote out such commonalities as "Both got rid of something" or "Both ended an association." As in this example, rerepresentation can result in the extraction of common higher order relational structure that was formerly embedded in a specific domain. In Clement, Mawby, and Giles' (1994) terms, latent relational commonalities become manifest via rerepresentation.

Structure-Mapping in Development

This process model has implications for the course of comparison in learning and development. First, because matches at all levels enter into the maximal alignment, the easiest and most inevitably noticed similarity comparisons are those in which there is rich overall (literal) similarity between the situations. Thus, the match between a dachshund chasing a mouse and another dachshund chasing a mouse is inescapable; even a small child can see such a match. In this case the comparison process runs off easily, because the matching information is obvious and the object matches support the relational match. This means that a strong overall similarity match produces one clear dominant interpretation. In contrast, the match between a dachshund chasing a mouse and a shark chasing a salmon is not obvious on the surface. It is easy enough for a learner who understands ideas of predation and carnivorous behavior, but these ideas may not be available to a novice such as a small child. This suggests that novice learners and children should perceive overall similarity matches before they perceive partial matches, and indeed, this is the case (Gentner & Loewenstein, 2002; Halford, 1987, 1993; Kotovsky & Gentner, 1996; Smith, 1983, 1989).

Early in learning, children (and other novices) have rich knowledge of objects and sparse knowledge of relations; thus they initially make object matches and overall similarity matches. With increasing domain knowledge, children's relational representations become richer and deeper; it becomes possible to perceive and interpret purely relational matches. Thus, there occurs a relational shift with increasing knowledge (Gentner, 1988; Gentner & Rattermann, 1991). This shift is not linked to any particular Piagetian stage. Rather, the timing of the shift varies across domains and appears driven by changes in domain knowledge (Rattermann & Gentner, 1998a; see also Siegler, 1989). Evidence that the relational shift is due to gains in relational knowledge comes in three varieties: (a) the relational shift occurs at different ages for different domains and tasks; in particular, even very young children can show considerable analogical ability in highly familiar domains (Goswami & Brown, 1989; Gentner & Rattermann, 1991; Rattermann & Gentner, 1998a); (b) within a given age, children (and adults) who possess deep relational knowledge perform better in tasks than those who do not; and (c) children's analogical performance can be aided by the introduction of relational language.

It is much easier for young children to carry out a literal similarity comparison than to carry out a purely relational comparison, as just discussed. But our research suggests that overall literal similarity, though it may seem obvious to adults, can be informative to children. We have repeatedly found that very close similarity comparisons potentiate relational insight and far transfer among children (Gentner & Namy, 1999, 2004; Kotovsky & Gentner, 1996; Loewenstein & Gentner, 2001). In particular, in some of our work, young children have shown substantial gains in relational insight after they were led to make close comparisons. Before describing the research, a brief introduction is in order.

Within this theoretical framework, we can contrast two kinds of analogical learning: *projective analogy* and *analogical encoding*. Both involve the same basic processes of alignment and projection. In *projective analogy*, the learning results chiefly from inference projection. A well-understood situation (the base) is aligned with a less understood situation (the target), and inferences are mapped from the base to the target (Gentner, 1983; Gentner & Markman, 1997; Holyoak & Thagard, 1995). Projective analogy plays an important role in learning and instruction (Bassok, 1990; Ross, 1987) as well as in scientific discovery (Dunbar, 1993; Gentner, 2002; Nersessian, 1992). But although projective analogy is important in learning, it cannot explain how the process gets started. If children learn by analogy with prior situations, then what happens in early learning, before the child has amassed a store of well-understood situations? A second limiting factor for projective analogy, even among adults, is the frequent failure of analogical remindings. When dealing with a given problem people often fail to think of a prior analogous case, even when it can be demonstrated that they have retained a memory of the case (Gentner, Rattermann, & Forbus, 1993; Gick & Holyoak, 1980, 1983). Thus even for adults, something more is needed to explain spontaneous relational learning.

This brings us to a second kind of analogical learning, *analogical encoding* or *mutual alignment* (Gentner, Loewenstein, & Thompson, 2003; Gentner & Namy, 1999; Loewenstein & Gentner, 2001; Loewenstein, Thompson, & Gentner, 1999; Namy & Gentner, 2002). Analogical encoding occurs when two analogous situations are present simultaneously and are compared to one another. Here the key process is not directional projection of information (though inferences can occur), but aligning, rerepresenting, and abstracting commonalities. If inferences are drawn they may be bidirectional, with both examples serving as bases as well as targets.

At first glance, it may seem that nothing new could come out of such a process. Indeed, it might seem that the result will simply be confusion, especially in early learning, when the two cases are only partially understood. On the contrary, the answer that emerges from our research is that comparison between two partially understood situations can lead to a better grasp of the relational structure and a deeper understanding of both situations. We have found this pattern with both children (Gentner & Namy, 1999;

Loewenstein & Gentner, 2001; Namy & Gentner, 2002; Waxman & Klibanoff, 2000) and adults (Loewenstein et al., 1999; Kurtz, Miao, & Gentner, 2001), in a variety of domains. Our results suggest that analogical encoding is a pervasive and important bootstrapping process. This process is initially quite conservative, particularly when driven by the child's own experience. Early in learning children spontaneously notice only very close matches, and spontaneously make only very small adjustments to the representations to achieve alignment. But a small gain in relational salience may pave the way for larger insights. In our research, we have found that even close alignments potentiate more distant (more analogical) alignments.

For example, in one line of studies, Kotovsky and Gentner (1996) showed that experience with concrete similarity comparisons can improve children's ability to detect cross-dimensional similarity. In these studies, 4-year-olds' ability to perceive cross-dimensional matches (e.g., a match between a size symmetry figure and a brightness symmetry figure) was markedly better when the cross-dimensional trials were presented *after* blocks of within-dimension trials (blocks of size symmetry and blocks of brightness symmetry) than when the two kinds of trials were intermixed. Note that this manipulation was extremely subtle. Both groups of children received the same trials, half within-dimension and half cross-dimension. Their task throughout was simply to choose which of two alternatives was most similar to the standard. No feedback was given at any time. Our intent was to mimic the results of experiential learning from fortuitous runs of examples that are either easy or difficult to compare.

The materials were all perceptual configurations, each made up of three geometric figures. A schematic example of a within-dimension size symmetry trial would be v-V-v (standard) with alternatives o-O-o versus O-o-o. In the cross-dimension trials, the alternatives were depicted along a different dimension as the standard (e.g., size vs. brightness): for example, v-V-v (standard) with ooo versus ooo (alternatives). In all cases, the two alternatives were made up of the same geometric figures (which were different from those in the standard), but only one alternative matched the standard's relational configuration. Thus the best answer could only be determined by a relational match.

The results showed that experience with concrete similarity comparisons can facilitate children's subsequent ability to detect cross-dimensional similarity. Specifically, 4-year-olds' ability to perceive cross-dimensional matches (e.g., size symmetry/brightness symmetry) was markedly better after experience with blocked trials of concrete similarity (blocks of size symmetry and blocks of brightness symmetry), as compared to a group who received the same set of trials intermixed. This result is perhaps surprising; it might have been supposed that comparing two highly similar examples would lead to the formation of a narrow understanding. After all, the rule of thumb in animal learning is "narrow training leads to fast learning, but narrow transfer." But in human learning, aligning highly similar examples can facilitate later perception of purely relational commonalities.

Kotovsky and Gentner (1996) suggested that the superior performance of the close-to-far sequence group results from *progressive alignment*: The within-dimension comparisons, being strong overall matches, are easy to align. But each time such an alignment occurs, the common structure is highlighted. Thus repeated experience on the within-dimension pairs acts to make the higher order relation of symmetry (or monotonic increase) more salient. When these children then encounter a cross-dimensional match, they experience a near-match of two highly salient relational structures, and this prompts a rerepresentation process, as discussed earlier. The remarks of one articulate 8-year-old in the mixed condition epitomize the process of alignment and rerepresentation.[6] On her first six trials, she responded correctly to all three within-dimension trials and incorrectly to all three cross-dimension trials. On the latter, she expressed a fair degree of frustration, with comments like "It can't be the size, because those two are the same size. It can't be color." Finally, on her seventh trial, she exclaimed excitedly, "Even though the smaller ones come first and the big one's in the middle, it's exactly the same—but different!" She went on to choose correctly for the remainder of the study. This comment "It's exactly the same—but different!" captures the essence of a rerepresentational insight.

These findings show that close alignment can potentiate far alignment. Making even very easy comparisons can increase children's insight into relational similarities. This pattern is consistent with research by Chen and Klahr (1999), who taught children the strategy of *control of variables*—normally a difficult idea to communicate—by giving them intensive practice in three different domains. Within each domain (e.g., pendulum motion), children naturally engaged in many close comparisons such as varying string length, weight, and so on. When children were transferred to a second, and especially to a third domain, they learned the strategy far faster than they had in the first domain. I speculate that one reason that children learned so well in this study was the concerted experience they received within each domain before encountering the next domain.

The claim that comparison enables spontaneous learning receives support from observations of young children's spontaneous learning. Many of the insights gained in this way are rather small, and to an adult may seem prosaic; but they are new to the child.

- Emma (at about 19 months) had a little cut on one hand. She looked back and forth between her two hands and then, looking at the hand with

[6]Unlike the 4-year-olds, 6- and 8-year-olds often succeeded at cross-dimensional matching even in the mixed condition.

the cut she said, "Yes booboo." Then she looked at the other hand and said "No booboo." She proceeded to look back and forth and say "Yes booboo; no booboo" at least half a dozen times, proud of herself for figuring this one out (J. Loewenstein, personal communication, January, 2003).

- Ricky (at about 22 months) was walking through the zoo with his mother and grandmother: "In the carnivore house we watched a magnificent Siberian tiger cleaning its paws with a huge pink tongue. Ricky stuck out his own tongue and handled it. Then he wanted to see my tongue and his mother's tongue in turn, as if to compare their sizes to his own and the tiger's. The tiger began to swing his tail slowly back and forth. Pointing to it, Ricky's mother said, 'Look at the tiger's long tail.' Ricky watched attentively, and then turned to look over his shoulder and felt his bottom to see the status of his tailedness" (M. Shatz, personal communication, April 1989; see also Shatz, 1994, p. 54).
- Sophie (3 years, 2 months) told that her father had gone to Jimmy's, at first thought it was a gym (as in "Gym-ies"). When her mother explained that Jimmy's was the name of a bar owned by a man named Jimmy, she said "Oh, I know. It's like Leona's (an Italian restaurant)." She later became interested in the fact that another restaurant, called Barbara Fritchie's, is *not* owned by Barbara Fritchie (A. Woodward, personal communication, March, 2003).
- At the same age, Sophie, having been told earlier that day about how baby teeth fall out and adult teeth grow in, advanced the theory that as she grows, her old "small" bones will leave her body and new "bigger" bones will grow in to replace them. She was very taken with the theory and resistant to correction. She also asserted that her small bones would be used by babies after she got her big bones—possibly by analogy to what happens with clothing she outgrows (A. Woodward, personal communication, March, 2003).

These examples illustrate several points: First, many of the early insights gained from spontaneous comparison are rather tiny from the adult point of view. Second, as with adults, candidate inferences may turn out to be wrong. Children do not have tails like tigers, and baby bones do not get discarded like baby teeth. But as in adult analogy, rejecting an analogical inference can be an occasion for learning. The child who checks to see whether he has a tail is following a precise morphological analogy; both the insight as to *where* his tail would have to be, if he had one, and the explicit noticing of the fact that he does not are advances in his understanding. In the restaurant case, by pursuing the analogy, Sophie learns the correct rule, that many, but not all, restaurants with human names are owned by the corresponding person. A third point, brought out by the tiger analogy, is that children do not have to discover all these useful comparisons

by themselves. Adults often suggest a comparison, and the child then takes it further.

This raises the question of what kinds of situations invite children to engage in comparison process, especially early in learning, when (as just discussed) children's relational encodings are situated in specific contexts. One way that the "invitation to compare" can come about is through high surface similarity and spatiotemporal juxtaposition, as in the Kotovsky and Gentner studies. These kinds of close comparisons abound in some domains—for example, early object manipulation and support. In these domains, children naturally encounter closely juxtaposed high similarity pairs often enough to materially advance their understanding. For example, when young children repeat the same actions over and over (e.g., building and knocking down a block tower), this may reflect their delight at comparing and learning from small variations in the action.

Some relational categories may be helped along by this kind of experiential learning, with progressive alignment from close matches to less similar instances of the same relational structure. For example, the category *gift* may be initially learned by comparing literally similar exemplars at birthdays and holidays. The child's first representation of *gift* may typically include wrapping paper and a bow. Over time, further comparisons could lead the child to notice and extract the relational commonality that a gift must pass from one person to another. But many relational categories lack such "training wheels"—closely similar, frequently juxtaposed pairs that can seed the alignment of relational structure. Here culture and language step in. One way children notice nonobvious relational similarity is by receiving a direct signal to compare, as in "Look, these are alike" or "See this? It's kind of like a robin" (Callanan, 1990). Another way that children are invited to compare things that are not obviously similar is by their having the same linguistic label—what Gentner and Medina (1998) called "symbolic juxtaposition."

We have found evidence in our prior research that linguistic labels can invite relational concepts (Gentner & Loewenstein, 2002; Gentner & Rattermann, 1991; Loewenstein & Gentner, in press; Rattermann & Gentner, 1998b). For example, Rattermann and Gentner (1998b, 2001) found that introducing relational language helped 3-year-olds to carry out a relational mapping. In these studies, the relational pattern was *monotonic increase in size* across a line of objects; the correct answer was based on matching relative size and position. The mapping was made difficult by including a cross-mapping (Gentner & Toupin, 1986), such that the correct relational match was in conflict with a high-similarity local object match. We found that children who heard language conveying a monotonic relational structure (either *Daddy-Mommy-Baby* or *big-little-tiny*) performed far better than those who did not. The effect of language was dramatic: 3-year-olds in the

No-label condition performed at chance (32% correct); in contrast, a matched group given relational labels performed at 79% correct. Their performance was comparable to that of the 5-year-olds in the No-label condition. Further, children were able to transfer this learning to new triads even with no further use of the labels by the experimenters. Finally, children who receive relational language maintained high performance even 4 to 6 weeks later. Rattermann and Gentner suggested that the use of relational labels invited attention to the common relation of monotonic change and made it possible for the children to carry out a relational alignment (see Loewenstein & Gentner, in press, for related findings). These results are evidence for a facilitating effect of common language on children's appreciation of relational similarities.

The Role of Relational Language in Learning Relational Categories

The preceding discussion suggests that analogical comparison might provide a route by which children can learn relational categories—categories like *gift* and *accident*. For the many relational categories whose members do not share intrinsic similarity, I speculate that relational language plays a crucial role in acquisition, by inviting comparisons among exemplars of the categories (Gentner & Loewenstein, 2002). In this case, symbolic juxtaposition via common labels (rather than by spatiotemporal juxtaposition and overall similarity) is the impetus for comparison.

Gentner and Klibanoff (2001) investigated the acquisition of such categories by 3-, 4-, and 6-year-olds. As noted previously, relational categories (whose membership is determined by common relations rather than with common object properties) are learned rather late. This pattern fits with the general pattern of a relational shift in children's understanding, as discussed earlier (Gentner & Rattermann, 1991). More specifically, in the word learning task, children's well-documented focus on objects and object categories when learning nouns (E. Markman, 1989; Waxman, 1990) might be expected to interfere with their ability to learn relational nouns, which disregard the intrinsic properties of the object named. In a sense, children have to overcome their prior object-naming strategies to learn relational nouns.

Gentner and Klibanoff (2001) asked (a) whether children could derive a new relational abstraction over two examples; and (b) whether receiving a novel word to describe the common relation would help or hurt. We expected that older children—6-year-olds—would benefit from hearing a word used, while younger children—3-year-olds—would seek object reference meanings and would thus be less willing to derive a relational meaning if a word were used than without words.

We used a combination of comparison and labeling. Children aged 3, 4, and 6 were shown two pairs of picture cards (e.g., a knife and a watermelon, followed by an ax and a tree). In the Word condition, the experimenter used the same novel relational noun in both of these parallel contexts: for example, "Look, the knife is the *blick* for the watermelon. And see, the ax is the *blick* for the tree." Then the children were asked to choose the referent of the new relational term in a third context: for example, "What would be the *blick* for the paper?" They were given three alternative pictures: a pair of scissors (*relational response* [correct]), a pencil (*thematic response*), and another piece of paper (*object response*). A No-word control group saw the same series of analogous examples without the novel word: for example, "The knife goes with the watermelon, and the ax goes with the tree the *same way*. What would go with the paper the same way?" Note that the No-word group received overt comparison information (that the pairs go together in "the same way"), whereas the language group did not. If common labels are especially effective at encouraging deeper comparison processes, then children in the Word condition should outperform those in the No-word condition. However, if children take the word to be the name of an object, then children in the Word condition should perform worse than those in the No-word condition.

Not surprisingly, there was an effect of age: 4- and 6-year-old children gave many more relational responses than 3-year-olds, who performed at chance. More interestingly, as predicted, 6-year-olds, and also 4-year-olds, who heard novel relational nouns were more likely to choose the *same-relation* card than were their counterparts in the No-Word condition. In these studies, we did not see a depressive effect of a novel word; the youngest group (3-year-olds) performed at chance in both conditions. These studies suggest that at least by the age of four, (a) the meanings of novel relational categories can be learned through comparison by abstracting common relations across situations; (b) the use of a common label invites comparison processes; and (c) a direct statement that the situations are alike also prompts comparison, though (at least here) not as effectively as using the same word. These results show that common language can actually be a stronger invitation to compare than even direct statements like "this one goes with this one *in the same way*."

Learning Relational Aspects of Object Categories: Effects of Language and of Comparison

We have discussed children's learning of relational categories. But what about the relational aspects of object categories? Relational structure is central in adult category representation, and our relational knowledge about ordinary categories is often quite rich. For example, many studies have demonstrated the importance of causal structure in category learning and use. It has been shown that people's ability to learn and use categories is influenced by the number of causal links running in either direction (Rehder & Hastie, 2001) as well as by the causal status of properties as causes or effects (Ahn, 1998; Ahn, Gelman, Amsterlaw, Hohenstein, & Kalish, 2000). This work suggests that properties that enter into causal relations—perhaps especially those that cause other properties—take on greater salience and greater weight in membership judgments (Sloman, Love, & Ahn, 1998). Another indication of the importance of relational structure is that animate and artifactual categories appear to be differentiated in part by the nature of their causal relations (Carey, 1992; R. Gelman, Spelke, & Meck, 1983; S. Gelman, 1988; Keil, 1994). The importance of causal explanatory patterns in category representations is also underscored by the examples and arguments given in Murphy and Medin's (1985) theory view of categorization. For example, we classify a man in a tuxedo leaping into a swimming pool as a drunk person, even if that particular exemplar is entirely new, because our causal model of drunken partygoers fits his behavior.

How do children come to understand the relational aspects of ordinary entity categories—the causal and functional aspects of categories like *plate, tricycle,* and *umbrella*? Our research suggests that although preschool children may have considerable tacit knowledge of the relations that things participate in, this knowledge is often initially situationally embedded and difficult to access explicitly. To put it concretely, a child who knows that she can ride on a tricycle does not necessarily know that a tricycle can be categorized with other vehicles. Our research further suggests that comparison across examples is crucial in achieving a disembedded, portable knowledge of the relations that characterize ordinary categories.

Early in learning, children rely heavily on perceptual similarity in category extension (Baldwin, 1989; Imai, Gentner, & Uchida, 1994; Landau, Smith, & Jones, 1988). For example, they apply a novel word to objects that share shape or other distinctive features with the exemplar on which the label was learned; they call horses and cats "doggies," or any round shape a "ball" (Clark, 1973). This bias toward perceptual similarity as a basis for word extension might be a reasonable heuristic for young children, given their incomplete knowledge of causal and functional properties and how they enter into word meaning. The high correlation between perceptual similarity and conceptual similarity for basic level categories (which predominate in preschoolers' lexicons) means that perceptual similarity is often a good guide to a word's extension. However, this strategy is clearly not adequate over the long haul. Children must eventually come to appreciate the relational commonalities that loom large in the intensions of human categories. The question, then, is how children come to appreciate functional and relational aspects of categories.

Laura Namy and I proposed that structural alignment aids children's category learning by elevating the salience of relational knowledge that might otherwise remain situated and implicit (Gentner & Namy, 1999, 2004; Namy & Gentner, 2002). Specifically, we suggested (a) that hearing common labels applied to multiple entities invites children to engage in comparison processes; and (b) that the process of comparison highlights relational commonalities, including many that are not immediately evident on surface-level inspection. If comparison renders relations more salient, and if common labels encourage comparison, then this process might enable children to override compelling perceptual commonalities in favor of deeper conceptual ones.

To test these claims, we experimentally manipulated children's opportunity to compare objects from a given category and then tested their word extension, pitting a perceptual match against a perceptually dissimilar taxonomic match (Gentner & Namy, 1999). We used the standard word-extension method of applying a novel word (in "doggie language") to a standard, and asking the children to choose another exemplar from two (or three) alternatives. In our studies, we manipulated the opportunity to compare by presenting children with either a single instance or two perceptually similar instances of the category before eliciting their word extensions. The design of the materials allowed us to test whether comparison would highlight *relational* information or featural information. There were three groups of 4-year-olds: a comparison group and two solo (single-standard) groups. For example, children in the comparison group might see two instances of nonmotoric vehicles (a bicycle and a tricycle). One solo group would see the bicycle, and the other would see the tricycle. The task was a standard word-extension task. Children were told a novel puppet name for the standard(s) such as "blicket," and then asked to choose another blicket. The children chose from two alternatives: a perceptual alternative (a perceptually similar object from a different category, e.g., eyeglasses) and a taxonomic category alternative (a perceptually dissimilar object from the same category, e.g., a skateboard) (see Fig. 10.1).

Importantly, the standards were designed so that each of them was highly similar to the perceptual alternative. Indeed, when either standard was presented singly (in the solo conditions), the children chose the perceptually similar alternative. This allows a test of the structure-mapping account of the comparison process. If children compute similarity merely by concatenating surface commonalities, then if they select the perceptual match in the two solo conditions, they should be doubly likely to do so in the comparison condition. That is, on the behaviorist account of similarity, comparing two standards that are *both* more featurally similar to the perceptual choice than to the category choice should *increase* perceptual responding, relative to viewing a single standard. (Note also that the perceptual fea-

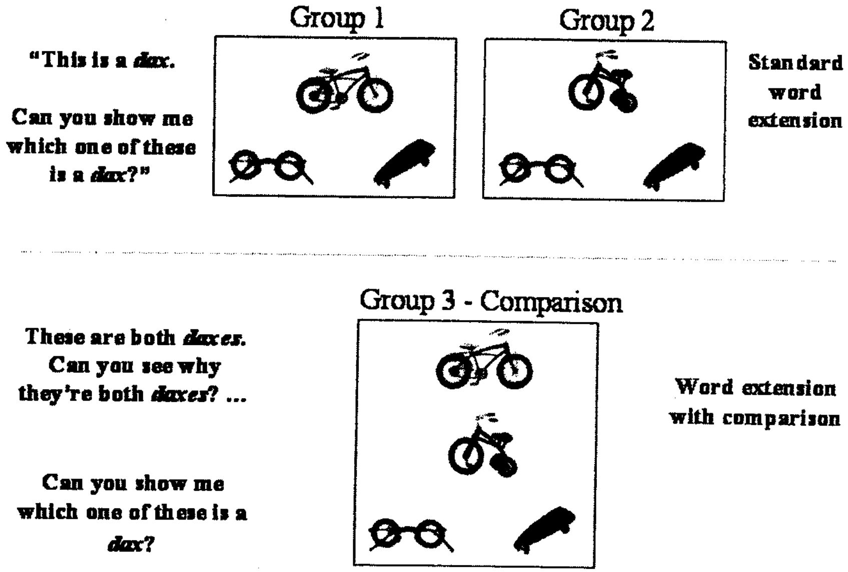

FIG. 10.1. Sample materials in the Gentner and Namy (1999) study, showing solo conditions (above) and the comparison condition (below).

ures that the two standards both share are the *same* features they each share with the perceptual alternative [e.g., two horizontally aligned circles]. Thus it is not a question of different matches diluting the featural predictions.) However, if comparing instances induces a structural alignment process, then children who view two standards may be led to focus on previously implicit common relational structure, such as how the objects are used and what causal activities they normally participate in. If so, then comparison should lead to a shift toward taxonomic category responding, despite the strong perceptual similarity between both of the standards and the perceptual alternative.

This is precisely what we found. When shown both standards *together*, children chose the conceptual match as the other "blicket" despite the fact that their agemates preferred the perceptual match for *either* of the standards presented singly. These results show that comparison can facilitate word extension on the basis of conceptual relations, and not merely perceptual features. More generally, at the theoretical level, these results provide critical evidence that comparison highlights common relations, *even when common salient object features are also available.* Even when both standards share salient perceptual commonalities, aligning the two can reveal a common relational structure that the child will then attend to. This suggests a route by which early perceptually driven word extensions can give rise to

conceptual understanding. Structure-mapping processes may thus be instrumental in guiding children to adult-like category knowledge.

The Role of Language in Inviting Comparison. How important is language in this process? Gentner and Namy (2003) found that hearing a common label encouraged children to engage in comparison. We compared 4-year-olds' performance on the word extension task when they were either given a novel label for the two standards, as already shown, or were simply asked to compare them and find another: "See this one, and see this one? See how these are the same kind of thing? Can you find another one that's the same kind as these?" Children gave more category responses when common labels were used than in the no-label condition. This suggests that having a common label helps to invite comparison processes, that (to paraphrase Roger Brown, 1958), words are invitations to make comparisons.

In subsequent studies, Namy and Gentner (2002) further probed the relation between common labels and alignment processes. Children (again, 4-year-olds) were assigned to either a Unifying Label or Conflicting Label condition. In both conditions, children were shown two standard objects from the same category. Those in the Unifying Label condition heard both standards labeled with the same novel word (e.g., "This is a blicket and this is a blicket!"). Those in the Conflicting Label condition heard the two standards labeled with different novel words (e.g., "This is a blicket and this is a daxen!"). Then both groups were shown two alternatives (a perceptual alternative and a category alternative) and asked "Can you tell me? Which one is the same kind as these?" As in our previous studies, children who heard a Unifying label reliably selected the category alternative. In contrast, children who heard Conflicting labels chose the perceptual alternative; they resembled the children who had seen only a single standard object in the previous studies.

These results, like the Gentner and Klibanoff results, show that alignment can be invited by hearing a common label for two exemplars. Our results also show that alignment processes are used to extend novel words to new instances. Thus, we suggest that the relation between alignment and word-learning is a true boot-strapping relation. Hearing a common term invites an alignment that is then used to extend the term to new exemplars.

In addition to showing effects of language, these findings show that initial attention to perceptual similarities in word extension can actually serve to promote attention to deeper commonalities. Even though children's attention may be initially drawn to surface commonalities between two exemplars, the full process of comparison will result in highlighting any further relational commonalities that may be present. Add to this the fact that perceptual commonalities are often highly correlated with deeper relational commonalities, particularly for basic-level terms, which young children are most likely to be hearing and learning. (For example, fins and gills are correlated with a different type of breathing apparatus than are legs and fur [e.g., Murphy & Medin, 1985].) Gentner (1989) described this correlation between perceptual and conceptual information as fitting the "world is kind" hypothesis. Because of these correlations, perceptual commonalities are far more likely to be helpful than harmful in guiding young children to deeper understanding. Surface commonalities act as initial invitations to compare, and thus point the way to deeper commonalities that naturally become salient during comparison.

Such an explanation is consistent with the striking findings of Samuelson and Smith (2000). They showed experimentally that inducing a shape bias actually increases children's rate of vocabulary acquisition outside the lab. They taught children to attend to shape by teaching them object names that were organized by common shape. Children given this experimentally induced shape bias showed greater gains in vocabulary (as assessed by the MacArthur CDI parental checklist) over the weeks following the training session than did matched children who either received no training or received training with varied patterns of input. These results would be baffling if we assumed that children stop at noticing perceptual commonalities. But in light of the earlier discussion, children's use of perceptual features as a basis for word extension may in fact be highly adaptive. When children lack knowledge about a category, perceptual commonalities may serve as way of gathering exemplars that can then be compared more deeply to extract relational commonalities[7] that can serve as core category knowledge.

Alignment in Early Learning

The analogical encoding process offers an escape from the conundrum that confronts us in trying to explain early experiential learning. Because it can lead to at least modest learning, even with only partially understood cases, it removes the need for a well-understood prior analog. Consistent with this line of reasoning, there is evidence that young infants can benefit from close comparisons. For example, Oakes and Ribar (in press) found that 4- to 6-month-old infants more readily form perceptual categories such as *dog* (and discriminate dogs from perceptually similar cats) when the infants are given the opportunity to view and compare objects in pairs than when the objects are presented one at a time. Namy, Smith, and Gershkoff-Stowe (1997) found that 18-month-old infants who were encouraged to directly compare items exhibited stronger categorical responding in a sequential touching-type task than did infants who were not encouraged to compare the items. Finally, Oakes and her colleagues found that prior opportunities to compare items facilitated the performance of 13- and 16-month-old infants on a sequential-touching task (Oakes & Madole, 2003). When given multiple exemplars of *people* and *sea animals*, intermixed on the table, the infants were more likely to show category-related touching if they had previously been presented with the two categories separated (thus facilitating within-category comparison).

[7]A similar interplay between initial perceptual object similarity and relational similarity has been observed in adult learning (Brooks, 1987; Ross, 1999; Ross, Perkins, & Tenpenny, 1990).

There is also evidence consistent with another claim made earlier: that early in development, comparison-based learning thrives on rich commonalities—both object commonalities and relational commonalities. Booth and Waxman (2002) showed that 14-month-olds given an optimal comparison sequence can learn better when given pairs characterized by both relational commonalities and property commonalities than when given pairs that only share object properties.[8] They taught infants categories like "oval green-blue objects with a loop on top, which can be swung from a hook." One group of infants had the opportunity to compare the objects carrying out their functions. Infants were first shown a pair of in-category objects, for each of which the experimenter demonstrated the function (swinging from a hook), saying, "Look what I can do with this one." After repeating this sequence, the experimenter showed another pair of in-category objects and repeated the same procedure. The experimenter next showed a contrasting object, a spool-shaped object with no loop on top, and demonstrated that the object could not perform the function. Next, the experimenter brought out the target object (another member of the initial category), demonstrated its function, and asked the child to "find another one." There were two alternatives, one from the initial category and one from a contrasting category. Under these rich comparison conditions, the 14-month-olds (and also a group of 18-month-olds) chose the same-category item about 70% of the time, significantly above chance. A second group of infants was given the same sequence, except that instead of demonstrating the function of the objects, the experimenter called them each by the same name. The 18-month-olds showed above-chance selection of the same-category object (roughly 60%), but the 14-month-olds did not.

These results suggest, first, that rich matches are easier than sparse matches for infants. The 14-month-olds did far better when the overall match was very strong, consisting of both object and relational commonalities, than when given a pure object match. This finding is consistent with the claim that early learning is conservative, and more importantly, with the

[8]Booth and Waxman's (2002) categories consisted of closely similar, but not identical, objects; there were small differences in color and shape. This probably made the categories more challenging to the 14-month-olds.

claim that relational learning can be bootstrapped by experience. Infants at this age would not succeed at a purely relational comparison (e.g., to learn a common function across a set of completely different objects), but they *can* learn a rich literal category that includes relational as well as property information. As countless studies have shown, once relational information is present, even if it is initially bound to particular object properties, it can be abstracted away via further comparisons (e.g., Gentner et al., 2003; Gick & Holyoak, 1983).

A second implication concerns the role of language. Although language comes to act as an invitation to compare, the results here suggest (not surprisingly) that rich overall similarity operates earlier in development as an incitement to compare. It is possible that common language comes to signal a comparison opportunity through the child's experience that highly similar objects have the same names.

GENERAL DISCUSSION

Relational categories are an important aspect of human cognition. My central thesis here is that these categories can be learned—although not by simple association models. I have suggested that comparison processes, either spontaneous or invited by language, drive much of this learning process.

This research suggests several conclusions. First, it supports the career of similarity thesis: Children begin with highly concrete overall similarity matches[9] and gradually become able to appreciate partial matches. Second, importantly, these early matches can involve relational as well as object commonalities. Third, among partial matches there is a relational shift from an early ability to match objects and object properties to a later ability to perceive purely relational commonalities. Fourth, this development is driven in large part by changes in domain knowledge. Fifth, by the second year of life, comparison processes can be invited not only by high similarity and close spatiotemporal juxtaposition, but also by the presence of common linguistic labels. The first of these represents alignment through experiential juxtaposition; the second, alignment through symbolic juxtaposition.

An important aspect of this proposal is that the increase in the sophistication of children's concepts does not result from global changes in logical processes or processing capacity. Such changes may occur, but they are not

[9]Early similarity is often described as "holistic." My sense of what this means is that early in development, alignment can only succeed with very strong overall similarity, in part because the infant's representations are idiosyncratic, lacking a uniform set of dimensions and properties. As infants' knowledge becomes more stable, they become able to match things that are only partially alike.

necessary to explain the relational shift. Rather, the shift comes about through gains in relational knowledge (sometimes gained via comparison). Numerous simulations have demonstrated that the relational shift can be modeled by applying the same structure-mapping processes to knowledge that varies in relational depth (e.g., Gentner, Rattermann, Markman, & Kotovsky, 1995; Loewenstein & Gentner, in press). The structure-mapping process grades naturally from highly concrete, literally similar comparisons to purely relational comparisons. Thus it can span the developmental course from overall similarity to relational similarity and abstract mappings.

Correspondence and Coherence

The contrast between *object categories* (defined chiefly by intrinsic features) and *relational categories* (defined chiefly by relations to other concepts) is related to the distinction between correspondence-based concepts (defined by reference to things in the world) and coherence-based concepts (defined by their relations with other concepts). Like the distinction between object categories and relational categories, this distinction is actually a continuum, but here too it is useful to contrast the end points. The correspondence-based view figures prominently in theories of concepts that conceive of word meaning as pointers to the world (e.g., Frege, 1892/1980; Russell, 1905/1956). The coherence-driven view is epitomized in Saussure's writings on language; for example, in the idea that "Language is a system of interdependent terms in which the value of each term results solely from the simultaneous presence of the others . . ." (Saussure, 1916/1966, p. 114). It seems obvious that both coherence and correspondence enter into human conceptual structure (see Markman, 1999). What is more interesting is the way they enter in. I suggest that concrete object concepts are heavily correspondence driven. They arise naturally from the way our perceptual capacities operate on the experiential world. Relational terms, like verbs and relational nouns, are relatively more coherence driven; there is a greater role for cultural and linguistic structures in determining the way information is organized into concepts.

The Role of Language

A constant theme in this chapter has been the importance of relational language in driving the development of relational category knowledge (Gentner, 2003; Gentner & Loewenstein, 2002). In this chapter I focused chiefly on the role of common language in inviting comparison, thereby promoting the highlighting of relational commonalities and a more uniform encoding of exemplars through rerepresentation. By fostering uniform relation representation, common relational language promotes transfer to new situations (Clement et al., 1994; Forbus, Gentner, & Law, 1995). As children gradually shed the concrete details of their initial representation—learning, for example, that taxis do not have to be yellow, nor islands sandy—they become able to transfer their relational concepts more broadly.[10]

Summary

Learning relational categories, and the relational information that belongs to object categories, is a challenge to learners. Relations are not preindividuated by our cognitive systems. They can be encoded in dozens of different ways, and children have to learn the ways that work in their physical, cultural, and linguistic environment. But this learning can be achieved via general learning mechanisms available to humans; in particular, by structure-mapping processes. These processes are used in concert with other general processes such as associative learning and attentional mechanisms, and are guided by linguistic labels and social interactions that augment experiential knowledge and serve to invite comparison when intrinsic similarity is not enough. Progressive alignment promotes representational uniformity, and increases the likelihood that the learner will encode new relational situations in the same way across different situations. This process, which Gentner and Rattermann (1991) referred to as "the gentrification of knowledge," entails some loss of the child's immediate perception of the world. But the resulting gains in representational uniformity are important in achieving stable conceptual structure and a systematic and portable set of relational abstractions.

ACKNOWLEDGMENTS

This research was supported by NSF grant SBR-95-11757 and by NSF-ROLE award 21002/REC-0087516. The computational research was supported by ONR contract N00014-92-J-1098. I thank Melissa Bowerman for extremely helpful comments on an earlier draft, and Ken Forbus, Jeff Loewenstein, Ken Kurtz, Art Markman, Mary Jo Rattermann, Phillip Wolff, and the Analogy and Similarity group at Northwestern University for discussions of these issues. I also thank Kathleen Braun for assistance on the research and Jonathan Cohen for help in preparing the manuscript.

[10]This gradual abstraction need not entail losing the initial concrete representations. My assumption is that human representations are pluralistic; that, for example, we retain the concrete notion of taxis as yellow cabs along with the more abstract taxi concept that allows such uses as "water taxi."

REFERENCES

Ahn, W. K. (1998). Why are different features central for natural kinds and artifacts?: The role of causal status in determining feature centrality. *Cognition, 69,* 135–178.

Ahn, W. K., Gelman, S. A., Amsterlaw, J. A., Hohenstein, J., & Kalish, C. W. (2000). Causal status effect in children's categorization. *Cognition, 76,* 35–44.

Asmuth, J. A., & Gentner, D. (2004). *Context sensitivity of relational nouns.* Unpublished manuscript, Northwestern University.

Au, T. K., Dapretto, M., & Song, Y. (1994). Input vs. constraints: Early words acquisition in Korean and English. *Journal of Memory and Language, 33*(5), 567–582.

Baldwin, D. A. (1989). Priorities in children's expectations about object label reference: Form over color. *Child Development, 60,* 1289–1306.

Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory & Cognition, 15,* 397–418.

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

Booth, A. E., & Waxman, S. R. (2002). Object names and object functions serve as cues to categories for infants. *Developmental Psychology, 38,* 948–957.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization* (pp. 141–174). Cambridge, England: Cambridge University Press.

Brown, R. (1958). *Words and things: An introduction to language.* New York: The Free Press.

Callanan, M. A. (1990). Parents' descriptions of objects: Potential for children's inferences about category principles. *Cognitive Development, 5,* 101–122.

Carey, S. (1992). The origin and evolution of everyday concepts. In R. N. Giere & H. Feigl (Eds.), *Cognitive models of science: Minnesota studies in the philosophy of science* (pp. 89–128). Minneapolis, MN: University of Minnesota Press.

Caselli, M. C., Bates, E., Casadio, P., Fenson, J., Fenson, L., Sanderl, L., & Weir, J. (1995). A cross-linguistic study of early lexical development. *Cognitive Development, 10,* 159–199.

Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the Control of Variables Strategy. *Child Development, 70*(5), 1098–1120.

Clark, E. V. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 65–110). New York: Academic Press.

Clark, E. V. (1993). *The lexicon in acquisition.* Cambridge, England: Cambridge University Press.

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

Croft, W. (2001). *Radical construction grammar.* Oxford: Oxford University Press.

Dunbar, K. (1993). Concept discovery in a scientific domain. *Cognitive Science, 17,* 391–434.

Earles, J. L., & Kersten, A. W. (2000). Adult age differences in memory for verbs and nouns. *Aging Neuropsychology & Cognition, 7*(2), 130–139.

Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation* (pp. 157–196). Los Altos, CA: Kaufmann.

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

Forbus, K. (2001). Exploring analogy in the large. In D. Gentner, K. J. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 23–58). Cambridge, MA: MIT Press.

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

Frege, G. (1980). On sense and meaning. In P. T. Geach & M. Black (Eds.), M. Black (Trans.), *Translations from the philosophical writings of Gottlob Frege* (3rd ed.). Oxford, England: Basil Blackwell. (Original work published 1892)

Gelman, R., Spelke, E. S., & Meck, E. (1983). What preschoolers know about animate and inanimate objects. In D. Rogers & J. A. Sobaoda (Eds.), *The acquisition of symbolic skills* (pp. 297–326). New York: Plenum.

Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology, 20,* 65–95.

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

Gentner, D. (1982). Why nouns are learned before verbs: Relativity vs. natural partitioning. In S. A. Kuczaj (Ed.), *Language development: Syntax and semantics.* Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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

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

Gentner, D. (2002). Analogy in scientific discovery: The case of Johannes Kepler. In L. Magnani & N. J. Nersessian (Eds.), *Model-based reasoning: Science, technology, values* (pp. 21–39). New York: Kluwer Academic/Plenum.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and cognition* (pp. 195–235). Cambridge, MA: MIT Press.

Gentner, D. (in press). Nouns and verbs, revisited. *APA PsycEXTRA.*

Gentner, D., & Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 215–256). New York: Cambridge University Press.

Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 343–382). San Mateo, CA: Kaufmann.

Gentner, D., & Klibanoff, R. S. (2001). *On acquiring* gift*: The acquisition of relational nouns.* Unpublished manuscript, Northwestern University.

Gentner, D., & Kurtz, K. J. (in press). Learning and using relational categories. In R. Goldstone & W. Ahn (Eds.), *Categorization inside and outside the lab: A Festschrift in honor of Dr. Douglas Medin.*

Gentner, D., & Loewenstein, J. (2002). Relational language and relational thought. In J. Byrnes & E. Amsel (Eds.), *Language, literacy, and cognitive development* (pp. 87–120). Mahwah, NJ: Lawrence Erlbaum Associates.

Gentner, D., Loewenstein, J., & Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology, 95*(2), 393–408.

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

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

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

Gentner, D., & Namy, L. (2004). The role of comparison in children's early word learning. In D. G. Hall & S. R. Waxman (Eds.), *Weaving a lexicon* (pp. 597–639). Cambridge, MA: MIT Press.

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., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability and inferential soundness. *Cognitive Psychology, 25,* 524–575.

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

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

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

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

Gopnik, A., & Choi, S. (1995). Names, relational words, and cognitive development in English and Korean speakers: Nouns are not always learned before verbs. In M. Tomasello & W. E. Merriman (Eds.), *Beyond names for things: Young children's acquisition of verbs* (pp. 63–80). Hillsdale, NJ: Lawrence Erlbaum Associates.

Goswami, U., & Brown, A. L. (1989). Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition, 35,* 69–95.

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

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

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

Hall, D. G., & Waxman, S. R. (1993). Assumptions about word meaning: Individuation and basic-level kinds. *Child Development, 64*(5), 1550–1570.

Holyoak, K. J., & Thagard, P. R. (1995). *Mental leaps: Analogy in creative thought.* Cambridge, MA: MIT Press.

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

Imai, M., Gentner, D., & Uchida, N. (1994). Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive Development, 9,* 45–75.

Keil, F. C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind* (pp. 234–254). New York: Cambridge University Press.

Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning & Verbal Behavior, 23,* 221–236.

Kersten, A. W., & Earles, J. L. (in press). Semantic context influences memory for verbs more than memory for nouns. *Memory & Cognition.*

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

Kurtz, K. J., & Gentner, D. (2001). Kinds of kinds: Sources of category coherence. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society,* 522–527.

Kurtz, K. J., Miao, C., & Gentner, D. (2001). Learning by analogical bootstrapping. *Journal of the Learning Sciences, 10*(4), 417–446.
Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3,* 299–321.
Langacker, R. W. (1987). Nouns and verbs. *Language, 63,* 53–94.
Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development, 2*(2), 189–219.
Loewenstein, J., & Gentner, D. (in press). Relational language and the development of relational mapping. *Cognitive Psychology.*
Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review, 6*(4), 586–597.
Loewenstein, J., Thompson, L., & Gentner, D. (2003). Analogical learning in negotiation teams: Comparing cases promotes learning and transfer. *Academy of Management Learning and Education, 2*(2), 119–127.
Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Lawrence Erlbaum Associates.
Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431–467.
Markman, A. B., & Gentner, D. (2000). Structure-mapping in the comparison process. *American Journal of Psychology, 113*(4), 501–538.
Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Intelligence, 13,* 329–358.
Markman, E. M. (1989). *Categorization and naming in children: Problems of induction.* Cambridge, MA: MIT Press.
Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100*(2), 254–278.
Morimoto, Y. (1999). Loan words and their implications for the categorial status of verbal nouns. In S. S. Chang, L. Liaw, & J. Ruppenhofer (Eds.), *Proceedings of the twenty-fifth annual meeting of the Berkeley linguistics society, Feb. 12–15, 1999: General Session and Parasession on Loan Word Phenomena (BLS 25)* (pp. 371–382). Berkeley: Berkeley Linguistics Society.
Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289–316.
Namy, L. L., & Gentner, D. (2002). Making a silk purse out of two sow's ears: Young children's use of comparison in category learning. *Journal of Experimental Psychology: General, 131,* 5–15.
Namy, L. L., Smith, L. B., & Gershkoff-Stowe, L. (1997). Young children's discovery of spatial classification. *Cognitive Development, 12,* 163–184.
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, MN: University of Minnesota Press.
Oakes, L. M., & Madole, K. L. (2003). Principles of developmental change in infants' category formation. In D. H. Rakison & L. M. Oakes (Eds.), *Early concept and category development: Making sense of the blooming, buzzing confusion* (pp. 132–158). New York: Oxford University Press.
Oakes, L. M., & Ribar, R. J. (in press). A comparison of infants' categorization in paired and successive presentation familiarization tasks. *Infancy.*
Pae, S. (1993). *Early vocabulary in Korean: Are nouns easier to learn than verbs?* Unpublished doctoral dissertation, University of Kansas, Lawrence.
Pavlicic, T., & Markman, A. B. (1997). The structure of the verb lexicon: Evidence from a structural alignment approach to similarity. *Proceedings of the 19th Annual Conference of the Cognitive Science Society,* Stanford, CA.
Rattermann, M. J., & Gentner, D. (1998a). More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task. *Cognitive Development, 13,* 453–478.

Rattermann, M. J., & Gentner, D. (1998b). The effect of language on similarity: The use of relational labels improves young children's performance in a mapping task. In K. Holyoak, D. Gentner, & B. Kokinov (Eds.), *Advances in analogy research: Integration of theory & data from the cognitive, computational, and neural sciences* (pp. 274–282). Sofia, Bulgaria: New Bulgarian University.
Rattermann, M. J., & Gentner, D. (2001). *The effect of language on similarity: The use of relational labels improves young children's analogical mapping performance.* Unpublished manuscript, Northwestern University.
Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General, 130,* 323–360.
Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(4), 629–639.
Ross, B. H. (1999). Postclassification category use: The effects of learning to use categories after learning to classify. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*(3), 743–757.
Ross, B. H., Perkins, S. J., & Tenpenny, P. L. (1990). Reminding-based category learning. *Cognitive Psychology, 22,* 460–492.
Russell, B. (1956). On denoting. In R. C. Marsh (Ed.), *Logic and knowledge* (pp. 39–56). London: Routledge. (Original work published 1905)
Samuelson, L. K., & Smith, L. B. (2000). Grounding development in cognitive processes. *Child Development, 71*(1), 98–106.
Saussure, F. de. (1966). *Course in general linguistics* (C. Bally & A. Sechehaye, Eds., Baskin, Trans.). New York: McGraw Hill. (Original work published 1916)
Shatz, M. (1994). *A toddler's life: Becoming a person.* New York: Oxford.
Siegler, R. S. (1989). Mechanisms of cognitive development. *Annual Review of Psychology, 40,* 353–379.
Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22,* 189–228.
Smith, L. B. (1983). Development of classification: The use of similarity and dimensional relations. *Journal of Experimental Child Psychology, 36,* 150–178.
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.
Sobin, N. (1982). Texas Spanish and lexical borrowing. In J. Amastae & L. Elias-Olivares (Eds.), *Spanish in the United States: Sociolinguistic aspects* (pp. 166–181). Cambridge: Cambridge University Press.
Talmy, L. (1975). Semantics and syntax of motion. In J. Kimball (Ed.), *Syntax and semantics* (Vol. 4, pp. 181–238). New York: Academic Press.
Tardif, T. (1996). Nouns are not always learned before verbs: Evidence from Mandarin speakers' early vocabularies. *Developmental Psychology, 32*(3), 492–504.
Tardif, T., Gelman, S. A., & Xu, F. (1999). Putting the noun bias in context: A comparison of English and Mandarin. *Child Development, 70*(3), 620–635.
Tversky, A. (1977). Features of similarity. *Psychological Review, 84*(4), 327–352.
Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development, 5,* 123–150.
Waxman, S. R., & Klibanoff, R. S. (2000). The role of comparison in the extension of novel adjectives. *Developmental Psychology, 36,* 571–581.
Yan, J., Forbus, K., & Gentner, D. (2003). A theory of rerepresentation in analogical matching. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Meeting of the Cognitive Science Society,* 1265–1270.