
Why We're So Smart

Dedre Gentner

8.1 Introduction

Human cognitive abilities are remarkable. Our mental agility has allowed us to adapt to a vast range of environments, and even to adapt our environments to suit ourselves. We've been clever enough to eradicate most of our competitors and to spread over most of the earth. Indeed, we now find ourselves in the ironic position of striving to preserve a few of our formerly fearsome predators.

What is the nature of this ability? A list of the cognitive skills that distinguish us would include

- The ability to draw abstractions from particulars—to generalize experience and store regularities across vastly different cases
- The ability to maintain hierarchies of abstraction, so that we can store information about Fido, about dachshunds, about dogs, or about living things
- The ability to concatenate assertions and arrive at a new conclusion
- The ability to reason outside of the current context—to think about different locations and different times and even to reason hypothetically about different possible worlds
- The ability to compare and contrast two representations to discover where they are consistent and where they differ
- The ability to reason analogically—to notice common relations across different situations and project further inferences
- The ability to learn and use external symbols to represent numerical, spatial, or conceptual information

Our language abilities are equally outstanding. They include our ability to learn a generative, recursive grammar, as well as a set of semantic-conceptual abilities:

- The ability to learn symbols that lack any iconic relation to their referents
- The ability to learn and use symbols whose meanings are defined in terms of other learned symbols, including even recursive symbols such as the set of all sets
- The ability to invent and learn terms for abstractions as well as for concrete entities
- The ability to invent and learn terms for relations as well as things

To what do we owe these powers? Are they multiple separate abilities, or is there a core set of abilities that engenders them all? At least three sources of our superiority have been proposed. One is innate domain theories: perhaps our starting knowledge state is qualitatively superior to that of other animals. Another possibility is innate processing abilities: we might possess larger processing capacity and/or more powerful learning mechanisms than other animals. A third possibility is that it is participation in human language and culture that gives us our edge. I will argue for the latter two possibilities. This is not to deny the considerable evidence that human infants are born with built-in attentional capacities and tacit expectations about the physical world and about social interactions. However, it's likely that many of those capacities are shared with other higher animals, particularly social animals.¹ The question here is what makes us smarter than the rest.

My thesis is this: what makes humans smart is (1) our exceptional ability to learn by analogy, (2) the possession of symbol systems such as language and mathematics, and (3) a relation of mutual causation between them whereby our analogical prowess is multiplied by the possession of relational language. My argument has three parts. First, relational concepts are critical to higher-order cognition, but relational concepts are both nonobvious in initial learning and elusive in memory retrieval. Second, analogy is the mechanism by which relational knowledge is revealed. Third, language serves both to invite learning relational concepts and to provide cognitive stability once they are learned. In

short, analogy is the key to conceptual learning, and relational language is the key to analogy.

My case for the importance of language rests on an account of higher-order learning that my colleagues and I have been developing for the last two decades. I begin by laying out a general account of cognitive development, emphasizing what we have called "the career of similarity" (Gentner and Rattermann 1991). I then discuss the symbiotic development of analogy and relational language.

8.2 The Career of Similarity

In his career-of-similarity hypothesis, Quine (1960) proposed that over development children move from perceiving only "brute" perceptual similarity to perceiving more sophisticated likenesses—"theoretical" similarity. The career of similarity has wide ramifications. Virtually every cognitive process, from categorization to transfer, is influenced by explicit or implicit similarity comparisons.

Gentner and Rattermann (1991) amplified this account to propose a developmental progression (1) from simply responding to overall similarity to attending to selective similarity; and (2) within selective similarity, from a focus on object similarity to a focus on relational similarity, and from perceptual commonalities to conceptual commonalities. We reviewed evidence to suggest that a major driver of this relational shift in similarity is changes in children's knowledge—particularly the acquisition of higher-order relational knowledge.

On this account, the career of similarity exists in a relation of mutual causation with the nature of children's representations. Children's initial knowledge representations differ from adult representations in being (1) more situation-specific, (2) more perceptual, and (3) more variable. There is abundant evidence for the claim that early representations contain relatively more situation-specific perceptual information than do adult representations (e.g., Rovee-Collier and Fagen 1981; Landau, Smith, and Jones 1988). The third claim, early variability, requires some explanation. What I mean is that early representations are highly variable across contexts, even within the same child. That is, different representations will be invoked at different times for situations that an

adult would encode in like terms. I hypothesize that even something as "stable" in adult life as the neighbor's dog coming to the fence may be represented differently on different occasions by a very young child.

These representational claims have implications for the career of similarity. The instability of early representations implies that children's earliest similarity matches should be highly conservative: that is, babies should perceive similarity only when there is a large degree of overlap. This is often stated as the claim of holistic similarity in babies and thought to arise from perceptual specificity. I suggest that the holistic similarity arises only in part from early perceptual specificity—it mainly results from the variability of early representations. For if different construals are possible even for the same object, then only if there is an extremely high degree of potential overlap will the child's representations overlap enough for him to perceive similarity.

There is considerable evidence for the claim of early conservative similarity (Gentner and Medina 1998; Gentner and Rattermann 1991; Smith 1989). For example, Chen, Sanchez, and Campbell (1997) found that 10-month-old infants could learn to pull on a cloth to reach a toy, but they failed to transfer to a new pulling situation unless it was highly similar to the previously experienced situation. By 13 months, infants were able to transfer with less concrete similarity. In studies of infants' causal reasoning, Oakes and Cohen (1990) found evidence both for early conservatism and for an early focus on objects over relations. They investigated 6- and 10-month-old infants' perception of launching events by varying spatial and temporal features that should render such events either causal or noncausal. The results showed that infants at 10 months, but not 6 months, discriminated the events on the basis of causal relations. The younger infants appeared to respond on the basis of the individual objects in the events, but not the causal relationship between objects. Their results suggest that infants' perception of causal relations appears gradually and that it is initially very conservative—that is, specific to the kinds of objects included in the event.

The claim that infants are extremely conservative learners may seem wildly implausible in view of the rapidity of human learning. I suggest that far from being a disadvantage, early conservatism is necessary in order for humans to be appropriately flexible learners. The fact that

early comparisons are extremely conservative allows for emergent abstractions. This brings me to the second part of the causal interaction: the influence of similarity comparisons on representation.

I began by stating that analogical processing is central in human cognition. By "analogical processing" I do not mean only the perception of distant similarity in which only the relations match. Rather, I include the kind of mundane similarity comparisons that involve common entities as well as common relations. Such comparisons are a driving force in children's learning. The course of emergent abstractions depends crucially on the way structure-mapping operates: on which similarity matches are relatively easy and inevitable, and on the results of carrying out comparisons. I will give a brief review of structure-mapping in section 8.3, but for now I focus on three key points.

First, the structure-mapping process is sensitive to both object similarity and relational similarity, but favors common relational structure because of a tacit preference for systematicity (connectivity) and depth in the matching process. Second, the same process of structural alignment and mapping is used for mundane literal similarity as for analogy. Literal similarity (overall similarity) comparisons are easy to compute, because the object and relational matches are all mutually supportive. Analogical matches are more difficult, because the relational correspondences are not supported by object matches, and may even be opposed by them.² Third, carrying out any comparison—even a literal similarity comparison—tends to render common relations more salient; thus, even a literal comparison facilitates carrying out a later analogy that is based on the same relational structure.

Overall (literal) similarity matches are the easiest matches to notice and process. Because both object correspondences and relation correspondences enter into the one consistent alignment, such matches can readily be aligned even early in learning, when representational variability is high. This, I believe, is the underlying reason that young children rely on holistic similarity. When representations are variable, only a rich, overdetermined match has the redundancy necessary to guarantee finding an alignment, and hence strong overall matches constitute the earliest reliable similarity matches (Kempner 1983; Smith 1989). The privileges of overall literal similarity do not end with childhood. There is

evidence that adults also process literal similarity matches faster than purely relational matches (Gentner and Kurtz, in preparation; Kurtz and Gentner 1998) and high-similarity matches faster than low-similarity matches (Wolff and Gentner 2000). Further, rich concrete matches, such as two identical dachshunds, are perceived as more similar than sparse concrete matches, such as two identical circles, by both children and adults (Gentner and Rattermann 1991; Tversky 1977).

Comparison processes can be prompted in several ways. Some comparisons are invited explicitly, when likenesses are pointed out by adults: for example, "That's a wolf. It's like a dog, except it's wilder." Some comparisons are invited by the fact that two situations have a common linguistic label (e.g., "These are both lamps"). Some arise from the child's spontaneous noticing of similarity. Some are engineered by the child, as in the circular reaction (Piaget 1952). An infant notices something interesting and then tries to repeat it again and again. This fascination with immediate repetition, I suggest, is a manifestation of comparison in learning. Such close repetition with variation may provide an ideal learning experience for infants.

Although comparison is an inborn process, its manifestation—whether a sense of sameness is perceived for a given pair of potential analogues—depends on how the situations are represented, and this in turn depends on experience. The conservative learning thesis implies that most of children's early spontaneous comparisons are mundane by adult lights. Early in learning, comparisons are made only between situations that match overwhelmingly. These close comparisons yield small insights; they render small differences between the situations salient; and they result in marginally more abstract representations that can then participate in more distant comparisons. As similarity comparisons evolve from being perceptual and context bound to becoming increasingly sensitive to common relational structure, children show an increasing capacity to reason at the level of abstract commonalities and rules.

Comparison is a major means by which children go beyond their early situated representations. Comparisons among exemplars—initially concrete, but progressively more abstract—promote abstraction and rule learning. Such learning provides a route by which children learn the

theory-like relational information that informs adult concepts. For example, children come to know that both tigers and sharks are carnivores, while deer and hippopotamuses are herbivores, that tigers prey on deer, and so on; they learn that a taxi is not defined as a yellow car but as a vehicle that can be hired (Keil and Batterman 1984). The structure-mapping process is central in this evolution in part because it allows learners to discover commonalities. More importantly, as noted above, the structure-mapping process promotes relational commonalities over common object properties. This is important because objects are more cognitively and perceptually salient than relations in the information structure of the perceived world (Gentner 1982; Gentner and Boroditsky 2001). Objects (or more precisely, complex structural objects) are relatively easy to individuate; they are learned early; and even adults are swayed by object matches in contexts where relational matches would clearly be more useful. The great value of analogy—and of structure-mapping processes even when applied to literal comparisons—lies in creating a focus on common relational systems and thus lifting a relational pattern away from its object arguments.

I have argued so far (1) that relational learning is important to the development of cognition and (2) that it proceeds via structure-mapping processes. Because many of the specific processing claims made here require an understanding of structure-mapping, before turning to the role of language I review the theory and simulation here.

8.3 Structure-Mapping: A Brief Review

Structure-mapping theory postulates that the comparison process is one of alignment and mapping between structured conceptual representations (Falkenhainer, Forbus, and Gentner 1989; Gentner 1983, 1989; Gentner and Markman 1997; Goldstone 1994; Goldstone and Medin 1994; Markman and Gentner 1993; Medin, Goldstone, and Gentner 1993). The commonalities and differences between two situations are found by determining the maximal structurally consistent alignment between their representations. A structurally consistent alignment is characterized by one-to-one mapping (i.e., an element in one representation can correspond to at most one element in the other representation) and

parallel connectivity (i.e., if elements correspond across the two representations, then the elements they govern must correspond as well). When more than one structurally consistent match exists between two representations, contextual relevance and the relative systematicity of the competing interpretations are used. All else being equal, the richest and deepest relational match is preferred (the systematicity principle). The alignment process favors deep systems over shallow systems, even if they have equal numbers of matches (Forbus and Gentner 1989). Finally, predicates connected to the common structure in the base, but not initially present in the target, are proposed as candidate inferences in the target. Because these inferences are the structural completion at the best match between the terms, and because the best match is highly likely to be a deep relational match, the candidate inferences are often causally informative. Thus, structure-mapping processes can lead to spontaneous but informative inferences.

Sentences (1)–(5) show different kinds of similarity matches:

- (1) The dog chased the cat.
- (2) The coyote chased the lynx.
- (3) The shark chased the mackerel.
- (4) Amalgamated Tire Co. made a takeover bid for Racine Ironworks.
- (5) The cat chased the mouse.

Because matches at all levels enter into the alignment process, the easiest comparisons should be those of rich overall (literal) similarity. A concrete match like {(1) and (2)}—in which both the objects and the relations match—is intuitively easier to process than a less similar abstract match like {(1) and (3)}, or—yet more challengingly—{(1) and (4)}. For pairs like (1) and (2), the comparison process is easy, because the matches are mutually supporting, yielding one clear dominant interpretation. As noted above, overall similarity matches—in which the object match and the relational matches are mutually supporting—are easier to process than purely analogical matches like {(1) and (3)}, in which there are often stray object-attribute matches that are inconsistent with the maximal relational match. A particularly difficult case is a cross-mapped analogy like {(1) and (5)}. A cross-mapped analogy (Gentner and Toupin

1986) contains an object match (e.g., cat → cat) that is inconsistent with the relational match. Such matches are more difficult for children to map than either literal similarity matches or standard analogies, because there are competing modes of alignment. However, despite the greater difficulty of a cross-mapping, older children and adults normally resolve them in favor of the relational match—evidence of a tacit preference for systematicity in analogical alignment (Gentner and Toupin 1986; Markman and Gentner 1993).³ Even adults often choose the object match when asked to state correspondences in cross-mapped pictorial scenes (Markman and Gentner 1993), especially under a high processing load (Kubose, Holyoak, and Hummel, in preparation). Finally, Goldstone (1994) found evidence that local matches dominate early and relational matches later in processing cross-mapped matches. A computational model of analogical mapping, the Structure-Mapping Engine (SME), provides a model of the processes of alignment and mapping.

This process model has important implications for learning and development. SME's alignment process, taken as a model of human processing, suggests that the act of carrying out a comparison promotes structural alignment and renders the common structure, especially relational structure, more salient (Gentner and Bowdle 2001; Markman and Gentner 1993; Wolff and Gentner 2000). It is also important that structure-mapping is accomplished with a process that begins blind and local. Achieving a deep structural alignment does not require advance knowledge of the point of the comparison. (If it did, it would be relatively useless as a developmental learning process.)

There are at least five ways in which the process of comparison can further the acquisition of knowledge: (1) highlighting and schema abstraction; (2) projection of candidate inferences—inviting inferences from one item to the other; (3) re-representation—altering one or both representations so as to improve the match (and thereby, as an important side effect, promoting representational uniformity); (4) promoting attention to relevant differences; and (5) restructuring—altering the domain structure of one domain in terms of the other (Gentner and Wolff 2000). These processes enable the child to learn abstract commonalities and to make relational inferences.

Highlighting commonalities may seem like a rather trivial learning process, but this is not true in the case of common relations. Considerable evidence shows that the process of promoting common relational structure invited by mutual alignment between closely similar items promotes learning and transfer. Because the alignment process renders common relational structure more salient, structural alignment promotes the disembedding of hitherto nonobvious common relational systems (Gentner and Namy 1999; Gick and Holyoak 1983; Loewenstein, Thompson, and Gentner 1999). Indeed, I suggest that comparison—that is, the process of structural alignment and mapping—is a learning mechanism powerful enough to acquire structured rulelike knowledge (Gentner and Medina 1998).

8.4 Why Relational Language Matters

Relational terms invite and preserve relational patterns that might otherwise be fleeting.⁴ Consider the terms in table 8.1, which range from spatial relations to causal relations to social communicative relations. A language lacking such terms would be unimaginably impoverished. In such a language, it would be prohibitively cumbersome to express complex predictions, conjectures, dichotomous chains of thought, hypothetical arguments, and so on—not to mention counterfactuals such as “If this language lacked relational terms, it would be more difficult to communicate ideas like this one.”

The sample of relational terms in table 8.1 suggests the range and utility of relational language. It includes verbs and prepositions—members of classes that are dedicated to conveying relational knowledge and that contrast with object reference terms on a number of grammatical and informational dimensions (Gentner 1981). However, it also includes a large number of relational nouns (e.g., *weapon*, *conduit*, and *barrier*), and these pose an interesting learning problem to which I return below.

However, although relational concepts are important, they are often not obvious. One reason that relational language is important in higher mental life is that, unlike object concepts, relational concepts are not automatically learned. Relational concepts are not simply given in the

Table 8.1
Some relational terms

Relational nouns	
<i>General relation terms</i>	<i>Terms incorporating similarity and logical relatedness</i>
cause	twin
prevention	equivalence
source	identity
result	converse
advantage	inverse
bane	prediction
ally	contradiction
accident	
<i>Terms of communication</i>	<i>Terms that range from concrete to abstract usage</i>
threat	weapon
lie	gift
promise	target
excuse	haven
pretext	screen
dispute	filter
	barrier
	conduit
	leeway
Verbs, prepositions, and general connectives	
cause	however
prevent	nevertheless
foster	therefore
engender	accordingly
permit	contrary
inhibit	except
deter	
accelerate	
force	

Table 8.2
Relational versus object-referential terms in the domain of biology

Relational terms	Object-reference terms
carnivore	tiger
scavenger	snail
prey	deer
parent	gorilla
weed	dandelion
parasite	flea

natural world: they are culturally and linguistically shaped (Bowerman 1996; Talmy 1975). This malleability is expressed in the relational relativity principle—that the parsing of the perceptual world into a relational lexicon differs more across languages than does that for object terms (Gentner 1982; Gentner and Boroditsky 2001). To bring home this second point, table 8.2 contrasts relational nouns with ordinary referential nouns within the domain of biology.

If relational language bears a nonobvious relation to the world, it follows that relational terms should be harder to learn than terms such as concrete nominals whose referential relations are more transparent. Indeed, there is considerable evidence that relational terms are hard to learn. One indication of the relative difficulty of learning relations is that verbs and prepositions enter children's vocabularies later than do concrete nouns (Gentner 1982; Gentner and Boroditsky 2001; Goldin-Meadow, Seligman, and Gelman 1976). Another indication is that the full meanings of verbs and other relational terms are acquired relatively slowly (Bowerman 1996; Olguin and Tomasello 1993). Words like *if* and *because* (Byrnes 1991; Scholnick and Wing 1982) or *buy*, *sell*, and *pay* (Gentner 1978) may not be fully understood until 8 or 9 years of age.

The difficulty of learning relational terms relative to object terms can be seen not only across form class—in the advantage of nouns over verbs—but also *within* the nominal class, in the acquisition of relational nouns. Relational nouns sometimes denote relations directly: for example, *symmetry*. More commonly, they denote categories whose membership is determined by a particular relation (either temporary or

enduring) that category members have with another entity or category: for example, *gift*, *weapon*, *friend*, *sister*, and *home*. Children often initially interpret relational terms as object reference terms, and only later come to appreciate the relational meaning (Gentner and Rattermann 1991). For example, kinship terms are often understood initially in terms of characteristics of individuals, and only later in terms of relational roles (Clark 1993). Likewise, Keil and Batterman (1984) found that 4-year-olds conceive of an *island* as a place with sand and palms, and of an *uncle* as a nice man with a pipe. Only later do they learn the relational descriptors, that an *island* is a body of land surrounded by water, and an *uncle*, any male in a sibling relationship with one's mother or father. 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 [referring to a doll] is a *blicket* BECAUSE IT IS RIDING IN A CAR," children tended to interpret the novel noun as referring to the object category and extended it to a similar-looking doll.

Gentner and Klibanoff (in preparation) tested preschool children's ability to learn relational meanings, using a combination of comparison and labeling to underscore the relational structure. In this study, 3-, 4-, and 6-year-olds were shown picture cards and heard a novel relational noun used in two parallel contexts: for example, "The knife is the *blick* for the watermelon, and the ax is the *blick* for the tree." Then they were asked to choose the "*blick*" in a third context: for example, "What would be the *blick* for the paper?" The children chose among three picture cards: *same relation* (correct; e.g., a pair of scissors), *thematic* (e.g., a pencil), and *same nominal category* (e.g., another piece of paper). Both 4- and 6-year-olds correctly chose the *same relation* card. However, 3-year-olds performed at chance in this task, despite the extensive guidance.

However, although relational language is hard to learn, the benefits outweigh the difficulty. Gentner and Loewenstein (2002) discuss several specific ways in which relational language can foster the learning and retention of relational patterns:

- *Abstraction.* Naming a relational pattern helps to abstract it—to de-situate it from its initial context. This increases the likelihood of seeing

the pattern elsewhere in another situation. We have obtained this effect in studies of mapping, as discussed later (Gentner and Rattermann 1991; Loewenstein and Gentner 1998, 2002; Rattermann and Gentner 1998, 2002).

- *Initial registration.* Hearing a relational term applied to a situation invites children to store the situation and its label, even before they fully understand the term's meaning. This is just to say that Roger Brown's "language as an invitation to form concepts" applies to relational concepts as well as to object concepts. Then, when further exemplars with the same label are encountered, there is a chance that comparison with the prior instance may promote a relational meaning, even when none was initially obvious. Hearing a relational term used across contexts invites abstracting its meaning. By giving two things the same name, we invite children to compare them, whether or not they occur in experiential juxtaposition. Thus, relational language creates symbolic juxtapositions that might not occur in the physical world.

- *Selectivity.* Once learned, relational terms afford not only abstraction but also selectivity. We focus on a different set of aspects and relations when we call a cat a *pet* from when we call him a *carnivore*, or a *good mouser*, or a *lap warmer*. Selective linguistic labeling can influence the construal of a situation. For example, a labeling manipulation can influence the degree of "functional fixedness" in an insight task (Glucksberg and Danks 1968; Glucksberg and Weisberg 1966).

- *Reification.* Using a relational term can reify an entire pattern, so that new assertions can be stated about it. A named relational schema can then serve as an argument to a higher-order proposition. For example, consider this sentence from the *New York Times Book Review*: "The economic adversity caused by droughts or floods far exceeds the direct impact on the food supply." The economy of expression made possible by the relational nouns *adversity*, *drought*, *flood*, and *impact*, as well as the higher-order connecting relations *cause* and *exceed*, makes it possible to state a complex embedded proposition compactly. Expressing such complex assertions as the above would be prohibitively awkward without such relational compaction.

- *Uniform relational encoding.* Habitual use of a given set of relational terms promotes uniform relational encoding, thereby increasing the

probability of transfer between like relational situations (Forbus, Gentner, and Law 1995). When a given domain is encoded in terms of a stable set of relational terms, the likelihood of matching new examples with stored exemplars that share relational structure is increased. Thus, habitual use of a stable system of relational language can increase the probability of relational reminding. In instructional situations, it can foster appropriate principle-based reminding and transfer, and mitigate the perennial bugaboos of retrieval: inert knowledge and surface-based retrieval. The growth of technical vocabulary in experts reflects the utility of possessing a uniform relational vocabulary.

8.4.1 Uniform Relational Structure, Retrieval, and Transfer

The claim that uniform relational language aids analogical retrieval is important, because analogical retrieval is generally quite poor. People routinely fail to be reminded of past experiences that are relationally similar to current experiences, even when such reminders would be useful in their current task, and even when it can be demonstrated that they have retained the prior knowledge (Gentner, Rattermann, and Forbus 1993; Gick and Holyoak 1980; Keane 1988; Ross 1989). There is evidence from studies in mathematics that this "inert knowledge" problem is less severe for experts than for novices (Novick 1988). Although Novick did not investigate the encoding vocabulary of the two groups, other studies of similarity-based retrieval have found a relation between the quality of the encoding (as assessed in participants' summaries of the materials) and the likelihood of relational retrieval (Gick and Holyoak 1983; Loewenstein, Thompson, and Gentner 1999). These results are consistent with the conjecture that one benefit of expertise is better, less idiosyncratic relational representations, which, as noted above, would promote relational retrieval.

More direct evidence that uniform relational language promotes transfer comes from studies by Clement, Mawby, and Giles (1994), who gave adults passages to read and later gave them new passages that were structurally similar but different in their specific characters and actions—the classic situation in which poor retrieval abilities have been demonstrated. For some learners, the parallel structure in the two matching passages was expressed using relational terms that had the same mean-

ings (e.g., X *ate* Y and A *consumed* B). For others, the parallel structure was expressed using nonsynonymous relational pairs (e.g., X *munched on* Y and A *gobbled up* B). This was a fairly subtle manipulation; the differing relational pairs were partly overlapping in meaning, so that they could readily have been aligned had the passages been seen together. However, even this minimal manipulation made a difference: people who received synonymous terms—such as *ate* and *consumed*—were more likely to retrieve the initial passage given the probe than those who received nonsynonymous pairs. Clement, Mawby, and Giles concluded that the use of common relational encoding can promote analogical retrieval in adults.

Does language—especially, use of uniform relational language— influence children’s memory retrieval? Some researchers have suggested that conversations with adults might be important in shaping children’s memories (Nelson 1996). Herbert and Hayne (2000) studied 18-month-old infants in a deferred imitation task. Children were shown how to rattle by putting a ball into a cup and shaking it. The key variable was what kind of language children heard during the first session: empty narration (e.g., “Let’s have a look at this ...”), actions only (e.g., “Push the ball into the cup ...”), goals only (e.g., “We can use these things to make a rattle. Let’s have a look at this ...”), or actions plus goals (e.g., “We can use these things to make a rattle. Push the ball into the cup ...”). The latter two groups also received a prompt before the test, reminding them that they could use these things to make a rattle. After four weeks, children were retested to see if they could still reproduce the actions. Only the group that heard action-plus-goal language was able to reproduce the action at above-baseline rates.

8.4.2 Relational Language in Cognitive Development

Relational language both invites comparison and preserves the results as a (relational) abstraction. Jeff Loewenstein, Mary Jo Rattermann, and I have sought empirical evidence for this claim. We have focused on spatial relations like *on*, *in*, and *under* (Loewenstein and Gentner 1998) and *symmetry* and *monotonicity* (Kotovsky and Gentner 1996; Rattermann and Gentner 1998, 2002). These kinds of spatial terms satisfy three criteria for an arena in which to investigate possible effects of language

on cognitive development: (1) they show substantial crosslinguistic variation, (2) they lend themselves to objective testing, and (3) they are accessible to children. The logic of our studies is first, to establish a challenging spatial relational task, and then to test whether language for spatial relations can improve children’s performance.

Rattermann and I tested the power of relational labels to promote relational insight, using a very simple mapping task (Gentner and Rattermann 1991; Rattermann and Gentner 1998, 2002). Children aged 3, 4, and 5 saw two triads of objects, the child’s set and the experimenter’s set, both arranged in monotonically increasing order according to size. As in DeLoache’s (1987, 1995) search studies, children watched as the experimenter hid a sticker under an object in the experimenter’s triad; they were told that they could find their sticker by looking “in the same place” in their triad. The correct response was always based on relational similarity: that is, the child had to find the object of the same relative size and position (smallest → smallest; middle → middle; etc.). Children were always shown the correct response after making their guess.

When the two sets were literally similar, 3-year-old children readily learned the mapping. But when the objects were shifted to a cross-mapped pattern, as in figure 8.1, so that the object matches were inconsistent with the best relational alignment (Gentner and Toupin 1986), the children had great difficulty grasping the relational match, particularly when the objects were rich and detailed. Indeed, in the rich-object cross-mapped versions of the task, 3- and 4-year-old children performed at chance (32%) even though they were shown the correct response on every trial (14 trials total).

Having thus established a difficult relational task, we then investigated whether providing relational language could help children perform this relational alignment. Before children carried out the cross-mapping task, they were provided with a brief training session in which we modeled using the labels *daddy*, *mommy*, and *baby* (or in other studies, *big*, *little*, *tiny*) for the characters in the two triads. We chose these family labels because they are often used spontaneously by preschool children to mark monotonic change in size (Smith 1989). The reasoning was that applying these labels to the three members of each triad would invite the child to

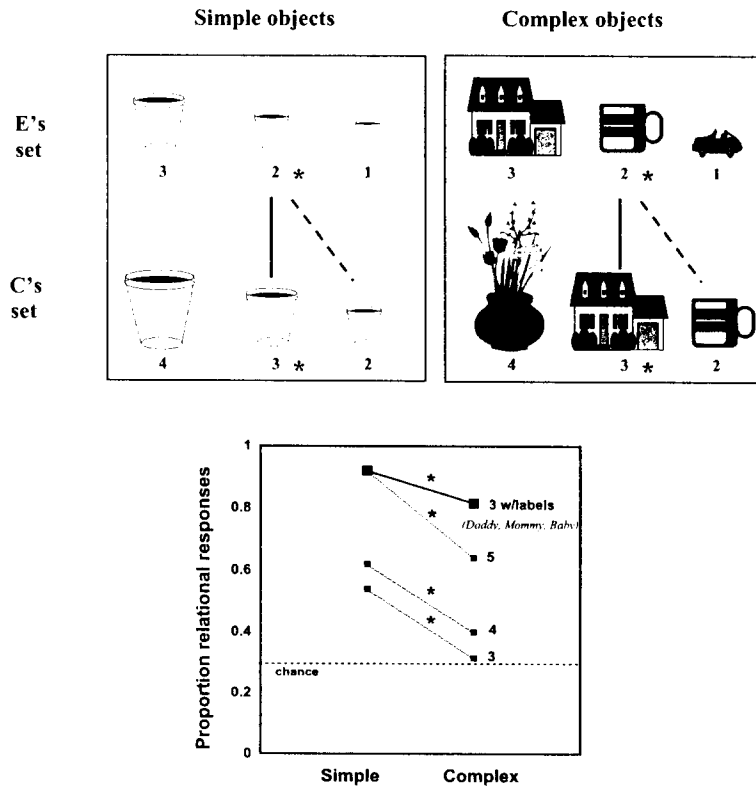


Figure 8.1

Materials and results for the Rattermann and Gentner spatial mapping task with cross-mapped objects, in which object matches compete with relational matches. In the top figures (*top*), asterisks and the solid line denote the correct match; the dashed line denotes the incorrect object match. The graph (*bottom*) shows the results for 3-, 4-, and 5-year-old children in the baseline condition and for 3-year-old children given the relational labels *daddy*, *mommy*, *baby*.

highlight the higher-order relational pattern of monotonic increase that forms the essential common system to align.

The results of the labeling manipulation were striking. The 3-year-olds given relational language performed well in the cross-mapping task on both the sparse (89% relational responding) and rich (79% relational responding) stimuli, as compared to performance rates of 54% and 32% without relational language. In fact, 3-year-olds given relational language performed on a par with 5-year-olds in the baseline condition. Further, 3-year-old children were fairly able to transfer their learning to new triads with no further use of the labels by the experimenters. That the improvement was specific to relational labels and was not just some general attentional effect of using language is shown by the fact that other relational labels denoting monotonic size-change (e.g., *big*, *little*, *tiny*) also improved performance, while neutral object labels (e.g., *jiggy*, *gimli*, *fantan* or *Freddy*, *Max*, *Bobby*) did not. Finally, when the children were brought back to the laboratory four to six weeks later, the group with relational language experience continued to show benefits of having represented the higher-order relational structure; they were better able to carry out the mapping task than their counterparts without relational language training. We suggest that the use of common relational labels prompted children to notice and represent the common higher-order relation of monotonic increase—in other words, that this facilitated making the relational alignment.

More evidence that language can foster higher-order relational structure comes from research by Loewenstein and Gentner (1998, 2002). We tested the effects of spatial language on spatial mapping ability, using the spatial prepositions *on*, *in*, and *under*—three particularly early spatial terms (Bowerman 1989; Clark 1974; Johnston 1988)—as well as the locatives *top*, *middle*, and *bottom*. As in the Rattermann and Gentner studies, we first established a difficult spatial analogy task and then tested whether labeling the relevant relations would improve performance. The child was shown two identical tall boxes, a hiding box and a finding box. Each box had a shelf in the middle so that it had three salient placement locations, as shown in figure 8.2: on top, in the middle, and under the box. Each box had three identical plastic cards, one in each

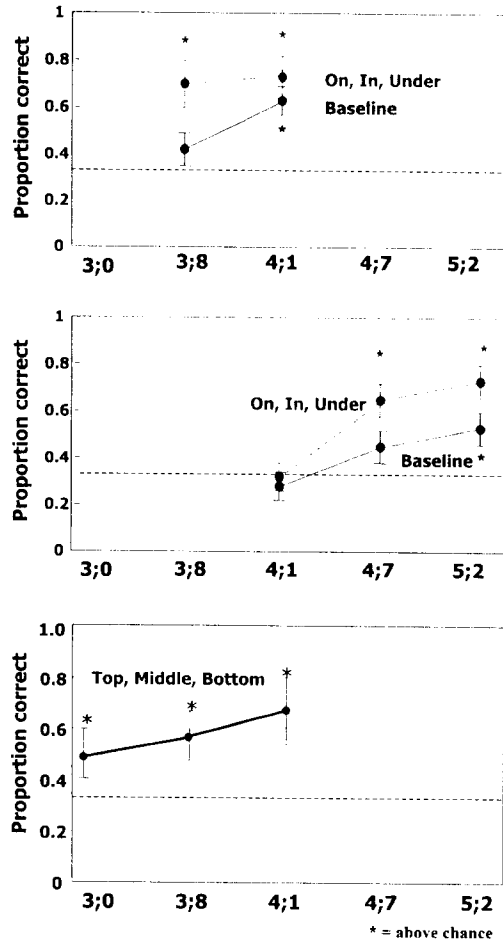




Figure 8.2 Results of the Loewenstein and Gentner spatial mapping studies, showing benefits of overt spatial language that diminish with age but reappear for more difficult tasks

position. One card had a star on its back, making it the “winner.” Preschool children were shown the location of the winner at the hiding box and had to find the winner in the corresponding location at the finding box.

In some respects, the task is a relatively easy version of the search task used in DeLoache’s (1987, 1995) and in our own (Loewenstein and Gentner 2001) model-room studies. The hiding and finding models are nearly identical and are placed close together so that they can be viewed simultaneously. However, the box task is considerably more difficult than the standard model-room task in one key respect: it cannot be solved by object correspondences. Because all the cards look alike, to solve the task the child must find corresponding spatial relations between the hiding box and the finding box.

For half the children, spatial relational language was used to describe the initial hiding event (e.g., “I’m putting this *on* the box”). For the other half (the control condition), the experimenter simply said as he placed the winner in its spot, “I’m putting it *here*.” In both cases, the child was asked, “Can you find the winner in the very same place in the finding box?” The experimenter put the winner at one of the three locations in the hiding box as the child watched, and the child searched for the corresponding winner in the finding box.

Loewenstein and I noted five predictions that should follow if spatial relational language leads to forming articulated spatial representations that support the relational mapping process: (1) young children should perform better when overt spatial relational language is used; (2) older children, who have internalized the relational system, will not need overt language; (3) if the task is made more difficult, older children will again show benefits from language; (4) the benefits of language should be predictable from the semantics of the terms (as opposed to there being some general attentional effect of labeling); and (5) the benefits should be retained over time. These predictions were borne out. At age 3;6, children who had heard the box locations described in terms of the spatial relations *on*, *in*, and *under* performed substantially better on the mapping task than control children, who performed at levels just better than chance. By age 4;0, children no longer needed to hear the relational language to succeed at the mapping task (figure 8.2, top panel). However, if

cross-mapped objects were used, placing object similarities in competition with the current relational correspondences, then 4-year-olds performed at chance in both conditions. With this more difficult task, in keeping with prediction (3), still older children (ages 4;7 and 5;2) showed significant benefits of relational language (figure 8.2, middle panel).

To test the claim of semantic specificity, we compared the terms *top/middle/bottom* (which form a connected relational system) with the terms *on/in/under* (which each express a separate relation between a figure and a ground). If children's representations reflect the semantics of the terms, then they should be better able to maintain a relational mapping with the deeper relational system conveyed by *top/middle/bottom*. Indeed, this was the outcome: even 3-year-olds were able to carry out the relational mapping when the connected system of *top/middle/bottom* was used. Thus, hearing relational language facilitated children's ability to encode and map on the basis of spatial relations. The benefits of language-guided encoding were maintained when children were brought back to the laboratory a few days later and asked to "play the game again" (with no mention of the spatial terms). It appears that the language experience led to a genuinely different encoding (and not to some momentary attentional benefit). This result is evidence that overt use of relational language can invite children to represent and use higher-order relational structure.

8.5 Symbol Use in Other Primates

Studies of the role of language in human thought are hampered by the fact that there is no comparison group of otherwise normal humans who lack a language. However, there is an indirect approach. We can compare nonhuman primates who have been taught symbol systems to otherwise matched animals who have not (Gentner and Rattermann 1991; Kuczaj and Hendry, this volume). There are several ways in which language appears to make a difference. I focus on two arenas: numerical competency and relational matching.

8.5.1 Number

Boysen and her colleagues have carried out an intriguing set of studies of quantity judgments among chimpanzees (Boysen and Berntson 1995;

Boysen et al. 1996). In their studies, a chimpanzee is shown two arrays of candy differing in quantity (e.g., one vs. three candies). The animal points to one of the arrays and then is given the other. Clearly, the best strategy is to point to the smaller number of candies, thereby garnering the larger set. This strategy turned out to be extremely difficult for the chimpanzees, even though all the animals tested had been given cognitive training with number symbols.⁵ Even after many trials, they continued to fail the task, repeatedly pointing to the larger amount and receiving the smaller amount. Not surprisingly, they readily succeeded when the task was simply to point to the array they wanted; but they could not master the reverse strategy of pointing to the array they did not want.

However, the situation changed when the same chimpanzees were tested with numerical symbols. In this case, they readily selected the smaller Arabic numeral, thus garnering the (larger) quantity of candies represented by the unselected numeral. In subsequent trials, they consistently mastered the correct strategy with numerals and failed it with actual arrays. Why do the animals perform so much better with numerals? The numbers do not add new quantity information—indeed, the animals are responding all too strongly to the quantity difference in the concrete situation. It appears that the advantage of abstract symbols is that they allow the chimpanzees to process the quantities at a level of abstraction removed from the rich sensory power of the actual food.

This pattern is reminiscent of studies of human development. In the Rattermann and Gentner mapping task, children are better able to resist a tempting (incorrect) object match when the objects are perceptually sparse than when the objects are richly detailed and thus far more compelling as similarity matches. Likewise, DeLoache has found that preschoolers do better in a model-room mapping task when given photographs rather than three-dimensional models. In the case of the chimpanzees, numerals served as the ultimate "abstract objects." Using numerals allowed them to select and compare only the property of magnitude, leaving behind the sensory qualities that were their undoing in the concrete choice task.

8.5.2 Relational Labeling and Relational Matching

Across species, relational matching is an uncommon ability. While many animals can succeed in learning a match-to-sample task with objects such

as that shown in (6), the ability to succeed at a relational matching task like the one shown in (7)—that is, at analogical matching—is much rarer (Premack 1983):

(6) A

A B

(7) AA

BB CD

Oden, Thompson, and Premack (2001) have carried out a fascinating set of studies that suggests that symbol training is crucial to relational matching (see Kuczaj and Hendry, this volume). When chimpanzees were taught to choose a particular symbol for two identical objects and another symbol for two nonidentical objects—that is, symbols for *same* and *different*—they readily generalized these symbols to relations between objects (Thompson, Oden, and Boysen 1997). For example, having learned to choose *same* for A/A and *different* for A/B, they can then solve a relational match-to-sample task. That is, when given triad (7), they can choose BB if asked to choose the *same* one, and CD if asked to choose the *different* one. To do this, the chimpanzee must apply *same/different* at the relational level as well as the object level. It is as though the chimp succeeds only when she can construct representations with relational predicates, yielding the triad in (8):

(8) *same*(A,A)

same(B,B) *different*(C,D)

Symbol training appears to be necessary for success on the relational matching task. However, it is not sufficient. There are species differences in the ability to learn relational similarity, even when symbols are given. Macaque monkeys given the same training with *same* and *different* symbols as the chimpanzees were eventually able to master object matching, but not relational matching (Washburn, Thompson, and Oden 1997). Interestingly, infant chimpanzees—but not infant macaques—show a kind of implicit relational matching. After handling a series of pairs of identical objects, they show more interest in a nonidentical

pair, and vice versa (Oden, Thompson, and Premack 2001). Thus, infant chimpanzees show an implicit capacity that can become an explicit cognitive ability with the support of symbols.

Premack (1983; see also Oden, Thompson, and Premack 2001) interprets such findings in terms of two codes: an imaginal code closely tied to the perceptual properties of objects and a propositional code. He suggests that only animals who have learned a symbolic communication system use a propositional code. Chimpanzees are born with the capacity for implicit relational matching, but whether they ever realize their full potential for analogical thinking depends on whether they learn relational language.

8.6 Summary and Discussion

To the question I began with—“What makes humans so smart relative to other species?”—I have given two answers: (1) analogical ability, (2) language. First, humans are endowed with a greater degree of analogical ability than other species. Although we are not the only animal with analogical ability, the difference in degree of ability is so great that it stands as a qualitative difference. We are roughly similar to other intelligent species in our ability to form associations and to engage in statistical learning. Indeed, in many arenas, such as navigation and spatial memory, our powers are inferior to those of other animals. Structure-mapping processes are where we most differ from other species in our cognitive powers.

The second contributor to our intelligence is language and other cultural systems, which multiply our cognitive resources. Language augments our cognition in a number of ways. Externally, it allows each new generation to learn from and build further on the knowledge of past generations. Internally, as argued in this chapter, language provides cognitive tools. It augments the ability to hold and manipulate concepts and sets of concepts—in particular, systems of relations. Thus, although structure-mapping may be a species-wide innate ability, its deployment is influenced by language and culture. The results reviewed here suggest that the acquisition of relational language is instrumental in the development of analogy. It follows, then, that the acquisition of relational language contributes importantly to the development of cognition.

8.6.1 Structural Alignment and the Career of Similarity

Are structure-mapping processes innate? Evidence provided by Gomez and Gerken (1999) and Marcus et al. (1999) suggests that the answer is yes: the ability to notice and abstract relational regularities across exemplars is in place even in 7- and 8-month-olds. In Marcus et al.'s studies, infants received 16 three-syllable strings with the same pattern—either ABA (e.g., *pa-ti-pa*, *go-di-go*) or ABB. After three repetitions of these 16 strings, the infants were tested on strings consisting of new syllables—half in the trained pattern, half in the nontrained pattern. Infants dishabituated significantly more often to sentences in the nontrained pattern than to sentences in the trained pattern, indicating that they had abstracted the common structure from the training set.

My colleagues and I have successfully modeled Marcus et al.'s results using a system (SEQL) that compares examples (via SME) and sequentially abstracts their common structure (Gentner, Kuehne, and Forbus, in preparation; Kuehne, Gentner, and Forbus 2000; Skorstad, Gentner, and Medin 1988). Unlike most simulations of these findings, it requires only the set of 48 sentences given to infants (whereas some connectionist simulations require thousands of trials). These findings are consistent with the possibility that structure-mapping processes are responsible for the infants' grammar-learning process.

At this point, a challenge might reasonably be raised: if structure-mapping processes are present at birth, then why is the normal course of development so slow? Or to put it another way, how can the same process explain both the results of the infant grammar studies, in which babies show rapid structural abstracting, and the lengthy process of normal children's grammar learning? The resolution lies in when and whether comparisons are made. Structural alignment processes are extremely powerful at aligning and revealing common relations *when they are brought to bear*. Even adults miss many potential comparisons. As noted earlier, in memory experiments, adult participants routinely fail to retrieve past exemplars that are analogous to current exemplars. We are often not reminded of prior experiences that are potentially analogous to current experiences, and this is particularly true for novices, whose representations are more idiosyncratic and less likely to match each other than those of experts. Children's early representations are highly

idiosyncratic and context specific. Thus, in order to notice a match, they require either very high overall similarity or very close temporal juxtaposition—ideally, both. In the infant grammar studies, the babies receive the latter—repeated close comparisons that allow progressive alignment of the common structure.

In habituation experiments like Marcus et al.'s studies, babies receive an optimal learning experience, from the vantage of structure-mapping theory.⁶ In the ordinary course of learning, the application of structure-mapping processes is largely constrained by the luck of environmental juxtapositions. In habituation experiments, luck is in the hands of a benevolent experimenter, who can guarantee optimal juxtapositions.

8.6.2 Are Symbols Necessary?

Several recent schools of thought—including dynamic systems theory, situated cognition, and distributed connectionism—have generated an interest in subsymbolic or nonsymbolic representations.⁷ In the extreme, some theorists have argued that symbolic representation, or structured representation, or even representation in general, has no role in human cognition. The evidence presented here suggests that human cognition arises not only from the world as directly perceived, but also from learned symbol systems that facilitate the apprehension of relational structure.

I reviewed two lines of evidence for the claim that learned symbol systems contribute to cognitive ability. The first line examines the effects of acquiring language. In our studies, children's performance on mapping tasks benefited from hearing the terms *top*, *middle*, *bottom* (Loewenstein and Gentner 1998, 2002) or the terms *daddy*, *mommy*, *baby* (Rattermann and Gentner 1998, 2002). In our studies, the overt use of relational language aids children's performance on analogical mapping tasks across a wide range of age and task difficulty. These findings are most naturally explained by assuming (1) that symbolic relations are used in carrying out analogical mapping tasks, and (2) that the acquisition of relational language plays a role in the development of symbolic representations.

The above line of argument has the disadvantage that all normal children can eventually perform the tasks in our studies without the overt

use of relational language. This fact is not necessarily inconsistent with my claims—for example, it could be that older children have internalized relational symbols learned from language—but it raises the possibility that the effects of language in our task are transient and perhaps epiphenomenal. However, there is a second line of evidence for the importance of symbols that cannot be explained away in this manner—namely, studies comparing other great apes who either do or do not possess symbol systems. Boysen's chimps can master the task of pointing to the nondesired pile of candy if and only if they have a symbolic code for numbers that lets them rise above the concrete situation. It is only when the perceptual-motor affordances of real foods are replaced by abstract symbols that the animals can reason clearly enough to choose the best strategy. Direct perception in this case is working against them; it is abstraction that allows them to succeed.

To ask whether similar benefits accrue in human learning, one avenue of inquiry is the acquisition of technical language. Because any given technical vocabulary is learned by some but not all humans, we can compare “haves” with “have nots” as in the chimpanzee studies. There is some evidence that the acquisition of technical language can confer new cognitive possibilities. Koedinger, Alibali, and Nathan (in preparation) argue that the acquisition of algebraic notation allows children to move from concretely grounded representations of word problems to symbolic representations; and further, that although grounded representations are more effective for simple problems, symbolic representations are better for complex problems.

All this suggests that although situated or embodied cognition may be a natural mode of human processing, there are many cases where what is needed is the opposite: representations that are de-situated or disembodied. Symbolic representations lose some of the richness of embodied cognition, but they open possibilities that cannot be imagined without them. One function of language may be to augment natural modes of cognition with an alternative representational scheme that permits abstract cognition.

8.6.3 Language and Thought

It is useful to contrast the view taken here with other views on language and thought. The strong version of the Sapir-Whorf hypothesis holds

that (1) languages vary in their semantic partitioning of the world; (2) the structure of one's language influences the manner in which one perceives and understands the world; (3) therefore, speakers of different languages will perceive the world differently. Past efforts to demonstrate the strong version of the Whorfian position have produced mostly negative results (Pinker 1994; however, see Hunt and Agnoli 1991; Kay and Kempton 1984; Lucy 1994; Lucy and Shweder 1979). Current research continues to find mixed results, as demonstrated by the chapters in this volume.

Vygotsky's (1962) theory also gives language a major role in cognition. However, his theory focuses chiefly on the general effects of learning a language, rather than on the specific conceptual construals invited by a given language. According to Vygotsky, with the advent of language children augment their prelinguistic cognitive abilities—reactive attention, associative learning, and sensorimotor intelligence—with new capacities for focused attention, deliberate memory, and symbolic thought (see also Dennett 1993). On this view, acquiring a language gives the child control over his own mental processes: the ability to direct attention, to choose a course of thought, and to formulate mental plans.

Thus, the Sapir-Whorf view has it that the grammatical structure of a language shapes its speakers' perception of the world, and the Vygotskian view emphasizes that possessing an internal language permits speakers to guide their own mental processes. I am suggesting a third, hybrid position: that learning specific relational terms and systems provides representational resources that augment our cognitive powers. On this account, language is neither a lens through which one forever sees the world, nor a control tower for guiding cognition, but a set of tools with which to construct and manipulate representations.

Whereas tests of the Whorfian hypothesis have generally involved between-language comparisons, the cognitive tools view can be tested within a language. We can compare outcomes when different subsets of symbolic terms are provided to different groups (as in our studies) or are acquired by different populations (as in the case of technical vocabularies). Of course, the cognitive resources view I espouse also suggests possible crosslinguistic differences. Languages that have different lexicalizations of relational information offer their speakers different options for representation and reasoning. Indeed, relational terms are the most

likely arena in which to find linguistic influences on thought, for two reasons. First, relational terms are more variable crosslinguistically than object reference terms (as discussed earlier). Obviously, semantic differences are necessary (though not sufficient) for there to be resulting cognitive differences. Second, relational terms—including spatial relational terms and verbs—provide framing structures for the encoding of situations and events. Hence, semantic differences in these categories could reasonably be expected to have cognitive consequences.

But despite the obvious importance of crosslinguistic studies, I have argued here that there are important issues that apply within a single language. Relational labels invite the child (or adult) to notice, represent, and retain structural patterns of elements, and therefore to transfer relational patterns and to reason fluently over combinations of relations. Even within a single language, the acquisition of relational terms provides both an invitation and a means for the learner to modify her thought.

8.6.4 Challenges and Limitations

First, a few clarifications are in order. I am not suggesting that all culturally learned concepts are relational; concepts like “fruit” and “shard” are counterexamples (for different reasons). I am also not claiming that all abstract concepts are relational. Counterexamples include concepts like “idea” and “entity.” (However, I’d guess that a large percentage of abstract terms are relational.) There are also many abstract concepts whose representations include both relational information and intrinsic information—for example, “mammal” and “reptile,” which are abstract. Another important clarification is that although I have focused on language, there are other acquired systems that make us smart—among them, numbers (Carey 1998; Spelke, this volume), maps (Uttal 2000), and other artifacts (Norman 1993).

Turning to a deeper issue, the proposal that learning language can invite new conceptual representations runs immediately into a classic objection. Fodor’s circularity challenge is that “... one cannot learn a language unless one has a language. In particular, one cannot learn a first language unless one already has a system capable of representing the predicates in that language *and their extensions*” (1975, 64; italics

original). Thus, we can’t learn a word’s meaning unless we already have the representational resources necessary to understand the concepts to which it refers. The learning hypothesis might still be saved by making strong assumptions about the innate set of representational resources—for example, by assuming that we begin with a set of primitives out of which semantic representations are built. However, although the empirical evidence concerning semantic primitives can be debated, this is clearly a troublesome move, particularly in the absence of a viable candidate set of primitives. Thus, Fodor concludes that learning cannot give us new concepts. He therefore proposes that humans are born with an innate language of thought, in terms of which they learn the overt language of their community.

This is the kind of argument that makes psychologists want to say, “Oh, go away—can’t you see we’re busy doing experiments?” But the question of what we start with is important. The challenge, then, is that *learners need a prior conceptual understanding of what a word means in order to attach a word to that meaning*. I do not have a complete answer. But I believe one part of the answer lies in the distinction between implicit and explicit understanding. Learning words provides explicit internal labels for ideas that were previously merely implicit, and this gain in explicitness has cognitive consequences. Likewise, carrying out an analogy lays bare common structure that was previously invisible, embedded in the richness of particular exemplars. In Boysen’s studies of chimpanzees, infant chimpanzees show implicit sensitivity to identity relations between objects. But they cannot cash in this sensitivity, even as adults, without language training. Only if they are given symbols for *same* and *different* can they reliably detect sameness and difference over relations. I suggest (1) that the relational symbols invite explicit representation of the relations, and (2) that this explicitness makes the *same/different* relations more portable—it allows the chimpanzees to go beyond object matches to a new level of application between relations.

Extrapolating to humans, this suggests that one result of language learning may be to change the internal language from a restricted implicit system to a more powerful explicit system. A reasonable question at this point is whether there are other relations, besides *same/different*, that might be implicitly present in humans prior to language learning.

Crosslinguistic patterns suggest that some relational terms are particularly easy to acquire (Choi and Bowerman 1991); although more remains to be done, this work may provide clues as to which relational concepts are implicitly formed prelinguistically.

Once some relational concepts are extracted, learning more words can occur by conceptual combination. For example, *forget* can be learned as not-remember, or *trade* as a reciprocal giving relation: *x* gives something to *y* and *y* gives something to *x*. Another way of deriving new meanings from old is by analogy. For example, suppose a child is told that fish breathe water with their *gills* (a new word). She is invited to map a causal chain from (8) to (9) (humans to fish):

- (8) PERMIT(EXTRACT(lungs(people), oxygen, air),
BREATHE(people, air))
- (9) PERMIT(EXTRACT(gills(fish), oxygen, water), BREATHE(fish,
water))

Of course, this is only the beginning. At this point, the child knows only the functional role of gills, not what they look like or how they work—but she has delineated the concept of gills and perhaps become curious to know more. That kind of focused curiosity is part of what makes language a potent force in learning.

Finally, the relational concepts provided by language and other cultural systems are a key starting set. But speakers are not limited to the set of existing lexicalized relations. As Bowerman (1981) and Clark (1993) have observed, children regularly invent new relational terms. Indeed, the propensity to invent symbols is a striking difference between humans and other apes. Further, new relational concepts can arise in a language through mechanisms such as metaphorical abstraction, by which concrete terms are extended into abstract meanings (see table 8.1) (Gentner and Bowdle 2001; Glucksberg and Keysar 1990; Kittay and Lehrer 1981; Wolff and Gentner 2000). Speakers constantly go beyond the current resources of their language to develop new relational abstractions. Extensions into progressively higher-order relational terms have characterized the history of science and mathematics. However, I suggest that systems of currently lexicalized relations frame the set of new ideas that can be readily noticed and articulated.

8.6.5 Conclusions

General learning mechanisms have come under heavy fire in the last few decades. Children's learning is seen as far too rapid to be accounted for by a general learning process. Further, children seem primed for learning in certain domains, such as mechanical causation, biology, and theory of mind. This has suggested to many researchers that humans possess special faculties for learning in privileged domains. By analogy with grammar, these other privileged domains are assumed to have built-in representations and processes that facilitate acquisition.

In the above account, the human advantage is a cognitive head start over other species. I suggest another perspective. The great evolutionary advantage of the human species is adaptability. We are at home in the tropics or in the Arctic. To design a superbly adaptable species, one might best create one that begins with few biases beyond those necessary for mammalian life, that has a powerful general learning mechanism that abstracts significant commonalities and differences, and that has a species-wide method of capture—namely, language—with which to preserve important cognitive discoveries so that they can be combined generatively and passed to the young.

I am not suggesting that humans are born without constraints. We appear to come equipped with the basic mammalian starting set of attentional biases and learning propensities, as well as others that stem from being social animals. There also appear to be attentional biases evolved specifically for language, such as a readiness to learn the voiced-nonvoiced distinction (Saffran 2001). But in contrast to theories that postulate that humans have more built-in knowledge and theory than other species, I suggest the reverse: if anything, we have less. Whereas the frog comes programmed to jump for looming shade and to flick its tongue for small moving objects, we come prepared to learn what is dangerous and what is edible. Far from being a disadvantage, our relatively unbiased initial state allows us to learn whatever comes our way.

This "less is more" proposal for the human endowment is not new, of course. It has a long history in evolutionary anthropology. But for the most part, general learning as an explanation for cognitive development has been out of favor in the last two decades of cognitive theorizing. In part, this resulted from the limitations of purely behaviorist approaches

to learning. But we now know of learning mechanisms that go beyond mere association and perceptual generalization. Structure-sensitive comparison processes, which occur even in infancy, can invite alignment and progressive abstraction of relational structures.

Finally, learning language is crucial to the development of cognition. Learned relational symbols provide representational tools with which to structure knowledge. These learned relational tools amplify the human capacity for structural alignment and mapping. For example, if a pattern discovered by analogy is named, it becomes easier to see as part of yet another analogy. This process of extracting relations via analogy and then preserving them via language acts to bootstrap learning and to create the structured symbolic representations essential for higher-order cognition.

Notes

The research on spatial relational terms reported here was supported by NSF-ROLE grant 21002/REC-0087516. The research on analogical learning and the development of the computer simulation of analogical processing were supported by ONR grant N00014-92-J-1098. I thank Susan Goldin-Meadow and Arthur Markman for comments on this chapter; Jeff Loewenstein, Mary Jo Rattermann, Ken Kurtz, and Arthur Markman for valuable discussions of these issues; and Kathleen Braun and Michelle Osmondson for help with the research.

1. It is surely not a coincidence that the species that show the most impressive cognitive abilities are social animals: apes, dogs, dolphins, crows, parrots.
2. Another factor is that there are typically multiple relational structures within any one representation. In literal similarity, most of these relational structures can be placed in correspondence, again strengthening the match. In analogy, typically only one or perhaps two relational structures match, so the maximal match must typically be discerned from among many local relational matches, most of which must eventually be discarded.
3. Younger children often make the object match instead of the relational match, presumably because they lack sufficiently elaborated relational representations to yield a relational alignment deep enough to prevail against the object match (Gentner and Rattermann 1991; Gentner and Toupin 1986).
4. *Relational terms* are terms that convey a relation: that is, a proposition taking at least two arguments. First-order relations take entities as their arguments. Nth order relations have at least one (N - 1)th relation as arguments.
5. Space does not allow a full description of chimpanzees' number achievements. However, Boysen and her colleagues have trained one animal, Sheba, to

point to a number on a screen to express the cardinality of a set of objects (up to at least five). In number tasks, Sheba often partitions objects and touches them sequentially, as do children learning to count. Matsuzawa (1991) taught a female chimpanzee (Ai) to name the number of items from one to six. Ai was able to transfer this skill to new objects.

6. Consideration of comparison processes also points up an important issue in the interpretation of habituation results. It is fair to conclude that the generalizations infants arrive at in habituation experiments are within their power to learn, but not that the knowledge was present before the experience of habituation. Thus, conclusions of the form "The babies understand that ..." should in many cases be replaced by "The babies can readily learn that ..."

7. See Markman 1999 and Markman and Dietrich 1999, 2000, for extended discussions of this point.

References

- Bowerman, M. (1981). The child's expression of meaning: Expanding relationships among lexicon, syntax and morphology. In H. Winitz (Ed.), *Native language and foreign language acquisition* (Vol. 379, pp. 172-189). New York: New York Academy of Sciences.
- Bowerman, M. (1989). Learning a semantic system: What role do cognitive predispositions play? In M. L. Rice and R. L. Schiefelbusch (Eds.), *The teachability of language* (pp. 133-168). Baltimore: Brookes.
- Bowerman, M. (1996). Learning how to structure space for language: A cross-linguistic perspective. In P. Bloom, M. A. Peterson, L. Nadel, and M. F. Garrett (Eds.), *Language and space* (pp. 385-436). Cambridge, MA: MIT Press.
- Boysen, S. T., and Berntson, G. G. (1995). Responses to quantity: Perceptual versus cognitive mechanisms in chimpanzees (*Pan troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes*, 21, 82-86.
- Boysen, S. T., Berntson, G. G., Hannan, M. B., and Cacioppo, J. T. (1996). Quantity-based interference and symbolic representations in chimpanzees (*Pan troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes*, 22, 76-86.
- Byrnes, J. P. (1991). Acquisition and development of *if* and *because*: Conceptual and linguistic aspects. In S. A. Gelman and J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 3-27). Cambridge: Cambridge University Press.
- Carey, S. (1998). Knowledge of number: Its evolution and ontogenesis. *Science*, 242, 641-642.
- Chen, Z., Sanchez, R. P., and Campbell, T. (1997). From beyond to within their grasp: The rudiments of analogical problem solving in 10- and 13-month-olds. *Developmental Psychology*, 33, 790-801.

- Choi, S., and Bowerman, M. (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, 41, 83-121.
- Clark, E. V. (1974). Non-linguistic strategies and the acquisition of word meanings. *Cognition*, 2, 161-182.
- Clark, E. V. (1993). *The lexicon in acquisition*. Cambridge: Cambridge University Press.
- Clement, C. A., Mawby, R., and Giles, D. E. (1994). The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language*, 33, 396-420.
- DeLoache, J. S. (1987). Rapid change in the symbolic functioning of very young children. *Science*, 238, 1556-1557.
- DeLoache, J. S. (1995). Early understanding and use of symbols: The model model. *Current Directions in Psychological Science*, 4(4), 109-113.
- Dennett, D. C. (1993). Learning and labeling. *Mind and Language*, 8, 540-548.
- Falkenhainer, B., Forbus, K. D., and Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Fodor, J. A. (1975). *The language of thought*. New York: Crowell.
- Forbus, K. D., and Gentner, D. (1989). Structural evaluation of analogies: What counts? *Proceedings of the Cognitive Science Society*, 11, 341-348.
- Forbus, K. D., Gentner, D., and Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141-205.
- Gentner, D. (1978). On relational meaning: The acquisition of verb meaning. *Child Development*, 49, 988-998.
- Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, 4, 161-178.
- Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj (Ed.), *Language development: Vol. 2. Language, thought, and culture* (pp. 301-334). Hillsdale, NJ: Erlbaum.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). Cambridge: Cambridge University Press.
- Gentner, D., and Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman and S. C. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 215-256). Cambridge: Cambridge University Press.
- Gentner, D., and Bowdle, B. F. (2001). Convention, form, and figurative language processing. *Metaphor and Symbol*, 16, 223-247.
- Gentner, D., and Klibanoff, R. S. (in preparation). *On acquiring gifts: The development of the ability to learn relational nouns*.

- Gentner, D., Kuehne, S., and Forbus, K. (in preparation). *A model of infant grammar learning using structural alignment*.
- Gentner, D., and Kurtz, K. J. (in preparation). *The mechanisms of mapping: Evidence from on-line analogy judgments*.
- Gentner, D., and Loewenstein, J. (2002). Relational language and relational thought. In J. Byrnes and E. Amsel (Eds.), *Language, literacy, and cognitive development* (pp. 87-120). Mahwah, NJ: Erlbaum.
- Gentner, D., and Markman, A. B. (1997). Structure-mapping in analogy and similarity. *American Psychologist*, 52, 45-56.
- Gentner, D., and Medina, J. (1998). Similarity and the development of rules. *Cognition*, 65, 263-297.
- Gentner, D., and Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, 14, 487-513.
- Gentner, D., and Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman and J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 225-277). Cambridge: Cambridge University Press.
- Gentner, D., Rattermann, M. J., and Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25, 524-575.
- Gentner, D., and Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277-300.
- Gentner, D., and Wolff, P. (1997). Alignment in the processing of metaphor. *Journal of Memory and Language*, 37, 331-355.
- Gick, M. L., and Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M. L., and Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Glucksberg, S., and Danks, J. H. (1968). Effects of discriminative labels and of nonsense labels upon the availability of novel function. *Journal of Verbal Learning and Verbal Behavior*, 7, 72-76.
- Glucksberg, S., and Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, 97, 3-18.
- Glucksberg, S., and Weisberg, R. W. (1966). Verbal behavior and problem solving: Some effects of labeling in a functional fixedness problem. *Journal of Experimental Psychology*, 71, 659-664.
- Goldin-Meadow, S., Seligman, M. E., and Gelman, R. (1976). Language in the two-year-old. *Cognition*, 4, 189-202.
- Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 3-28.

- Goldstone, R. L., and Medin, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 29–50.
- Gomez, R. L., and Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70, 109–135.
- Hall, D. G., and Waxman, S. R. (1993). Assumptions about word meaning: Individuation and basic-level kinds. *Child Development*, 64, 1550–1570.
- Herbert, J., and Hayne, H. (2000). The ontogeny of long-term retention during the second year of life. *Developmental Science*, 3, 50–56.
- Hunt, E., and Agnoli, F. (1991). The Whorfian hypothesis: A cognitive psychology perspective. *Psychological Review*, 98, 377–389.
- Johnston, J. R. (1988). Children's verbal representation of spatial location. In J. Stiles-Davis and M. Kritchevsky (Eds.), *Spatial cognition: Brain bases and development* (pp. 195–205). Hillsdale, NJ: Erlbaum.
- Kay, P., and Kempton, W. (1984). What is the Sapir-Whorf hypothesis? *American Anthropologist*, 86, 65–79.
- Keane, M. T. (1988). *Analogical problem solving*. New York: Halsted Press.
- Keil, F. C., and Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, 23, 221–236.
- Kemler, D. G. (1983). Holistic and analytical modes in perceptual and cognitive development. In T. J. Tighe and B. E. Shepp (Eds.), *Perception, cognition and development: Interactional analysis* (pp. 77–102). Hillsdale, NJ: Erlbaum.
- Kittay, E. F., and Lehrer, A. (1981). Semantic fields and the structure of metaphor. *Studies in Language*, 5, 31–63.
- Koedinger, K. R., Alibali, M. W., and Nathan, M. J. (in preparation). *Trade-offs between grounded and abstract representations: Evidence from algebra problem solving*.
- Kotovskiy, L., and Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797–2822.
- Kubose, T. T., Holyoak, K. J., and Hummel, J. E. (in preparation). *Strategic use of working memory in analogical mapping*.
- Kuehne, S. E., Gentner, D., and Forbus, K. D. (2000). Modeling infant learning via symbolic structural alignment. *Proceedings of the Cognitive Science Society*, 22, 286–291.
- Kurtz, K. J., and Gentner, D. (1998, November). *The mechanisms of mapping: Evidence from on-line judgments of analogy*. Poster presented at the meeting of the Psychonomic Society, Dallas, TX.
- Landau, B., Smith, L., and Jones, S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299–321.
- Loewenstein, J., and Gentner, D. (1998). Relational language facilitates analogy in children. *Proceedings of the Cognitive Science Society*, 20, 615–620.

- Loewenstein, J., and Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, 2, 189–219.
- Loewenstein, J., and Gentner, D. (2002). *Relational language fosters relational encoding and mapping*. Manuscript submitted for publication.
- Loewenstein, J., Thompson, L., and Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin and Review*, 6, 586–597.
- Lucy, J. A. (1994). *Grammatical categories and cognition*. Cambridge: Cambridge University Press.
- Lucy, J. A., and Shweder, R. A. (1979). Whorf and his critics: Linguistic and nonlinguistic influences on color memory. *American Anthropologist*, 81, 581–618.
- Marcus, G. F., Vijayan, S., Bandi, R. S., and Vishton, P. M. (1999). Rule-learning in seven-month-old infants. *Science*, 283, 77–80.
- Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Erlbaum.
- Markman, A. B., and Dietrich, E. (1999). Whither structured representations? *Behavioral and Brain Sciences*, 22, 626–627.
- Markman, A. B., and Dietrich, E. (2000). In defense of representation. *Cognitive Psychology*, 40, 138–171.
- Markman, A. B., and Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25, 431–467.
- Matsuzawa, T. (1991). Nesting cups and metatools in chimpanzees. *Behavioral and Brain Sciences*, 14, 570–571.
- Medin, D. L., Goldstone, R. L., and Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100, 254–278.
- Nelson, K. (1996). *Language in cognitive development: The emergence of the mediated mind*. Cambridge: Cambridge University Press.
- Norman, D. A. (1993). *Things that make us smart*. Reading, MA: Addison-Wesley.
- Novick, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 510–520.
- Oakes, L. M., and Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, 5, 193–207.
- Oden, D. L., Thompson, K. R., and Premack, D. (2001). Can an ape reason analogically? Comprehension and production of analogical problems by Sarah, a chimpanzee (*Pan troglodytes*). In D. Gentner, K. J. Holyoak, and B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 471–498). Cambridge, MA: MIT Press.

- Olguin, R., and Tomasello, M. (1993). Twenty-five-month-old children do not have a grammatical category of verb. *Cognitive Development*, 8, 245–272.
- Piaget, J. (1952). *The origins of intelligence in children*. New York: Norton.
- Pinker, S. (1994). *The language instinct*. New York: Morrow.
- Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences*, 6, 125–167.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Rattermann, M. J., and Gentner, D. (1998). 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, and B. Kokinov (Eds.), *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences* (pp. 274–282). Sophia: New Bulgarian University.
- Rattermann, M. J., and Gentner, D. (2002). *The effect of language on similarity: The use of relational labels improves young children's analogical mapping performance*. Manuscript submitted for publication.
- Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 456–468.
- Rovee-Collier, C. K., and Fagen, J. W. (1981). The retrieval of memory in early infancy. In L. P. Lipsett (Ed.), *Advances in infancy research* (Vol. 1, pp. 225–254). Norwood, NJ: Ablex.
- Saffran, J. R. (2001). The use of predictive dependencies in language learning. *Journal of Memory and Language*, 44, 493–515.
- Scholnick, E. K., and Wing, C. S. (1982). The pragmatics of subordinating conjunctions: A second look. *Journal of Child Language*, 9, 461–479.
- Skorstad, J., Gentner, D., and Medin, D. (1988). Abstraction processes during concept learning: A structural view. *Proceedings of the Cognitive Science Society*, 10, 419–425.
- Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 146–178). New York: 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.
- Thompson, R. K. R., Oden, D. L., and Boysen, S. T. (1997). Language-naive chimpanzees (*Pan troglodytes*) judge relations between relations in a conceptual matching-to-sample task. *Journal of Experimental Psychology: Animal Behavior Processes*, 23, 31–43.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352.
- Uttal, D. H. (2000). Seeing the big picture: Map use and the development of spatial cognition. *Developmental Science*, 3, 247–286.

- Vygotsky, L. (1962). *Thought and language*. Cambridge, MA: MIT Press. (Original work published 1934.)
- Washburn, D. A., Thompson, R. K. R., and Oden, D. L. (1997, November). *Monkeys trained with same/different symbols do not match relations*. Paper presented at the meeting of the Psychonomic Society, Philadelphia, PA.
- Wolff, P., and Gentner, D. (2000). Evidence for role-neutral initial processing of metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 529–541.