Developing Relations

Michael Gasser, Eliana Colunga and Linda B. Smith
Cognitive Science Program
Indiana University

Abstract

Relations lie at the center of humankind's most intellectual endeavors and are also fundamental to any account of linguistic semantics. Despite the importance of relations in understanding cognition and language, there is no well-accepted account of the origins of relations. What are relations made of? How are they made? In this chapter we address these questions. First, we consider past proposals of how relations are represented and the implications of these representational ideas for development. Second, we review the developmental evidence in the context of five psychological facts about relations that must be explained by any account of their origin. This evidence suggests that relational concepts are similarity based, influenced by specific developmental history, and influenced by language. Third, we summarize Gasser and Colunga's Playpen model of the learning of relations. This connectionist model instantiates a new proposal about the stuff out of which relations are made and the experiences that make them. Finally, we outline how the model explains the five psychological facts and consider the implications of this model for one of the questions addressed in this volume, the interface between conceptual structure and spatial representation.

Relations lie at the center of humankind's most intellectual endeavors — science, mathematics, poetry. Relations are also fundamental to any account of linguistic semantics. Indeed, a fundamental psychological distinction between objects and relations may be reflected in the universal linguistic distinction between nouns and verbs (Gentner, 1982; Langacker, 1987). Despite the importance of relations in understanding cognition and language and despite the dogged and continuing work of psychologists, linguists, and philosophers on the problem, there is no well-accepted account of the origins of relations. What are relations made of? How are they made? This chapter addresses this question by considering the psychological evidence on how children acquire relations and relational language and by proposing a computational model.

The plan of the chapter is as follows. First, we consider past proposals of how relations are represented and the implications of these representational ideas for development. Second, we review the developmental evidence in the context of five psychological facts about relations that must be explained by any account of their origin. This evidence suggests that relational concepts are similarity based, influenced by specific developmental history, and influenced by language. Third, we summarize Gasser and Colunga's Playpen model of the learning of relations. This connectionist model instantiates a new proposal about the stuff out of which relations are made and the experiences that make them. Finally, we outline how the model explains the five psychological facts and consider the implications of this model for the question addressed in this volume, the interface between conceptual structure and spatial representation.

Representing relations

Symbolic and connectionist theories seem at loggerheads throughout much of cognition, but in the domain of theories of relations, they are remarkably similar. Both classes of theories start with the same founding premise: Objects are prior to relations and atomic in the definition of relations. This makes sense: after all, ABOVE[1] means two objects in a particular relation, one to the other. To make sense of the idea of ABOVE one has to already have the idea of OBJECT. The starting problem, then, for all classes of relational theories has been how to represent the connection between related objects that specifies the relation. Two components have been taken as critical to these representations: (1) an element that characterizes the content and arity of the relation — that it is ABOVENESS for example and not BETWEENNESS that is being represented, and (2) a set of bindings that map the object arguments onto roles in the relation (Halford, Wilson, & Phillips, forthcoming). The binding of objects to roles is crucial in order to conceptually keep separate distinct situations, for example, BOOK ABOVE TABLE and TABLE ABOVE BOOK.

A brief consideration of the kinds of solutions offered to the binding problem suffices to make clear the uniformity of solutions to this representational problem. We begin with Figure 1, which offers a symbolic representation of ABOVE in which the relation term is represented by an explicit symbol, that is, the sequence of characters A, B, O, V, E. Binding is implemented by assigning particular positions in the representation to the roles of the relation and then by inserting representations of objects into these positions. This is the approach used in standard predicate-calculus notation: *Above (Book, Table).*

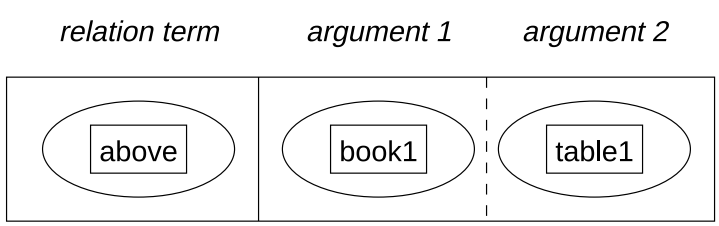

Figure 1. A relation represented using symbolic argument-style representation. The relation term and the arguments are symbols, the bindings are represented by the positions of the arguments.

Halford et al (1994) have proposed a connectionist version of the same kind of representation. We illustrate this in Figure 2: The relation term and the related objects are all activation vectors. They are fed into separate banks of units, places in the network, each

[1] In the chapter we will use SMALL CAPITALS to represent concepts and *italics* to represent expressions in natural language and other representational languages.

of which is dedicated to representing a particular component. The tensor product of these three vectors (for a binary relation) is computed to complete the binding process.

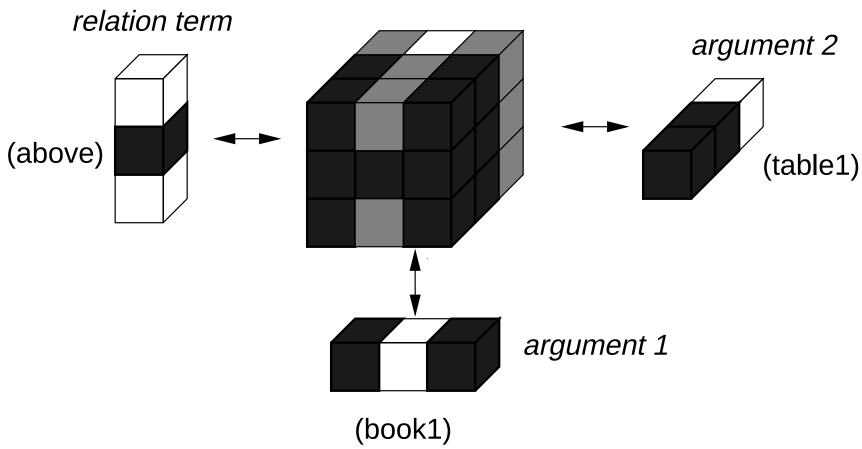

Figure 2. A relation represented using a connectionist argument-style representation. The arguments are fed to dedicated banks of units, and their bindings are represented using the tensor product.

Another solution to the binding problem involves pairing the objects with explicitly labeled role names (slots) rather than with places. A symbolic version of a slot-filler representation is illustrated in Figure 3. Here objects and roles are paired by concatenating the role-name symbol and the object symbol. One connectionist version of a slot-filler representation has been offered by Smolensky (1990). For each role-filler pair, a role-name vector and an object vector are fed into banks of role and filler units respectively and the tensor product of these vectors is calculated. Note that the relation term may be left out if it is completely specified by the role names; e.g., in place of ABOVE we have ABOVE-HIGHER and ABOVE-LOWER. This approach is illustrated in Figure 4.

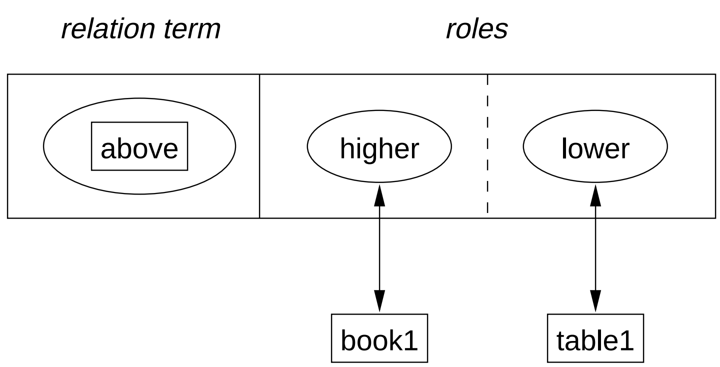

Figure 3. A relation represented using the symbolic explicit role representation. The binding is achieving by concatenating the role-name symbol and the filler object symbol.

In other connectionist approaches, separate role and filler units are somehow marked as belonging together rather than being placed on special purpose banks of units. In this approach, each unit in the network has an associated value (as well as an activation). When

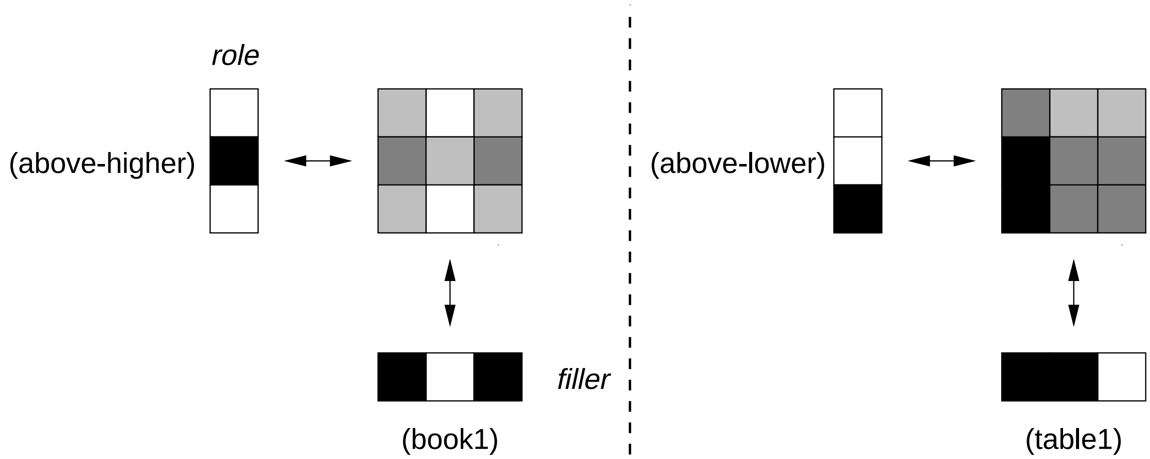

Figure 4. A relation represented using a distributed connectionist explicit role representation. The binding of a role and its filler is computed using the tensor product or convolution.

this value matches the value of another unit, they are bound together. In the dynamic binding approach (Hummel & Biederman, 1992; Hummel & Holyoak, 1997; Shastri & Ajjanagadde, 1993; Sporns, Gally, Reeke, & Edelman, 1989), units "fire" at particular times, and units whose firings are synchronized are considered bound. This localist approach is illustrated in Figure 5.

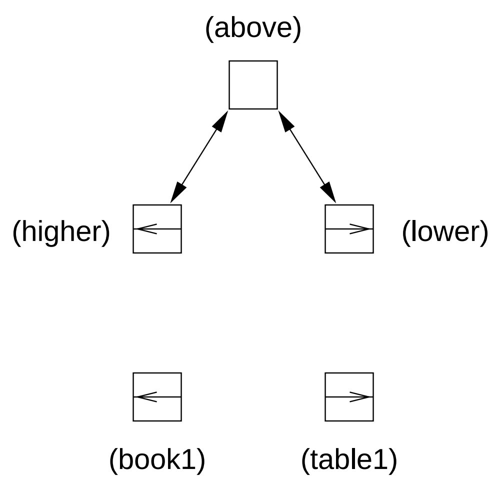

Figure 5. A relation represented using a localist connectionist explicit role representation. Binding is achieved through a value that is shared by the bound role and filler (arrows in the figure).

Even if not reducible one to the other, all of these ways of representing relations comprise a highly similar class. Table 1 summarizes the various approaches. All assume that the specification of how the objects in a relation are related is by explicitly labeling them as being in that relation. But where does this labeling come from? How do labeled representations interface with perception and actions on objects such that the experience of a particular book and particular table manages to engage the BOOK ABOVE TABLE representation? In all of the proposals about relational representation, the relations are just there, presumed a priori abstractions. This might be acceptable if there were a universal set of innate relations hardwired some way into biology. But the developmental and cross-language evidence on this point is clear: relations are learned. Their course of development is protracted and highly influenced by language learning.

Table 1: Approaches to the representation of relational knowledge.

		Relation term	Bindings
Symbolic	Predicate calculus	Symbol	Symbols in argument positions
	Slot-filler	Symbol	Role symbol + filler symbol
Connectionist	Argument style	Vector	Tensor product of relation and filler vectors
	Distributed, explicit role	(Implicit in bindings)	Sum of tensor product of role and filler vectors
	Localist	Unit	Role and filler units, synchronized

Five facts about relations

There are several extensive reviews of the development of relations (Bloom, Tinker, & Margulis, 1993; Gentner & Rattermann, 1991; Gentner, Rattermann, Markman, & Kotovsky, 1995; Smith, 1989). We do not attempt to duplicate these reviews here. Instead, we highlight five facts about development, facts about which there is remarkably little dispute in the usually contentious subfield of cognitive development. We take these as the facts that must be explained in an account of how and out of what relations are made.

Fact 1: Language matters.

Languages look very different from one another with respect to relations (Bowerman, 1996; Choi & Bowerman, 1992; Gentner, 1982). Even a cursory examination of the spatial relation expressions in a subset of languages reveals the variety of relational concepts possible. Consider some of the possibilities for encoding relations of CONTACT, SUPPORT, and CONTAINMENT between two objects (Landau, 1996). The roles (or slots) for these spatial concepts are the **trajector** — the thing being related, and the **landmark** — the thing to which the trajector is being related. Thus, the ball is the trajector and the cup is the landmark in THE BALL IS IN THE CUP.

Figure 6 presents four possible arrangements of a trajector (indicated by black) and a landmark (indicated by gray). Spanish uses a single word *en* for all of them. English uses one word, *on*, for the two situations in which containment is not involved and another, *in*, for situations in which the trajector is (at least partially) contained in the landmark. German distinguishes two kinds of situations for which English uses *on*: *auf* when the landmark is under the trajector and *an* when the trajector is fixed to a vertical surface of the landmark. Korean distinguishes two kinds of CONTAINMENT (and CONTACT) situations, those in which the trajector fits tightly with the landmark, for which *sok* is used, and those in which there is a loose fit, for which *ahn* is used. Clearly, languages "slice up" the relational space in different ways. This means at the least that the human capacity to represent and learn relations is a flexible one.

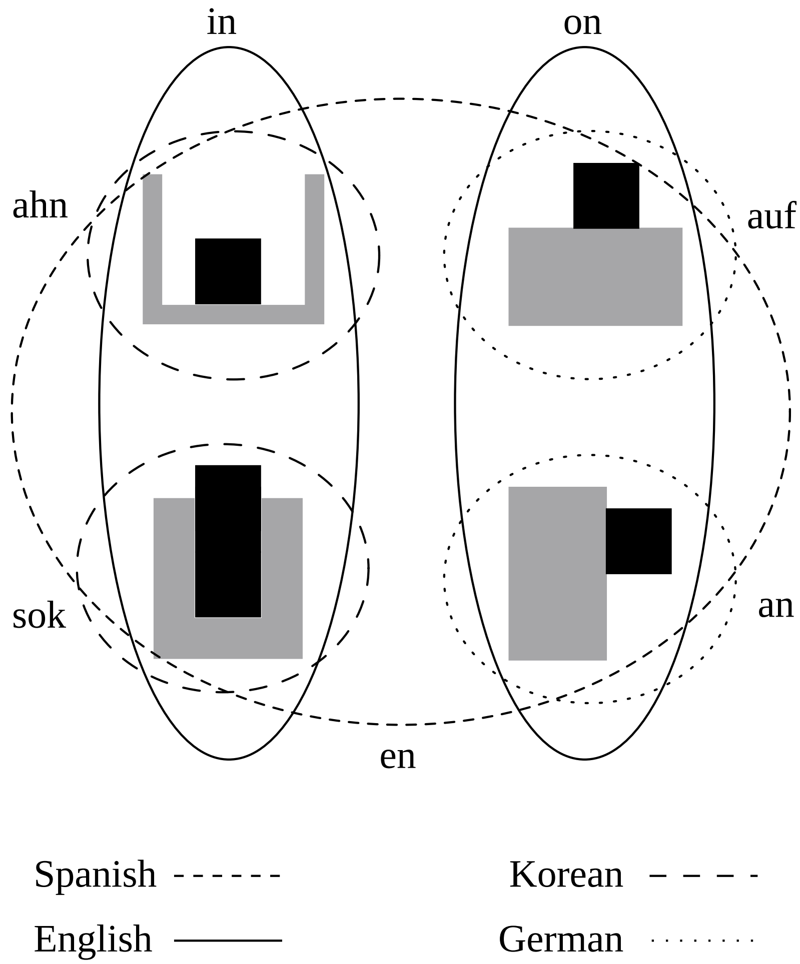

Figure 6. CONTAINMENT, SUPPORT, and CONTACT in four languages. Trajectors are indicated in black, landmarks in gray.

Of course, to be more accurate, languages do not "slice up" anything. Languages do not actually do anything at all; it is people that do. Each individual human being develops relational representations as a product of that individual's own activity — perceiving, acting, learning language. In learning the spatial categories that matter in their language, children could either be starting out with a universal set of fundamental categories (Jackendoff, 1996; Pinker, 1994) from which they select depending on their linguistic environment or learning the linguistic categories without the benefit of pre-existing knowledge. In either case, it is clear that the language being learned has much to do with the course of learning. Work by Bowerman and colleagues (Bowerman, 1996; Choi & Bowerman, 1992) on the acquisition of spatial terms by children learning different languages makes this clear. For children learning English (and many other languages as well), the ideas of CONTAINMENT and SUPPORT, the ideas conveyed by the words *in* and *on*, seem fundamental and early (Johnston & Slobin, 1979). But Korean children seem to make no use of these ideas in any obvious way in learning spatial terms (Choi & Bowerman, 1992). The global semantic categories of CONTAINMENT and SURFACE CONTACT/SUPPORT are not expressed in Korean in a transparent way and they are not used by Korean children. Instead, Korean children learn early and readily a distinction between TIGHT and LOOSE FIT, a distinction pertinent to their language. Any account of the development of relational representation must account for this diversity among developmental progressions.

But Korean children seem to make no use of these ideas in any obvious way in learning spatial terms (Choi & Bowerman, 1992). The global semantic categories of CONTAINMENT and SURFACE CONTACT/SUPPORT are not expressed in Korean in a transparent way and they are not used by Korean children. Instead, Korean children learn early and readily a distinction between TIGHT and LOOSE FIT, a distinction pertinent to their language. Any account of the development of relational representation must account for this diversity among developmental progressions.

Fact 2: Object categories are easier and earlier than relational categories.

Children's early vocabularies of their first 50 to 100 words provides one line of evidence suggesting that object categories are more easily learned that relational categories (see Gentner, (1982); Gentner & Boroditsky (in press) for reviews and discussion). Indeed, for a long time, object names were considered privileged at the start of word learning. Certainly common nouns dominate in English-speaking children's early vocabularies; relational terms (verbs and spatial terms) are rare (Gentner, 1982; Macnamara, 1982; Nelson, 1973). The universality of the noun advantage across languages is currently under attack (Bloom et al., 1993; Tardif, 1996; Gopnik & Choi, 1995). In particular, it has been suggested that in certain languages such as Korean (Gopnik & Choi, 1995) and Mandarin (Tardif, 1996) verbs are as or more prevalent in early vocabularies. There are many controversies in this literature and studies contradicting the claim of a verb advantage in Korean and Mandarin (Au, Dapretto, & Song, 1994; Gentner, 1982; Tardif, 1996). The controversies and ambiguities all revolve around how to measure the words in children's vocabularies. However, the best summary at present is that when proper nouns are included in the counts, the noun advantage over relational terms varies from a 4:1 noun advantage in early vocabulary by some measures in some languages (Au et al., 1994) to near 1:1 parity by other measures in other languages (Tardif, 1996; see Gentner & Boroditsky, in press for a more detailed discussion.) All in all, there would seem to be a bias toward learning nouns over relational terms in early word learning, but just as clearly this advantage depends to some degree on the language being learned and on parents' usual means of talking to their children.

The evidence on early vocabularies concerns the words children say without consideration of how they understand them. Stronger evidence for the claim that object concepts are easier and earlier than relational concepts derives not from children's word productions but from their comprehension of common nouns versus relational terms. There are extensive literatures in both areas although they are difficult to compare because of differing methods, ages of subjects, and empirical questions. These differences derive directly from the apparent ease with which children learn object categories as opposed to their difficulty in learning in relational categories. The key question for researchers who study early noun acquisition is how it is that children learn so many nouns so rapidly and with so few errors. The only errors consistently studied in this literature are the overextension errors in production typically noticed at about the time productive vocabulary growth begins to accelerate. However, these errors may not be category errors per se. Instead, these overextensions (for example, calling a zebra "doggie") may reflect pragmatic strategies or retrieval errors (Gershkoff-Stowe & Smith, 1997; Huttenlocher, 1974). Consistent with this idea is the rarity of overextensions in comprehension (see, for example, Naigles & Gelman, 1995).

In contrast, the key question for researchers who study the acquisition of relational terms is why they are so difficult to learn. The central phenomena are comprehension errors. Long after children begin to use relational terms, when they are as old as three, four, or even five years, their interpretations are errorful. Preschool children make such mistakes as misinterpreting *in front of* to mean 'near,' *put* to mean 'give,' *higher* to mean 'on top' (Johnston & Slobin, 1979; Bowerman, 1994; Clark, 1971; Gentner, 1975; Kuczaj & Maratsos, 1975; Smith, Cooney, & McCord, 1986)). Simply, common object categories are for the most part trivially easy for children to acquire whereas relational terms exhibit a protracted and errorful course of development.

This difference is also evident in artificial word learning studies. In these studies researchers present children with a novel object or event and label it. Children's interpretation of the label is measured by the kinds of other objects to which they generalize the newly learned label. Considerable evidence indicates that by 18 months (and quite possibly before), children systematically generalize novel nouns to new instances that are in the same taxonomic category (Markman, 1989; Waxman, 1994; Smith, 1995). There are fewer studies of children's generalizations of novel relational terms and most involve children at least 3 years of age and older. But the results of these studies are markedly different from those concerning object terms. First, young children's judgments are more variable and less systematic (Landau, 1996); they sometimes err by interpreting relational terms as labels for one of the objects entering in the relation (Kersten & Smith, 1998; Ryalls, Winslow, & Smith, in press); and they are much more conservative in their generalizations (Tomasello, Akhtar, Dodson, & Rekau, 1997).

In sum, one major fact to be explained in developing a theory of relational representations is the relative difficulty in acquiring words to talk about relations as opposed to words to talk about objects.

Fact 3. Understanding relations is dependent on the specific objects entering into those relations.

The developmental evidence indicates that there is not some magical point in development at which children become able to use relations. Rather, relational development appears to progress domain by domain — with children understanding relations in domains in which they are knowledgeable and reasoning poorly in domains in which they are relative novices (see Gentner & Rattermann, 1994, for a review). Thus, one sees in development the same developmental trend over and over in different domains — first children center on objects, then as they know more about the specific relational domain, they attend to relations presented in known contexts and with known objects, and ultimately, they attend to and/or reason about the relation across diverse kinds of objects and settings. That is, they progress from a more similarity-based to a more abstract understanding of relations within each domain.

This trend, for example, is evident in four- to seven-month-old babies' attention to the relations of OVER and UNDER. In one study, Quinn (1994) used a familiarization paradigm. This paradigm makes use of infants' increased attention to novelty. During familiarization trials, the infant is presented repeatedly with stimuli from one category and then on test trials is presented with novel stimuli that are either in that category or not. The reasoning is this: if infants perceive the within-category test stimulus to be like (in the same category as) the familiarization set, then they should find it boring in post-familiarization. In contrast, if infants see the out-of-category test stimulus as unlike (not in the same category as) the familiarization set, they should show increased attention to this stimulus because it will be perceived as novel. Figure 7 (top) shows one familiarization and test set used by Quinn. The infants were repeatedly shown the a and b familiarization sets — sets which depict the relation of OVER. On test they were presented four novel test stimuli, two of which also depicted the relation of OVER and two the relation of UNDER.

Quinn found that four-month-old infants looked longer at the UNDER test set, indicating that they saw these two stimuli as more different from the familiarization stimuli than the OVER test set. Analogously when infants were presented the familiarization stimuli depicting UNDER (bottom of Figure 7), they looked more at test events depicting OVER than test events depicting UNDER. These results show that four-month-olds, at least in this context, are able to attend to the relations among the line and dots and not to just the individual objects.

In a subsequent study, Quinn and colleagues (1996) varied the shapes of the components of the stimulus displays from familiarization to test. This variation among the objects involved in the relational display disrupted four-month-olds' performance. When the objects changed, all the test displays apparently looked new and therefore were attended to. In contrast, older infants, seven-month-olds, looked less at the test displays presenting the same relational configuration and more at the test displays presenting the different configurations even when the component objects in the displays varied. Thus, seven-month olds but not younger infants were able to generalize over different kinds of objects involved in the relation. The developmental trend is from object-centered conservative generalizations to ones apparently based on a more abstract representation of the relation.

Gentner & Rattermann (1991) review a number of other studies of infants' attention to relations that make the same point: attention to relations is at first highly dependent on the objects involved and becomes less so with development. They also review numerous studies of relational concepts and reasoning in much older children that again show the same trend from more object-based relations to more abstract ones. In an unpublished study by Ratterman, Gentner, and DeLoache cited in Gentner & Ratterman (1991) three- and four-year-olds were presented with the following task: The experimenter and child were each given three objects as illustrated in Figure 8. The experimenter selected one from her set as the "winner" and the child's task was to select the corresponding object. Young children had considerable difficulty choosing relationally and tended to choose the object from their

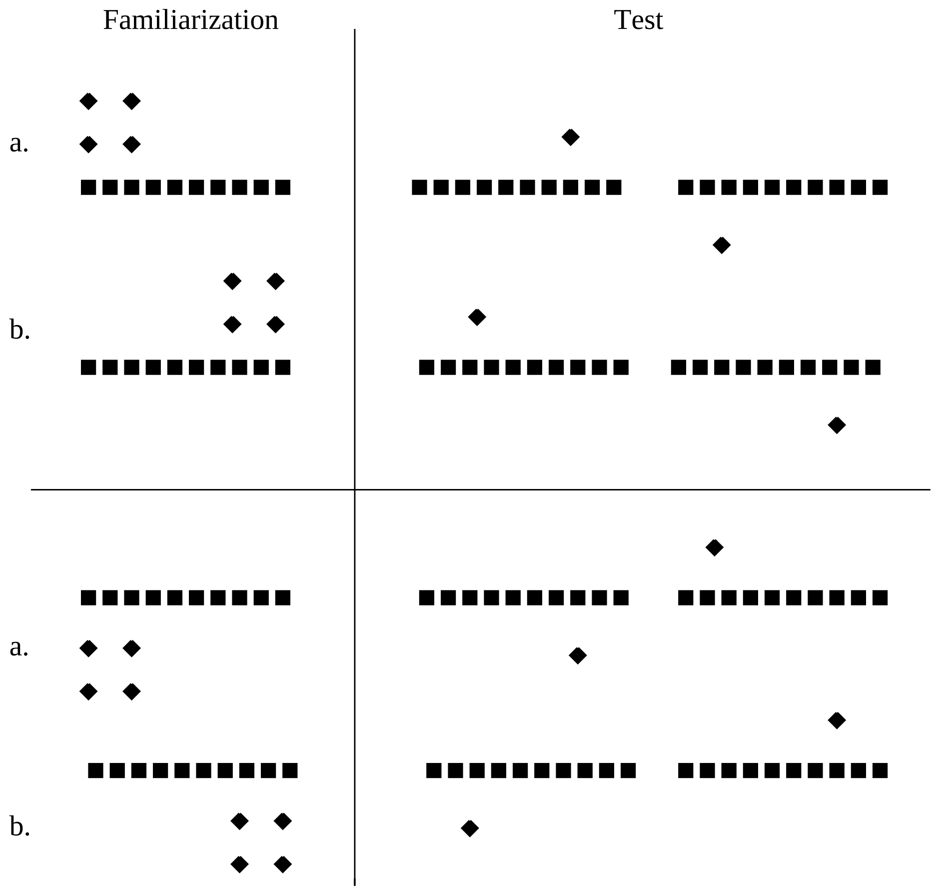

Figure 7. Stimuli used by Quinn to study infants' notions of OVER and UNDER. The infant is familiarized with patterns like those on the left and then presented patterns like those on the right. The subject should attend more to test stimuli which are perceived as novel.

set that matched the experimenters in actual size or in other object properties. However, in a subsequent experiment, it was shown that young children could respond relationally if the task was presented in a domain that they understood — specifically, if the objects in each set were presented as "the daddy one," "the mommy one," and "the baby one." This result, along with that of a study by Kotovski & Gentner (1996) in which infants were progressively trained to make more and more abstract relational inferences, demonstrate that the developmental change is driven more by experience than by maturation or age.

Fact 4. Object properties are relevant to understanding relations.

It is probably a good thing that infants and children begin learning about relations by attending to the objects and context in which they encounter a relation. Consider, for example, real world instances of CONTAINMENT and SUPPORT. What objects can go "inside" other objects, what objects can support other objects, and how one physically realizes these relations depends very much on the specific objects and their perceptible properties. Titzer, Thelen, & Smith (1998) recently demonstrated this point in a study of infants understanding of containment and support given transparent or opaque landmarks. The experimental

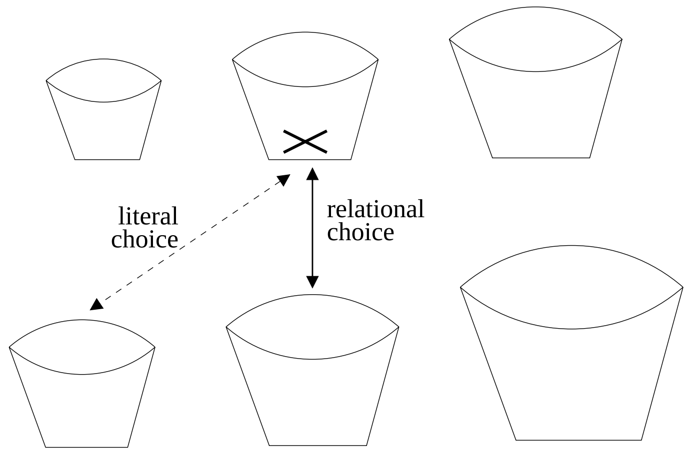

Figure 8. Stimuli used by Ratterman, Gentner, & DeLoache to study children's understanding of relations. The child's task is to select the object in the lower set which corresponds to the object in the upper set that has been selected by the experimenter.

procedure built on an earlier study by Diamond (1990). She showed that infants are more successful in retrieving desired toys from opaque rather than transparent containers. From one perspective, this is a perplexing result: an object in an opaque container cannot be seen and must be remembered whereas the desired object in a transparent container can be continuously seen. However, from the perspective of a concept of containment, Diamond's results are not surprising. Transparent containers present unique perceptual cues to surfaces and openings. And, indeed, the babies in Diamond's study seemed not to understand where the openings were in the transparent containers as they tried to reach through the transparent surface to retrieve the toy.

Titzer et al. (1998) tested the idea that specific experience with transparent containers may be essential to successful object retrieval from such containers. In a two-month training study, they exposed eight-month-old infants to transparent containers; control infants were given opaque containers for the same period. The containers were identical in both cases except for opaqueness and varied in shape and size from small cups to buckets large enough for a baby to put over his or her head or even to sit on. The infants were given no special training; parents were simply asked to give the infants the containers to play with daily for at least 10 minutes over the two-month experiment. When the infants were brought back to the laboratory at ten months and tested in Diamond's procedure, the infants experienced with transparent containers and the control infants performed differently. Specifically, only the infants who had played with transparent boxes knew how to rapidly retrieve objects from transparent boxes. Two months of perceiving and acting on transparent containers taught these infants the unique perceptual cues relevant for containment in a transparent receptacle.

The infants generalized their learning about transparent containers to problems concerning SUPPORT. Titzer et al. (1998) tested both the infants trained with transparent containers and those trained with opaque containers on an apparatus known as the visual cliff. This is a highly studied device invented by Gibson & Walk (1964) to test infants sensitivity to depth cues. In traditional testing, the apparatus is a transparent table top that sits over a substantial drop. Infants are placed on the shallow side near the drop and their behavior is observed. And the traditional result is that ten-month-old infants avoid the cliff and the deep end, crawling away from the visual cliff to a secure position on the shallow side. Titzer et al. (1998) observed the infants in the control condition, the ones who had only played with opaque containers, also retreat from the visual cliff, apparently believing — despite the solid transparent surface below them — that they might fall. In contrast, infants trained with the transparent containers acted unlike infants in any other study of the visual cliff; they confidently and happily crawled right over the visual cliff. These infants had apparently learned the visual cues to transparent surfaces and knew what typically developing infants with limited experiences with transparent surfaces do not: solid transparent surfaces support just as do opaque ones. Clearly, real-word experiences and specific object properties matter in understanding SUPPORT and CONTAINMENT in the context of transparency; we suspect that this is the same for understanding of SUPPORT and CONTAINMENT in the context of opacity. The real-world use and recognition of relations requires their grounding in specific object properties

Indeed, objects properties are integral to all spatial concepts. For example, when shown a block on a box while being told "the block is acorp the box" English speakers interpret *acorp* to mean 'on'. In contrast, when shown a stick on a box, English speakers interpret *acorp* to mean 'across' (Landau, 1996). The shape of the trajector and the shape of the landmark matter. The relevance of object properties is also apparent in judgments of containment. Consider, for example, panels A and B in Figure 9. Most people judge the apple not to be in the bowl in panel A but to be in the bowl in panel B. An apple on top of other fruit that is contained in a vessel is IN. The relevant properties are not just spatial ones; object categories matter as well. For example, the apple is in the bowl when sitting on other apples but is not in the bowl in panel C when sitting on blocks that are in the bowl.

Feist & Gentner (1998) showed in a recent study of adults' judgments of *in* and *on* that the way in which the very same object was categorized determined relational judgment. They varied the curvature of the landmark as illustrated in Figure 10 In some conditions, the displays were unlabeled; in others the trajector was labeled as an animate fly and the landmark as a bowl, plate, or dish. All variables mattered. Inanimate trajectors were more likely to be judged *in* than animate ones, more curved landmarks yielded more *in* judgments than less curved ones, and labeling the landmark as a bowl rather than as a plate yielded more *in* judgments. Clearly, object labels — not just their spatial properties — matter when adults make relational judgments. Thus, relational development must not consist so much of stripping away all object information, but must instead consist of learning the particular kinds of object properties relevant to particular relations.

Fact 5. Relational concepts have a category structure.

Other research indicates that relational concepts seem to be like object concepts in having a graded similarity structure. For example, Logan and Sadler (1996) asked adults to rate the goodness of representations such as those shown in Figure 11 as instances of the spatial terms *above*, *below*, *over*, *under*, *left of* and RIGHT OF. The adults' judgments were highly organized and consistent, and as shown in Figure 12; various instances of trajector-

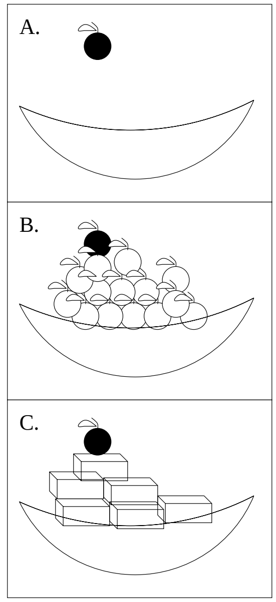

Figure 9. Effect of the properties of the trajector and landmark on judgments of a relation. An apple is construed as *in* a vessel when it is supported by other fruit which is in the vessel (B) but not when it is not supported by anything in the vessel (A) or even when it is supported by objects other than fruit (C).

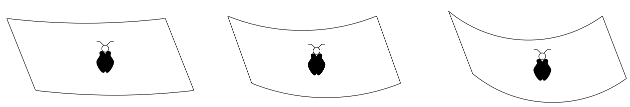

Figure 10. Stimuli used by Feist & Gentner to study adults' judgments of *in* and *on*. Curvature of the landmark was varied.

landmark relatedness are better and worse instances of the concepts.

Children's acquisition of relational terms also exhibit a graded category structure. Smith, Cooney & McCord (1986; see also Ryalls, Winslow & Smith, in press) investigated children's understanding of the words *higher* and *lower*. Three- and four-year-old children were presented with two objects at a time and asked to indicate "which is higher?" or on other trials "which is lower?" The objects were discs about a foot in diameter that were positioned (above the ground) at 1' vs. 2', 3' vs. 4' or 5' vs. 6'. The results are shown in Figure 13. When presented with the discs at the two highest locations (5' vs. 6'), young children chose correctly when asked "which is higher?" but chose incorrectly when asked "which is lower?" Conversely, when presented with the discs at the two lowest locations (1' vs. 2'), young children chose correctly when as "which is lower?" but chose incorrectly when asked "which is higher?" Judgments at the mid-height locations were intermediate. These results strongly suggest a categorical representation of *higher* and *lower* in which the best exemplar of *higher* is an object that is very high and the best exemplar of *lower* is an object that is very low. Notice also that these judgments suggest that children do not represent higher and lower as opposites — that they can know that A IS HIGHER THAN B

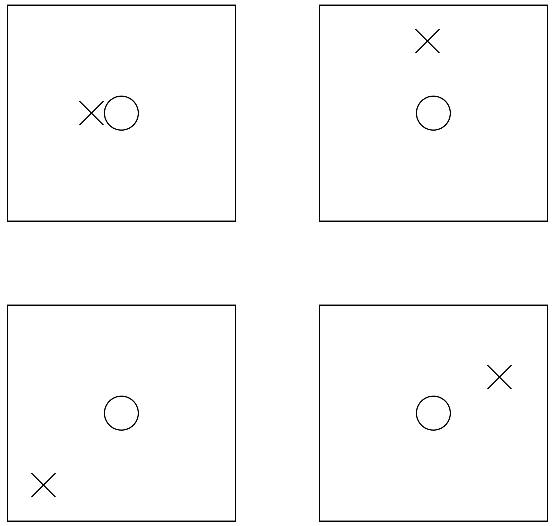

Figure 11. Stimuli used by Logan & Sadler to study adults' ratings of the goodness of spatial terms. The Xs represent trajectors, the circles landmarks.

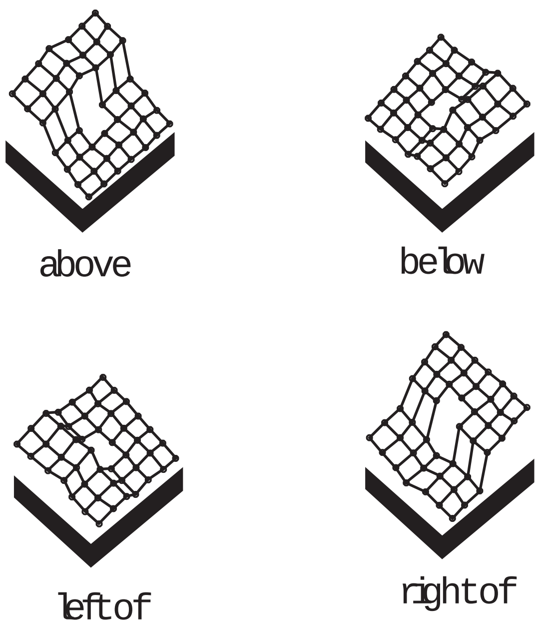

Figure 12. Average ratings for *above*, *below*, *left of*, *right of* in Logan & Sadler's goodness rating task.

without knowing that B IS LOWER THAN A.

It is not only children's representations that are categorical in this way. In comparative judgment tasks, adults show the same pattern in reaction time that children show in errors (Petrusic, 1992). For both children and adults, some instances of a relation are better instances.

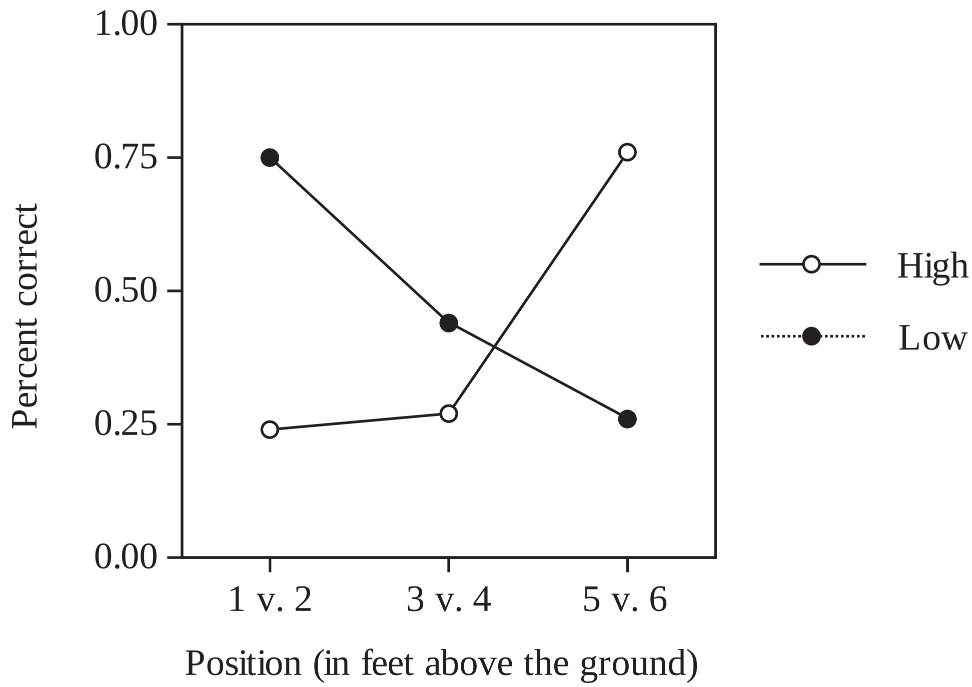

Figure 13. Children's accuracy in selecting the *lower* or *higher* of two discs placed at different heights (Smith, Cooney, & McCord, 1986).

Summary and implications

These five facts are not easily alignable with the kinds of relational representations summarized in Table 1. What kind of system can develop different relational concepts depending on its linguistic experiences in the world? Why does development go from being more object centered to more relational as a function — not of maturation or age — but of experience with that specific relation? How are the relevant object properties discovered and represented for different relations and how do these yield a graded category structure?

Past approaches all concentrate on the specification of how the objects in a relation are related to one another and all accomplish this by some explicit relation term (or explicit role name) together with a mechanism for binding the objects to the roles of the relation. But none of these approaches tells us where the relation term or roles come from.

The five empirical facts presented above suggest a new approach, one which seeks an explanation of the substrate and processes out which relational categories such as ABOVE are formed. At the center of this approach is the question of how specific relation instances are handled. A relation instance is an explicit association of a particular pair of objects, for example, the spatial relation between a book and a table above which the book is suspended. Just as object categories such as BOOK are generalizations over instances of the category, relational categories such as ABOVE are generalizations over instances of the category (see Kersten & Billman, 1997). In the next section, we present a new model of relational representation based on these ideas: Gasser & Colunga's Playpen model (Gasser & Colunga, 1997, 1998).

A new proposal about relational representations

We propose that relations are represented in terms of directly accessible relational associations which specify how objects are related to one another in terms of object features and inter-relations associations. In this approach, relation categories are built up from relational feature correlations learned on the basis of relation instances, much as object categories are built up from object feature correlations on the basis of object instances.

Representing objects

Objects are crucial in our account of relations just as they are in all accounts. For our purposes, object instances are cognitive entities (rather than entities in the external world) consisting of values on dimensions such as color, size, extent, material. Object categories such as BOWL take the form of ranges of values on each dimension. They are developed from repeated experiences with individual instances.

Playpen's representation of objects is illustrated in Figure 14. Each rectangle represents the pattern of activation on a single object dimension (e.g, vertical extent, horizontal extent, color). The particular values on these dimensions presented by this single object instance are indicated by the blackened region. Thus, Figure 14a might represent a particular object that is 8 inches tall, 20 inches long, and blue. The presented values on the three illustrated dimensions are associated by the illustrated connections among them. Figure 14b illustrates an object category which consists of a range of values (features) on the different dimensions defined by the correlations among features found over multiple instances. Thus, object-feature correlations begin with the inter-value connections that are created (or strengthened) with the presentation of an object instance.

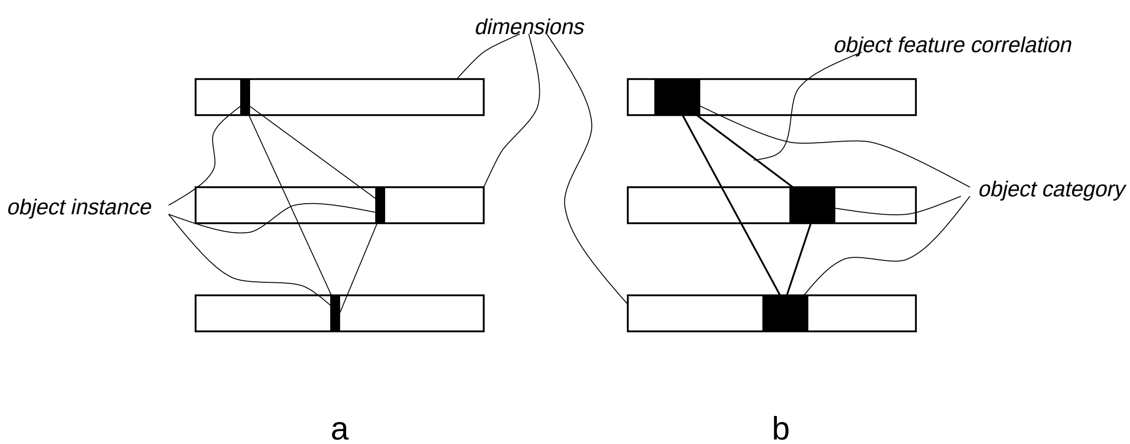

Figure 14. An **object instance** (a) consists of values on each of a set of dimensions. An **object category** (b) consists of *ranges* of values on the different dimensions defined by the correlations (bold lines) found over multiple instances.

Representing relations

When two objects are experienced simultaneously, there is the potential for an explicit connection between features of different objects. An explicit connection between features of different objects is a relation instance. Imagine for example the experience of a small round object on a flat surface. The connection between small and round would be created (or strengthened) within the object instance (and the object category) representation but the connection between small and flat could also be created (or strengthened); this would be a connection between features in different object instance representations. With the presentation of multiple relations instances with similar values for different dimensions, a relational correlation is created.

Consider first the case of a single dimension. Each of the two objects in a relation instance has a value on that dimension; one object, for example, may be big and the other little. Multiple relation instances of this sort, if repeated sufficiently, may lead to a one-dimensional relational correlation. Figure 15 illustrates a relation instance (left) and a one dimensional relational correlation (right). For example, the relation instance might be a small object and a large object represented by the specific values on the input dimension, the small object by the white region for example, and the large object by the black region. If such instances are experienced with regularity, a relational correlation would develop, consisting of the ranges of correlated values, as illustrated on the right. For this kind of a system to work, the relational correlations between distinct objects must be represented in ways that are distinguished from those used to represent the feature correlations presented by a single object instance or object category. We propose that these relational correlations are represented not by simple connections as feature correlations but by separate micro-relation units (MRUs). In the figures, these units appear as diamonds. Note that these units need have no built-in meaning, but are analogous to units in distributed representations of objects in PDP networks. That is, MRUs take on their significance as the weights connecting them to object feature units and other MRUs evolve in response to correlations in input events. In this way, the relational meaning of an event is similarity-based and object based: a generalization across multiple experiences of bundles of object features that co-occur.

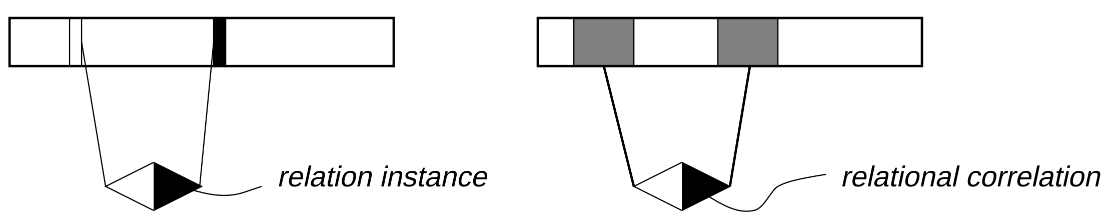

Figure 15. A **relation instance** (left) is two objects presented simultaneously. If the values of the two objects along a particular dimension correlate over several relation instances, a **one-dimensional relational correlation** can be created (right).

With experience, a learner may generalize from narrow regions of values for the two objects as instances of a relation to relative values across the whole dimension. For example, NEAR does not refer to two objects located in a range of specific absolute locations, but rather to the proximity of objects located anywhere. One way to represent such a relation is through the association of more specific absolute relational correlations with each other through a relational category unit, as shown in Figure 16 for NEAR. The category unit must point to each of the relational correlations rather than to the correlated values; thus each relational correlation must take the form of an explicit unit rather than a simple connection.

Notice that both the correlations and the category unit are MRUs.

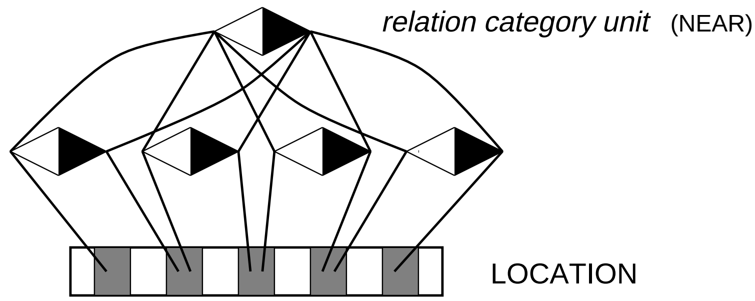

Figure 16. An example of how a relational category unit can be used to represent a relational category (NEAR) by connecting several units representing specific absolute relational correlations along a dimension (LOCATION).

As with a single dimension, a learner can generalize from absolute values to relative values across one or more of the dimensions. For example, the knowledge about the relationship between SIZE and LOUDNESS could take the form of the knowledge that relative size, wherever on the size scale, correlates with relative loudness, wherever on the loudness scale. We believe that many familiar spatial relational categories such as ON are actually learned in terms of cross-dimensional relational correlations of this type. Thus for ON, the relative location of the upper and lower boundaries of two objects, which seems to define the relation for us, correlates with the relative size and movability of the objects.

Playpen

In this section we summarize the main features of Playpen, focusing on those that are relevant for this paper. For technical details see Gasser & Colunga (1998).

Architecture

Playpen (Gasser & Colunga, 1997, 1998) is a connectionist model of the acquisition of word meaning. The network itself is a generalization of a continuous Hopfield network; that is, it consists of symmetrically connected simple processing units which respond by updating their activations when certain units are "clamped" to a particular input pattern. The long-term knowledge in the network is encoded in the weights on the connections; these are either hard-wired or learned as the network is presented with training patterns.

We divide the units into three layers (Figure 17) a VISION layer, a SPATIAL CONCEPTS layer, and a WORDS layer. In our work to date, we have treated the VISION and WORDS layers as input/output layers, and the SPATIAL CONCEPTS layer as a hidden layer. That is, when we present a pattern to the network for training or testing, it is the VISION or WORDS units which are clamped to the values in that pattern, and then it is the VISION and WORDS units which we treat as the network's response to the input pattern. The SPATIAL CONCEPTS units are hidden in the sense that they are not directly accessible "from the outside" (though we can of course observe their activations as the network runs).

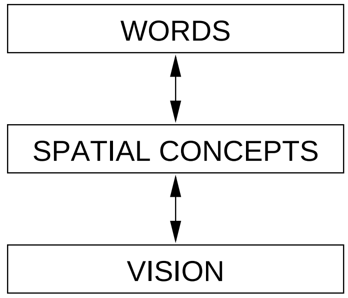

Figure 17. Basic Playpen architecture. Each rectangle represents a layer of MOUs and/or MRUs and each arrow a pattern of connectivity between layers.

As discussed above, units in the network are of two basic types. **Micro-object units (MOUs)** represent primitive features of objects. These are just the familiar processing units in networks of this general type except that in addition to an activation, each has a **phase angle**. Phase angles provide a solution to the **binding problem**; activated MOUs with similar phase angles represent features of a single object. MOUs which are in phase with one another affect each other's activation more strongly, units connected by positive weights tend to attract each other's phase angles, and units connected by negative weights tend to repel each other's phase angles.

Micro-relation units (MRUs) represent features of relations. While each MRU is a single unit in the sense that it has a single activation, it has two separate **micro-roles**, and separate connections from each micro-role to other MOUs and MRUs. Each micro-role has its own phase angle, corresponding to the phase angle of one of the objects that is being related. All else being equal, an MRU tends to be activated to the extent that it receives inputs on its two micro-roles which are maximally out-of-phase, that is, inputs representing two distinct objects.

Relational correlations take the form of connections between MRUs. Each connection is really a pair of connections, mapping the micro-roles of one MRU to those of the other. Positive connections cause one MRU to tend to activate the other and align its phase angles with it. Negative connections cause one MRU to tend to inhibit the other.

Processing

As noted above, the model is run on a pattern by clamping (fixing the activations and phase angles of) some of the units on the VISION or WORDS layers. The units in the network are then updated repeatedly — that is, their activations and phase angles are updated — until the changes have stabilized. The pattern of activation of the units on the VISION and WORDS layers at this point represents the network's response to the input pattern. Note that the model can be run in both the "comprehension" and the "production" directions, or even in combinations of the two, by clamping units on different layers.

Learning

Weights on the connections joining units are learned as the network is trained on a set of patterns. Learning is through a variant of Hebbian learning known as **Contrastive**

Hebbian Learning (Movellan, 1990). Learning takes place in two phases. During the positive phase, an input pattern and the appropriate pattern are both clamped, the network is allowed to stabilize, and the weights are adjusted in proportion to the correlation between the activations of the connected units. During the negative phase, only an input pattern is clamped, the network is again allowed to stabilize, and the weights are adjusted in proportion to the *anti-correlation* between the activations of the connected units. When the patterns have been learned, the two changes cancel each other out because the network's behavior in the two phases is identical. That is, the network produces the desired output for a given input during the negative phase.

Simulation

In this section we illustrate the behavior of the model with a simple simulation.[2] In this simulation a network learns both inter-MRU correlations across different dimensions and also simple relational categories. The simulation also illustrates how the ease of learning particular words can depend on the match between the words and the relational correlations that the system has already picked up on.

We first defined a set of correlations among non-linguistic dimensions and a set of correlations between the non-linguistic dimensions and words. There were two non-linguistic dimensions and four possible micro-relations within each of these dimensions. For example, micro-relation A represented a relation between a very low value for one object and a very high value for another object on Dimension 1, and micro-relation E represented a similar micro-relation on Dimension 2. There were correlations between these micro-relations across the two dimensions. That is, a pair of objects with particular values on one dimension tended to have particular values on the other. For example, micro-relation A on Dimension 1 correlated with micro-relations E and F, but not with micro-relations G and H, on Dimension 2. The relational correlations are shown in Figure 18a.

We defined two "languages," an *Easy* language, which agrees with the non-linguistic correlations, and a *Hard* language, which disagrees with the non-linguistic correlations, as shown in Figure 18b. Each language consists of two relational words. For the *Easy* language, the categories in the world agree with the categories associated with the words. That is, each of the two correlational clusters existing in the world is tied to one of the words in the language. For the Easy language, the values on both input dimensions are relevant to the word choice. For example, knowing either that a pair of input objects has micro-relation A or B on Dimension 1 or has micro-relation E or F on Dimension 2 tells us that Word 1 is appropriate. For the *Hard* language, the words cut across the two pre-linguistic correlational clusters in such a way that the word describing a pair of values along the two dimensions is determined by the value along Dimension 1 only. For example, according to the pattern of correlations between dimensions in Figure 18a, the pairs of values labeled A-F and B-E should be in the same category but they are assigned to different words in the Hard language: A-F is associated with Word 1, while B-E is associated with Word 2. In the Hard language, the value along Dimension 2 is not predictive of the linguistic category; knowing that an input pair of objects has micro-relation E on Dimension 2 tells us nothing about which word in the Hard language is appropriate. Note, however, that Dimension 1 is

[2]For details of this and two other related simulations, see Colunga & Gasser, 1998.

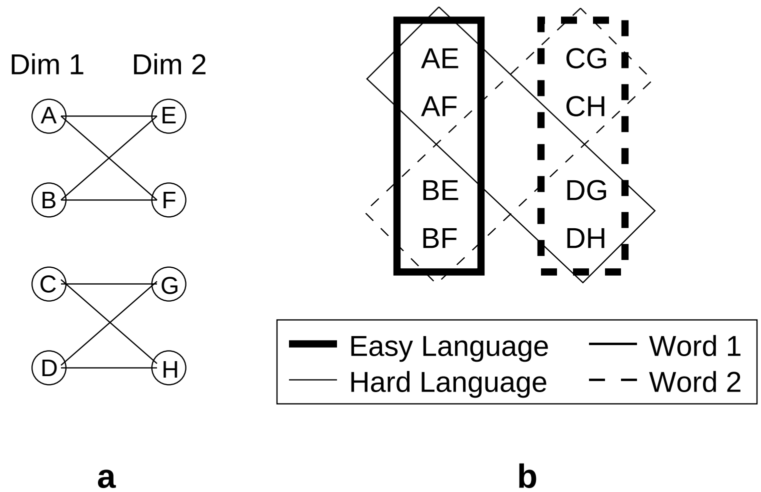

Figure 18. Correlations used in simulation. A, B, C, D, E, F, G, and H represent possible micro-relations between features within Dimensions 1 and 2. (a) The micro-relations correlate with each other across the dimensions in the two clusters shown. (b) Possible pairings of relations on the two dimensions are associated with one or the other of two words. In the Easy language, the words agree with the non-linguistic correlations; in the Hard language, the words correlate only with micro-relations on Dimension 1.

relevant for the Hard language; knowing that an input has micro-relation A on Dimension 1 tells us that Word 1 is appropriate.

The architecture of the network used in the simulation is shown in Figure 19.

The goal of the experiment is to see how the different correlational patterns both between dimensions and with the words affect the difficulty of learning the two languages. The network was trained and tested on two different tasks. Training began with a Pre-linguistic Phase in which the task was **Non-linguistic Pattern Completion** That is, for each trial the network was presented with a pattern on one of the visual dimensions and expected to produce an appropriate pattern on the other. (Note that there are always two possibilities for the appropriate pattern.) The network can learn to solve this task using the connections joining the VISION and SPATIAL CONCEPTS layers or the connections between the two SPATIAL CONCEPTS layers. This phase continued for 30 repetitions of the relevant training patterns (epochs). Next, during a Linguistic Phase, Pattern Completion training was discontinued, and the networks were trained on **Production** for seven epochs. For this task, the network was presented with a pattern on the VISION layer and expected to output a word. Training in the Linguistic Phase began with weights of 0.0 connecting the SPATIAL CONCEPTS to the WORDS units, so the network was initially unable to produce any words.

We predicted that the Easy language would be learned faster than the Hard language during the Production phase because the Easy language categories agreed with the non-linguistic categories.

During the Pre-linguistic Phase, the networks mastered the Pattern Completion task by learning weights between the two Hidden layers representing the non-linguistic correla-

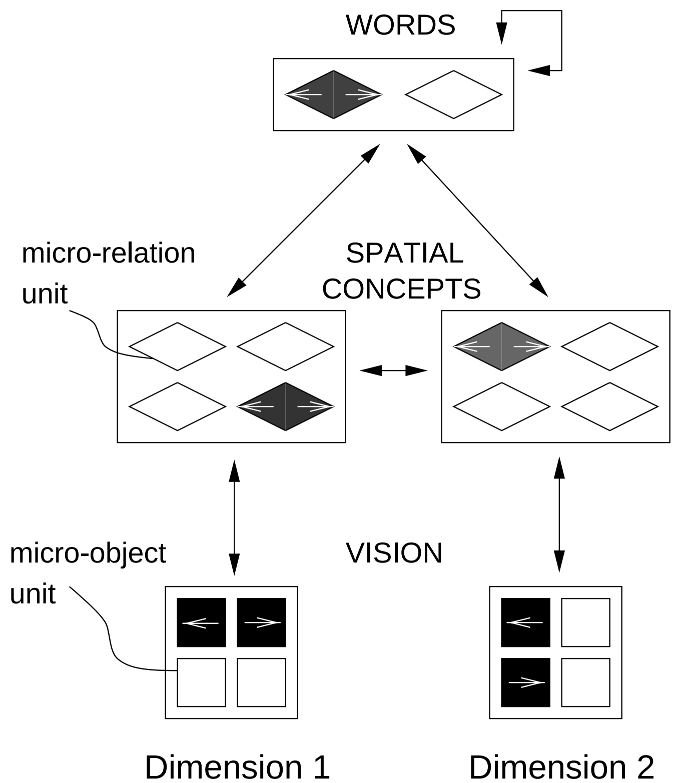

Figure 19. Network architecture. Micro-object units are represented by squares, micro-relation units by diamonds. Arrows indicate complete connectivity between layers. Each SPATIAL CONCEPTS MRU is associated with a pair of VISION MOUs. A possible pattern across the network is shown. Darkness indicates activation, and arrow direction indicates relative phase angle.

tions. Results for the Linguistic Phase are shown in Figure 20, starting with performance after one epoch.

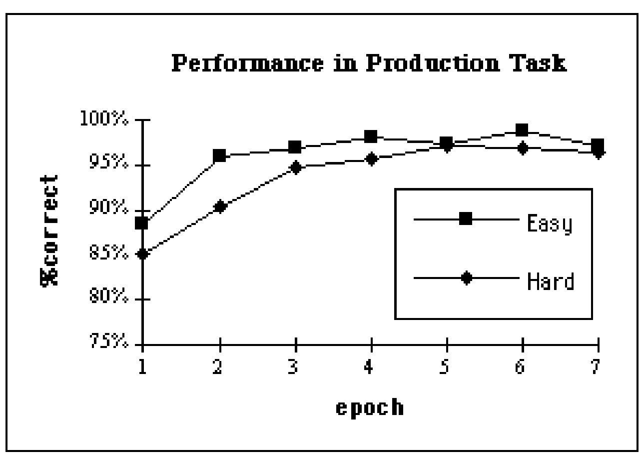

Figure 20. Results for simulation. The Easy language is learned faster than the Hard language.

The data were submitted to a 2(Language) $*$ 7(Epoch) analysis of variance for a mixed design. This analysis revealed a main effect of epoch, indicating that the networks get better as they receive more training. More importantly, as predicted, there is a main effect of language ($p < .001$). Thus, as predicted, the Easy language is learned faster than the Hard language, although by the end of the training the two networks have comparable performance. No interactions between language and epoch were found.

How Playpen fits the five facts

1: Language matters.

In Playpen language is crucially involved in the learning of spatial concepts. Spatial concepts are learned as the network picks up on correlations between spatial features of pairs of objects in the world. When spatial terms are presented, these take part in the correlations as well. Because spatial terms in different languages correlate differently with spatial features, the particular language being learned will affect the way spatial concepts are learned in the network. Particular correlations in the world may or may not agree with the words being learned, and this can result in two sorts of effects. Words that match the correlations in the world are learned faster than those that do not. When words agree with the values on a particular visual/spatial dimension, the network can use the language to help it process visual scenes, even in ostensibly non-linguistic tasks.

2: Object categories are easier and earlier than relational categories.

In Playpen, learning a category means learning the correlations among a set of features. For both objects and relations, this presupposes the preprocessing required for extracting the features, for example, for detecting the edges that are the boundaries of objects. For the learning of an object category all that is then required is that two or more object features be activated simultaneously. The co-occurrence of a particular texture and a particular color may be the beginning of a category such as YOGURT. For the learning of a relational category, there must also be a set of activated units, in this case relation units, but now there is the additional constraint that the phase angles of the corresponding micro-roles of the relation units match one another. That is, the relation units need to have been activated in response to the features of two objects, which the system must have distinguished from one another.

Note that the network can treat a scene consisting of more than one object as a single object, activating object units and learning feature correlations, but learning nothing relational about the scene because no relation units have been activated.

Because they do not presuppose the object segregation that is required for the learning of relational categories, object categories are easier and earlier than relational categories in Playpen.

The fact that each relation unit brings together two features or sets of features also means that, all else being equal, more relation units than object units are required to cover the space of possibilities. Thus, given the same amount of resources in the network for objects and relations, the space will be covered more sparsely by relation units. This should lead to slower learning of relations.

3: Understanding relations is dependent on the specific objects entering into those relations.

Learning in Playpen, as in neural networks in general, starts from specific instances. For objects, the first (weak) correlations that are learned correspond to individual objects or small sets of objects, a particular cup or the three cups most often presented. Only after the presentation of a number of exemplars does the pattern of connectivity come to reflect the more general associations we expect for categories such as CUP. This does not mean that specific associations are lost; as long as the network continues to be presented frequently with specific cups, those cups will continue to exist as "micro-categories" within the network.

For relations the same context-specifity applies, but now it is the specific object entering into the relation that are relevant. Given repeated presentation of the same doll in the same basket, the network's early representation of CONTAINMENT will be specific to that doll and that basket. When the basket is presented with different dolls in it, the relational concept becomes more general in one way. When different baskets are presented, it becomes more general in another way. Only after considerable variation among both the container and the contained objects would the network come to represent the set of associations characterizing CONTAINMENT (if ever). But the more specific subtypes of CONTAINMENT would not be lost.

Thus in their development and in their endstate, relations in Playpen are bound to the specific objects that were presented during training.

4: Object properties are relevant to understanding relations.

Rather than directly associating full-blown objects, relation units in Playpen associate the features of two objects. Thus relations in Playpen can "look into" their component objects in a way that is not possible in approaches where the components of relations are essentially symbolic, labels representing already categorized objects.[3] In the learning of the SUPPORT relation, for example, the shape and even the color of the supported object is potentially relevant to the relational associations that develop in Playpen, but this is independent of how the supported object is categorized by the system, that is, what word unit it turns on.

However, because relational categories involve correlations of all sorts, the categories of the component objects may also be relevant. If the objects in a spatial relation instance are categorized by the network (that is, they activate category label units), these categories may play a role in the particular relation units that are activated, just as primitive object features are.

5: Relational concepts have category structure.

In Playpen both object categories and relational categories have a graded similarity structure. As in other neural networks, a category *is* the instances that make it up. When an instance is presented, it results in changes in the weights between units representing correlations between features. Each of these weights is the combined result of all of the instances the network has seen, and each category is realized as a pattern of weights among a number of units. When a new object is presented to the network, it activates object units. To the extent that these units are strongly associated with others through a pattern

[3]Note, however, that the approach we are proposing does not exclude a representational system in which the primitive features of objects are built into the system and each input object unit has an inherent meaning and each relation unit associates a fixed set of these object units.

of learned weights, the object is a good instance of the category represented by those weights and will tend to activate the category label unit.

In the same way, relation instances are better or worse instances of relational categories to the extent that they activate the relation units which are associated with others through the learned weights that make up the relational category. A network trained on instances of HIGHER in which one object is far from the other will fail to activate the HIGHER unit when presented with two objects which are close together.

Conclusion

In this chapter we have adopted a very different perspective on the relationship between language and spatial cognition than have the other authors in this book. We have focused on the *emergence* of relations and language rather than the end-state. In doing so, we have been concerned with the very nature of relations and have presented an account of the "stuff" out of which relations are constructed. We believe that this groundwork is essential if we are to understand the full-fledged adult system because we believe the adult system is the product of more of the same sort of process we have proposed for the early emergence of relations, the learning of more and more abstract relational correlations.

In taking this low-level, developmental approach, we have ignored most of the data that are of interest to the other authors in the book, data that have led people such as Jackendoff to posit separate representational modules for spatial concepts, conceptual representations, and language, as well as interface modules connecting them. Thus it should be not be surprising that the picture we offer is considerably simpler than the one offered by others. In a sense our picture is one in which WORDS and VISION are mediated by a single "interface module", which we have called SPATIAL CONCEPTS, a layer of micro-object and micro-relation units which is simultaneously under the influence of units in the WORDS and VISION layers.

However, we are aware that the mapping from visual/spatial representations to words and syntax is not a simple one. The SPATIAL CONCEPTS layer in our model is not meant to be taken as a layer of completely connected relation units and object units. Rather we expect this layer in the adult system to consist of multiple sublayers, each performing some form of transformation on the patterns it receives from above or below, each connected to its neighbors by an "interface" in the form of a simple pattern of connectivity. As yet we have nothing specific to say about the ultimate internal structure of the SPATIAL CONCEPTS layer nor about how much of this structure is built in and how much develops as spatial categories and words are learned. The point is that our proposals are not incompatible with accounts in which different levels of representation reside in different architectural layers. Our contribution to the question of the putative modules and interfaces joining visual and linguistic representations is twofold. (1) The modularity may be a matter of degree; that is, the representations may simply be more and more linguistic (under the influence of words and syntax) as they appear closer to the strictly linguistic portion of the neural network and more and more visual (under the influence of specifically visual patterns) as they appear closer to the visual portion of the network. (2) Modularity (to whatever degree) can *develop*, and the interface between language and vision is a place where we might expect this. Therefore, a major goal should be elucidating the mechanisms by which development takes place in response to the visual and linguistic patterns it is exposed to.

Spatial cognition and spatial language are fundamentally relational. In fact, given the importance of space in the young child's world, space may be the place to look for the emergence of relational knowledge. A range of well-attested data are consistent with the view that we have proposed: relational categories, like object categories, take shape as the child picks up on the rich set of correlations available to her in the world and in language, and, like object categories in neural network models, they are built up out of primitive connectionist units.

References

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

Bloom, L., Tinker, E., & Margulis, C. (1993). The words children learn: Evidence against a noun bias in early vocabularies. *Cognitive Development*, *8*, 431–450.

Bowerman, M. (1994). Learning a semantic system: What role do cognitive predispositions play? In P. Bloom (Ed.), *Language acquisition: Core readings* (pp. 329–363). Cambridge, MA: MIT Press.

Bowerman, M. (1996). Learning how to structure space for language: A crosslinguistic perspective. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 385–436). Cambridge, MA: MIT Press.

Choi, S., & Bowerman, M. (1992). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, *41*, 83–121.

Clark, E. V. (1971). On the acquisition of the meaning of *before* and *after*. *Journal of Verbal Learning and Behavior*, *10*, 266–275.

Colunga, E., & Gasser, M. (1998). Linguistic relativity and word acquisition: A computational approach. *Annual Conference of the Cognitive Science Society*, *20*, 244–249.

Diamond, A. (1990). Developmental time course in human infants and infant monkeys, and the neural bases of inhibitory control in reaching. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 637–676). New York: National Academy of Sciences.

Gasser, M., & Colunga, E. (1997). *Playpen: Toward an architecture for modeling the development of spatial cognition* (Tech. Rep. No. 195). Bloomington, IN: Indiana University, Cognitive Science Program.

Gasser, M., & Colunga, E. (1998). *Where do relations come from?* (Tech. Rep. No. 221). Bloomington, IN: Indiana University, Cognitive Science Program.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition* (pp. 211–246). San Francisco: Freeman.

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: Lawrence Erlbaum.

Gentner, D., & Boroditsky, L. (in press). Individuation, relativity, and early word meaning. In *Language acquisition and conceptual development.* Cambridge: Cambridge University 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). Cambridge: Cambridge University Press.

Gentner, D., Rattermann, M. J., Markman, A., & 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 Associate.

Gershkoff-Stowe, L., & Smith, L. B. (1997). A curvilinear trend in naming errors as a function of early vocabulary growth. *Cognitive Psychology*, *34*, 37–71.

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

Halford, G. S., Wilson, W. H., Guo, J., Gayler, R. W., Wiles, J., & Stewart, J. E. M. (1994). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J. Barnden (Eds.), *Advances in connnectionist and neural computation theory, vol. 2: Analogical connections* (pp. 363–415). Norwood, NJ: Ablex.

Halford, G. S., Wilson, W. H., & Phillips, S. (forthcoming). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences.*

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480–517.

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

Huttenlocher, J. (1974). The origins of language comprehension. In R. Solso (Ed.), *Theories in cognitive psychology.* Potomac, MD: Lawrence Erlbaum.

Jackendoff, R. (1996). The architecture of the linguistic-spatial interface. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 1–30). Cambridge, MA: MIT Press.

Johnston, J. R., & Slobin, D. I. (1979). The development of locative expressions in English, Italian, Serbo-Croatian, and Turkish. *Journal of Child Language*, *6*, 529–545.

Kersten, A., & Smith, L. B. (1998). *Attention to novel objects during verb learning.*

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

Kuczaj, S. A., & Maratsos, M. P. (1975). On the acquisition of *front*, *back*, and *side*. *Child Development*, *46*, 202–358.

Landau, B. (1996). Multiple geometric representations of objects in languages and language learners. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 317–363). Cambridge, MA: MIT Press.

Langacker, R. (1987). *Foundations of cognitive grammar.* Stanford: Stanford University Press.

Logan, G. D., & Sadler, D. D. (1996). A computational analysis of the apprehension of spatial relations. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 493–530). Cambridge, MA: MIT Press.

Macnamara, J. (1982). *Names for things: A study of human learning.* Cambridge, MA: MIT Press.

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

Movellan, J. (1990). Contrastive Hebbian learning in the continuous Hopfield model. In D. Touretzky, J. Elman, T. Sejnowski, & G. Hinton (Eds.), *Proceedings of the 1990 Connectionist Models Summer School* (pp. 10–17). San Mateo, CA: Morgan Kaufmann.

Naigles, L. G., & Gelman, S. (1995). Overextensions in comprehension and production revisited: Preferential-looking in a study of *dog*, *cat*, and *cow*. *Journal of Child Language*, *22*, 19–46.

Nelson, K. (1973). Some evidence for the cognitive primacy of categorization and its functional basis. *Merrill-Palmer Quarterly*, *19*, 21–39.

Petrusic, W. (1992). Semantic congruity effects and theories of the comparison process. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 962–986.

Pinker, S. (1994). How could a child use verb syntax to learn verb semantics? *Lingua*, *92*, 377–410.

Quinn, P. C. (1994). The categorization of above and below spatial relations by young infants. *Child Development*, *65*, 58–69.

Quinn, P. C., Cummins, M., Kase, J., Martin, E., & Weissman, S. (1996). Development of categorical representations for above and below spatial relations in 3- to 7-month-old infants. *Developmental Psychology*, *32*, 942–950.

Ryalls, B., Winslow, E., & Smith, L. B. (in press). A semantic congruity effect in children's acquisition of high and low. *Journal of Memory and Language*.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations so systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, *16*, 417–494.

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 analogy* (pp. 146–178). Cambridge: Cambridge University Press.

Smith, L. B. (1995). Self-organizing processes in learning to learn words: Development is not induction. In *Basic and applied perspectives on learning, cognition, and development* (Vol. 28, pp. 1–32). Marwah, NJ: Lawrence Erlbaum Associates.

Smith, L. B., Cooney, N. J., & McCord, C. (1986). What is "high"? development of implicit reference points for "high" and "low". *Child Development*, *57*, 583–602.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, *46*, 159–216.

Sporns, O., Gally, J. A., Reeke, G. N., & Edelman, G. M. (1989). Reentrant signaling among simulated neuronal groups leads to coherency in their oscillatory activity. *Proceedings of the National Academy of Sciences*, *86*, 7265–7269.

Tardif, T. (1996). Nouns are not always learned before verbs: Evidence from Mandarin speakers' early vocabularies. *Developmental Psychology*, *33*, 492–504.

Titzer, B., Thelen, E., & Smith, L. B. (1998). *The developmental dynamics of learning about transparency: a new interpretation of the object retrieval task and the visual cliff.* (under review)

Tomasello, M., Akhtar, N., Dodson, K., & Rekau, L. (1997). Differential productivity in young children's use of nouns and verbs. *Journal of Child Language*, *24*, 373–387.

Waxman, S. R. (1994). The development of an appreciation of specific linkages between linguistic and conceptual organization. *Lingua*, *92*, 229–250.