3 Initial knowledge and conceptual change: space and number

Elizabeth S. Spelke and Sanna Tsivkin

Massachusetts Institute of Technology

How do humans build the rich and intricate systems of knowledge that are characteristic of our species? How variable are these knowledge systems across human cultures, and what are their universal properties? What accounts for the flexibility, adaptability, and open-endedness of human knowledge systems on the one hand, and the ease of acquisition of some systems on the other? Finally, what differences between humans and other animals, even our closest primate relatives, lead only humans to develop highly elaborated knowledge systems?

Traditional answers to these questions are incomplete at best. On one view, humans are endowed with a powerful capacity for learning, shaping new concepts and beliefs to fit the environment. This view might account for the flexibility and adaptability of human knowledge systems but cannot readily explain why humans develop knowledge rapidly in some domains but slowly, with great difficulty, in others (see Chomsky 1975, and Keil 1981). On a second view, humans are endowed with domain-specific, core cognitive systems around which elaborated knowledge grows. This view might account for the ease of acquisition of certain knowledge systems but not for humans' ability to develop systematic knowledge in genuinely new domains or to change conceptions in radical ways (see Carey 1985; Hatfield 1990).

In this chapter, we explore a third answer to our opening questions. Humans, we suggest, are endowed with a set of core systems of knowledge, but the systems have critical limitations. Initial knowledge systems are domain-specific (each applies to a subset of the entities that infants perceive and act upon), task-specific (each serves to solve a limited set of problems), informationally encapsulated (each operates on only part of the information that perceivers detect and remember), autonomous (one system cannot change its operation to accord with the states of other systems), and isolated (each system gives rise to distinct representations of the environment). Because these properties are hallmarks of modular cognitive systems, humans' initial systems of knowledge are, to a first approximation, modules (see Fodor 1983[1]; Karmiloff-Smith 1992).

Humans overcome some of the limits of their initial knowledge systems, we suggest, by conjoining the separate representations that those systems deliver to create new concepts of greater scope and power. Whereas the initial systems of knowledge may underlie children's rapid learning in specific domains, the processes for conjoining domain-specific representations may underlie humans' ability to extend their knowledge into novel territory. Many core knowledge systems may be found in other animals, moreover, but the ability to conjoin representations rapidly and flexibly so as to yield new representations may be uniquely human. This ability may account in large part for the richness and diversity of mature human belief systems.

Although the ability to conjoin distinct representations remains obscure, the research to be described suggests that it involves language. Language might serve as a medium for conceptual change because of two of its central features. First, a natural language allows the expression of thoughts in any area of knowledge. Natural languages therefore provide a domain-general medium in which separate, domain-specific representations can be brought together. Second, a natural language is a combinatorial system, allowing distinct concepts to be juxtaposed and conjoined. Once children have mapped representations in different domains to expressions of their language, therefore, they can combine those representations. Through these combinations, language allows the expression of new concepts: concepts whose elements were present in the prelinguistic child's knowledge systems but whose conjunction was not expressible, because of the isolation of these systems. *Pace* Fodor (1975), children who learn a natural language may gain a more powerful system of representation than any they possessed before.

We will not argue here that this picture of cognitive development is true; indeed, experiments have hardly begun to test it. Rather, we hope to show that the picture is plausible, and that Fodor's (1975) compelling arguments for the impossibility of genuine conceptual enrichment through learning, particularly language learning, deserve another look. We discuss these issues in two concrete cases, describing studies of developing representations of space and developing representations of number. Finally, we consider the implications of this view for questions of linguistic relativity: might speakers of different languages think incommensurable thoughts?

1 Spatial representation

The ideas explored in this chapter were suggested by the thinking and research of Linda Hermer-Vazquez (Hermer 1994; Hermer & Spelke 1994, 1996; Hermer-Vazquez, Spelke, & Katsnelson 1999). Hermer-Vazquez's research began with a question that seems far removed from the study of language or cognitive development, concerning the spatial abilities of humans and other animals. Comparative studies of navigation and spatial localization present a striking puzzle. Research in behavioral ecology, experimental psychology, and cognitive neuroscience provides evidence that a wide variety of animal species, from insects to mammals, maintain an exquisitely precise sense of their own position in relation to significant places in the environment. Animals update their spatial representations as they move around, and they draw on those representations in navigating through the layout, reorienting themselves, and locating objects (see Gallistel 1990, and McNaughton, Knierim, & Wilson 1994, for reviews and discussion). Evidence for spatial representations is so ubiquitous in animals that Gallistel captures this evidence with striking simplicity: "There is no creature so lowly that it does not know, at all times, where it is."

In contrast, even casual observation suggests that one species is an exception to Gallistel's rule: *Homo sapiens*. Many people living in modern, technological societies appear to retain very little sense of their position or orientation, or of the egocentric directions of significant objects and places, as they move. Perhaps as a consequence, people often navigate on strikingly inefficient paths, even through familiar environments. To be sure, some people do remain aware of their egocentric orientation, but this fact raises a further question: why are human spatial abilities so variable, compared to those of other species?

Hermer-Vazquez's initial approach to these questions was based on the hunch that the unique weaknesses of human spatial representations would be counterbalanced by unique strengths. In particular, perhaps the inaccuracies and errors of human navigators are compensated by their flexibility. To explore this possibility, her research focused on a situation in which other mammals have been found to form and use spatial representations *inflexibly*: when they are disoriented and must call on memories of their surroundings in order to reorient themselves.

Cheng (1986) and Margules & Gallistel (1988) investigated rats' reorientation abilities by exposing hungry rats to partially buried food in a testing chamber, removing the rats and disorienting them, and then returning them to the chamber where the food was now fully buried. The investigators assumed that on first exposure to the food, the rats would record its geocentric position: for example, a rat might represent the food as buried in the northeast corner of the chamber. In order for a disoriented rat to retrieve the food on its return to the chamber, therefore, it first had to reorient itself: in this example, it would determine its current heading and then compute the egocentric direction of Northeast. A wealth of information was provided to specify the rats' geocentric orientation, including the presence of distinctive odors and patterns in different corners and, in some studies, the relative brightness of different walls. Rats' search patterns suggested, however, that they used only one property of the chamber to reorient themselves: its shape. Because the chamber was rectangular, its shape specified the rats' orientation up to a 180° ambiguity. Disoriented rats betrayed their strong reliance on geometric information by searching for the food with high frequency at its true location and at the geometrically equivalent, but featurally quite different, opposite location. Indeed, rats that were fully disoriented searched these two locations with equal frequency, despite the wealth of nongeometric information that distinguished the locations (Margules & Gallistel 1988).

In further studies, Cheng and Gallistel showed that rats' failure to distinguish between featurally distinct but geometrically equivalent locations was not attributable to a failure to attend to or remember the room's nongeometric properties (see Cheng 1986, and Gallistel 1990, for discussion). Rather, rats appeared specifically unable to use their memory for nongeometric properties of the room in order to reorient themselves. These findings led the investigators to conclude that the rat's reorientation process was task-specific and informationally encapsulated: "a geometric module" (Cheng 1986).

In contrast to rats, humans' spatial representations appear to be more flexible: a disoriented person may use a wealth of nongeometric information to determine where she is. Emerging from a subway, for example, a person may determine her heading by searching for the names of streets or the numbers on buildings, by looking for shops or other landmarks, or even by asking directions. These evident abilities to use nongeometric information might testify to just the sort of flexibility that distinguishes humans from other animals.

To address this possibility, Hermer-Vazquez adapted Cheng's task for use with human children and adults, and she focused on the reorientation processes of very young children. In her first experiment (Hermer & Spelke 1994), 18- to 24-month-old toddlers were introduced into a rectangular room that was either entirely white or had one nongeometric directional cue: a bright blue wall. After a toy was hidden in one of the four corners of the room, the children's eyes were covered, they were lifted and rotated several turns, and then they were placed on the floor and urged to find the toy. If children, like rats, can reorient by the shape of the room, then they should have searched the two geometrically appropriate corners of the white room more than the other corners. Moreover, if children can reorient by the distinctive coloring of one wall, they should have searched the correct corner of the room with the blue wall more than any other corner.

The findings of this experiment were both clear and surprising. First, children searched the two geometrically appropriate corners of the entirely

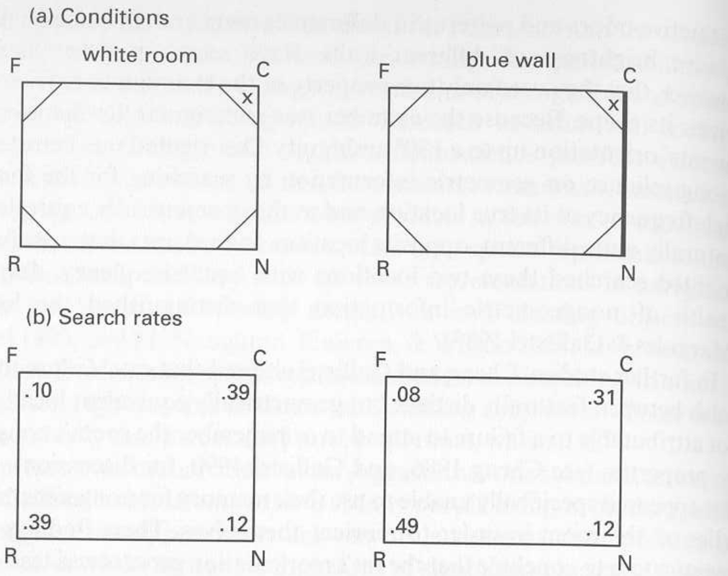

Fig. 3.1 (a) Overhead view of the chamber used in two conditions of a reorientation experiment. Although the object (indicated by x) was hidden in different corners for different subjects, the corner locations are rotated in this figure so that the correct search location is always depicted as the northeast corner (C) and the rotationally equivalent opposite location is always depicted as the southwest corner (R). The incorrect locations near and far from the correct corner are indicated by N and F. (b) The proportion of young children's search in corners C, R, N, and F in each condition of the experiment (after Hermer & Spelke 1996).

white room with high and equal frequency. This search pattern indicates that the children truly were disoriented, could not use any subtle, uncontrolled cues to locate the object, and reoriented in accord with the shape of the environment. Second, children utterly failed to use the blue wall to reorient themselves. In the room with one blue wall, children continued to search in geometrically appropriate locations, but they searched equally at the corners with appropriate vs. inappropriate coloring. Both qualitatively and quantitatively, children's performance closely resembled that of rats.

Because children's failure to use such a large and (to adults) salient landmark begged for explanation, subsequent analyses and experiments tested for a variety of possible sources of this failure. First, might children have searched only the corners that were visible when a trial began? A comparison of children's performance on trials which began with the correct corner in view vs. out of view ruled out this possibility, for children were almost

equally likely to search initially visible and invisible corners.[2] Second, might children have failed to notice the blue wall? In a follow-up experiment (Hermer & Spelke 1994), either the experimenter pointed to the blue and white short walls of the room until the child looked at them before the object was hidden, or the experimenter and child played with the blue fabric before the test began and together placed it on the wall. Neither manipulation affected search performance, which closely resembled that of the first study.

Third, would children reorient by the distinctive color of a wall if the geometric information for reorientation were reduced? To address this question, children were tested in a square room that contained one very bright, red satin wall (Wang, Hermer-Vazquez, & Spelke 1999). Interestingly, children's performance differed in the square room in one respect: this was the only experiment in which children tended to search in a constant egocentric direction relative to their facing position at the end of the disorientation procedure. This finding suggested that disorientated children did not attempt to reorient themselves in this geometrically impoverished environment and instead relied on an egocentric strategy for finding the object. Nevertheless, the use of a square room did not enhance children's ability to reorient themselves in accord with a nongeometric landmark. Search was no more accurate in the room with the red satin wall than in the room that was entirely white, even though the satin wall drew children's attention quite strongly at the start of the search session.

Since these findings suggest that children's reorientation process is quite impervious to wall coloring, further experiments investigated whether the reorientation process could take account of information specifying the categorical identity and properties of objects. In one study (Hermer, unpublished), a large multicolored plastic statue of a person was placed directly against one of the short walls in the rectangular room, and a blue trash can of similar global proportions was placed against the opposite wall. Children reliably confined their search to the two geometrically appropriate corners, indicating that they were sensitive to the lengths of the two short walls. Nevertheless, they failed to distinguish the correct from the opposite corner, suggesting that their reorientation process was insensitive to the identities of the objects at the center of those walls (see also Hermer & Spelke 1994: Exp. 3).

In the next study, the color and patterning of the object's hiding location served as the nongeometric information for reorientation. Disoriented children searched for an object that was hidden inside one of two containers with distinctive coloring and patterning but identical shapes, placed in geometrically indistinguishable, opposite corners of the rectangular room. Although children searched for the object by passing their hand directly into the distinctive container, they searched the correct and incorrect containers with equal frequency. Nongeometric information again failed to serve as a basis for reorientation (Hermer & Spelke 1994).

Why was children's reorientation impervious to nongeometric information? One possible explanation roots children's difficulty in a general failure to perceive, attend to, or remember nongeometric information: perhaps children are inattentive to colors, textures, and patterns in an enclosed environment, fail to retain this information during a disorientation procedure, or fail to access this information during an object search task. These possibilities were tested by allowing children to watch as a toy was hidden in one of two distinctive containers in the rectangular room, disorienting the children as in the previous experiments, moving both the children and the containers into a larger, geometrically distinctive space, and encouraging children to find the toy. This search required that children use the nongeometric properties of the containers to track the location of a displaced object but not to reorient themselves. For the first time, children succeeded in searching the box with the appropriate nongeometric properties, suggesting that children's previous failures to reorient by these properties did not stem from limits on attention, memory, or the perceptual guidance of action (Hermer & Spelke 1996; see also Hermer & Spelke 1994, Exp. 4).

A final experiment tested directly whether children, like rats, reorient by virtue of a task-specific, encapsulated process (Hermer & Spelke 1996, fig. 3.2). Children first watched as a toy was hidden in one of two distinctively colored and patterned containers, placed in adjacent corners of the room, and then their eyes were closed and the containers were quietly moved to the opposite two corners of the room such that their geometric and nongeometric properties were dissociated (i.e., if pink and green boxes originally appeared to the left and right of a short wall, respectively, the boxes subsequently appeared on the opposite wall in reversed left/right relations). In one condition, children were disoriented while the containers were moved; in the other condition, they were not disoriented. Note that in both conditions, children saw exactly the same environment and events throughout the study and were asked to engage in the same actions. Nevertheless, the tasks faced by children in the two conditions were deeply different. The disoriented children needed to relocate *themselves*, so that they could determine the egocentric direction of the hidden object. In contrast, the oriented children knew where they were and therefore could determine, by computing the hidden object's expected egocentric direction and encountering an empty corner, that the object had moved. Their task was to relocate the *object*.

Search patterns were quite different in the two conditions. Whereas the disoriented children searched primarily the container with the appropriate geometry, the oriented children searched primarily the container with the

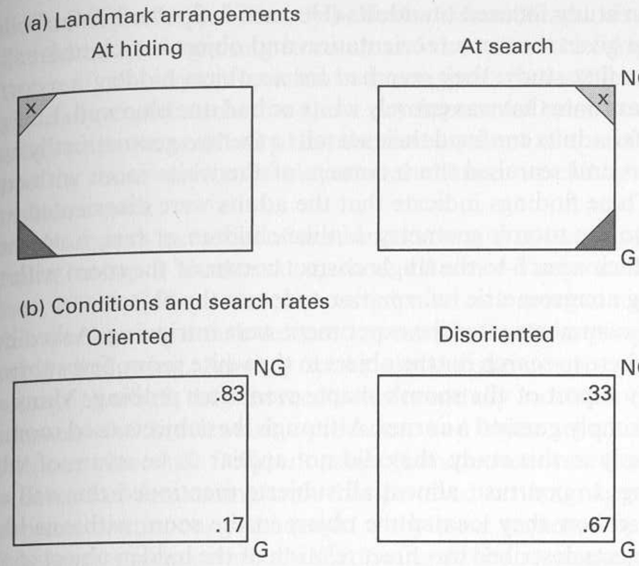

Fig. 3.2 (a) Overhead view of the chamber at the beginning (left) and end (right) of each search trial. For half the children, the geometrically correct location (G) was in the southeast corner as depicted; for the other children, the locations of that corner and of the nongeometrically correct corner (NG) were reversed. (b) The proportion of children's search in corners G and NG under conditions of orientation or disorientation (after Hermer & Spelke 1996).

appropriate color and markings. Importantly, the differing patterns were observed on the first search trial, before children could know what task they would face and what information they would need to remember. These findings provide evidence that young children perceive and remember both geometric and nongeometric properties of the environment, and that a task-specific, encapsulated reorientation process makes use only of a subset of these properties. Like rats, young children appear to reorient by a modular system sensitive only to geometry.

2 Developmental changes in spatial representation

Hermer-Vazquez's studies suggest a close correspondence between the reorientation systems of rats and humans, but they bring us no closer to answering the question with which we began: why are many people so bad at maintaining their orientation and navigating on efficient paths? To address this question, Hermer turned to studies of developmental changes in reorientation.

Her first study focused on adults (Hermer & Spelke 1994). College students were given the same reorientation and object search task as the children in the first study: they searched for an object hidden in a corner of a rectangular room that was entirely white or had one blue wall. Like children and rats, the adults confined their search to the two geometrically appropriate corners and searched those corners of the white room with equal frequency. These findings indicate that the adults were disoriented and were sensitive to the room's geometry. Unlike children or rats, however, adults confined their search to the single correct corner of the room with the blue wall, using nongeometric information to locate the object.

Adults' comments after the experiment were intriguing. Asked how they decided where to search for the object in the white room, few subjects mentioned any aspect of the room's shape, even after probing. Many subjects said they simply guessed a corner. Although the subjects used room geometry flawlessly in this study, they did not appear to be aware of what they were doing. In contrast, almost all subjects mentioned the wall coloring when asked how they located the object in the room with one blue wall. Many subjects described the direct relation of the hidden object to the wall: e.g. "I saw you hide it left of the blue wall."

These comments suggested that adults might have used the wall as a landmark for locating the object, rather than as a cue for reorienting themselves. A developmental experiment supported this suggestion. Three- to seven-year-old children were tested in the rectangular room with one blue wall (Hermer 1997). In one condition (an indirect task), the toy was hidden in a corner as in the previous studies: e.g., to the left of the blue wall. In the other condition (a direct task), the toy was hidden directly behind the center of a short wall: e.g., behind the blue wall. At six–seven years of age, children successfully located the object in both conditions. At three–four years of age, in contrast, children succeeded at the direct hiding task but not the indirect task. Because the blue wall provided the same information about the child's orientation in the two conditions, the younger children's performance suggests that they were not using the blue wall to reorient themselves but rather to specify the position of the hidden object. Such a specification would allow children and adults to locate the object while remaining in a state of disorientation.

Hermer's studies suggested that developmental changes in object localization were roughly correlated with changes in spatial language. Recall that adults both encoded and readily described the object's relation to the blue wall. In addition, young children's ability to encode the object's relation to the blue wall approximately coincided with the development of relevant spatial expressions: three- to four-year-old children command a variety of expressions, such as *in X* or *at X*, that could serve to specify the

location of an object hidden directly behind a blue wall, and, six- to seven-year-old children are beginning to command the relevant left/right terminology to specify the object's location in the original studies (Hermer 1997). Although many explanations for these developmental relations could be offered, they raise the possibility that the development of spatial language contributes in some way to developmental changes in children's performance on these tasks.

Hermer-Vazquez's last series of experiments begins to test this possibility with adults, using a dual-task method (Hermer-Vazquez, Spelke, & Katsnelson 1999). College students participated in a reorientation experiment while performing a task designed to interfere with language production: verbal shadowing. As they spoke aloud continuously, repeating a prose passage to which they listened over a centrally located loudspeaker, they underwent the disorientation procedure and object search task in the room with one blue wall. Different subjects, moreover, were tested with a different attention-demanding task that involved no language: a task in which subjects continuously shadowed a rhythmic sequence by clapping. The results were striking: verbally shadowing adults reoriented in accord with the shape of the room but not its coloring, searching equally at the two geometrically appropriate corners. In contrast, adults who engaged in the rhythmic shadowing task or in no shadowing located the object successfully at the single, correct corner. A task that interfered with language production also interfered with the adults' ability to localize the object in relation to the blue wall.

Hermer-Vazquez's findings encourage us to speculate how the development of language might lead to developmental changes in her tasks. Before the development of spatial language, children evidently form both representations of the geometry of the stable environmental layout (used for reorientation) and representations of the nongeometric properties of objects and surfaces in the layout (used for finding displaced, hidden objects). Because these representations are constructed by autonomous and encapsulated systems, they cannot be conjoined. The inability to conjoin geometric with nongeometric descriptions of the environment precludes children's representing an object's location as being in a certain geometric relation to a certain nongeometric environmental property: e.g., left of a blue wall. Nevertheless, children's modular systems permit them to represent all the ingredients of relations such as "left of blue". The geometric system preserves information about sense relations, allowing children to differentiate between corners of the room which differ only with respect to the left/right relation of the short and long walls. Other systems preserve information about nongeometric properties of the environment, allowing children to confine their search for a displaced object to a container with appropriate coloring and patterning.

Because both geometric and nongeometric relations are represented prelinguistically, children could learn to use language to conjoin these relations in either of two ways. First, children might learn terms such as *left* and *near* by mapping the terms directly to representations constructed by the geometric system, and they might learn terms such as *blue* or *toy* by mapping the terms to appropriate nongeometric representations. Once these mappings are learned, the domain-general, combinatorial properties of language would allow the child to interpret expressions such as *left of the blue wall* or *near the toy*.

A second way of using language to conjoin information from distinct, modular systems assumes that complex spatial expressions conjoining geometric and nongeometric information are learned as wholes. In this case, terms such as *left* and *near* would derive their meanings not from mappings to a single system of representation, but from simultaneous mappings to several distinct systems. For example, children might first learn the meanings of expressions such as *your left hand* or *the picture near the window*. Learning these expressions would require the simultaneous activation of (a) representations of token objects (i.e. "there's an object x and an object y"), (b) nongeometric representations of each object (e.g. "x is a window," "y is a picture"), and (c) geometric representations of the relation between the objects ("y is near x"). If spatial terms are learned only in the context of expressions that require multiple, simultaneously active representations for their satisfaction, each term will connect to both geometric and nongeometric representations and therefore will link these representations to one another.[3]

Both of these processes exploit two central properties of language. First, language is a domain-general system of representation, containing terms that refer to objects and relations whose primary representations are constructed by a diverse collection of modular systems. Second, language is a combinatorial system, allowing terms to be conjoined irrespective of their (domain-specific) content. Language therefore can expand the range of a child's concepts by conjoining terms that map to elements in distinct, nonverbal representations.

Returning to the case at hand, we suggest that the acquisition of spatial language allows the child to represent the position of a hidden object in new ways. The use of spatial language also might underlie the marked differences between the spatial behavior of humans and other species. People, we have noted, represent space more flexibly than other animals, capturing properties of the environment with words and maps (which conjoin both information about the shape of the environment and nongeometric information such as the names of streets and other landmarks and the nature of the terrain). Language may provide an important medium in which such information is organized, allowing people to use a wide range of representational resources to encode and remember routes through the environment and the location of objects and places.

On the negative side, western adults often make their way through environments with little sense of their geocentric orientation, traveling on inefficient paths. These limitations may stem from two properties of the way language represents space. First, languages such as English represent space independently of the geocentric positions of the self or objects. As a consequence, speakers of these languages do not need to maintain a sense of geocentric direction in order to talk (see Levinson 1996a).[4] Second, all natural languages appear to represent spatial relations crudely, with terms that capture schematic, categorical relations among objects irrespective of their metrical structure (Talmy 1983; Landau & Jackendoff 1993). Linguistic descriptions of the layout therefore lack the precision of nonverbal, metric representations.

Finally, the use of spatial language could account, in part, for individual differences in spatial performance within a single language community. Because language in principle allows for a multiplicity of conjunctions of information, it allows for a variety of representations of the environment. People may differ both in their degree of dependence on language-based conjunctions, and in the specific conjunctions that they use to guide their actions. If this possibility is correct, then we would expect substantial individual variation in spatial performance that relies on conjunctions of information and less variation in spatial performance that relies on purely geometric information. This prediction has not been tested.

3 Number

In the domain of space, uniquely human representational abilities appear to entail costs in precision and accuracy but yield gains in representational power. In other domains, the ability to conjoin distinct representations may lead both to gains in representational power and to increased precision. Here we consider one domain where both advances may occur.

Starting again from a comparative perspective, research on a wide variety of animals suggests that representations of numerosity are ubiquitous among vertebrates. Fish, birds, and mammals respond systematically to the rate at which food is provisioned in natural settings, and pigeons and rats have been trained in laboratory experiments to respond to the number of events (light flashes, sounds) in a sequence or the number of actions they have performed (for reviews, see Gallistel 1990; Boysen & Capaldi 1993; and Dehaene 1997). There is no obvious upper bound to the size of sets that animals represent. Nevertheless, animals' number representations are imprecise and their accuracy declines with increases in set size, in accord with Weber's law.

A number of mammals also have been shown to represent exact numerosity and to take account of effects of simple addition and subtraction. For example, a raccoon was trained to respond positively to a box displaying exactly three objects (Davis 1984), and both a parrot and a chimpanzee have been trained to give unique responses to sets of objects of numerosities varying from one to six or more (Matsuzawa 1985; Pepperberg 1987). Most interestingly, untrained monkeys and tamarins have been shown to compute exact additions and subtractions on small sets: if two objects are placed in succession behind an occluder, the subjects look longer if the screen is raised to reveal one or three objects than if it reveals two objects, suggesting that the monkeys anticipated seeing the correct number (Hauser, MacNeilage, & Ware 1996; Hauser, Carey, & Hauser 2000). For these representations, however, there appears to be a limit on set size. As set size increases, learning labels for a set of a given numerosity requires more and more training. Moreover, number encoding does not appear to be a readily accessible process for animals. If a chimpanzee has learned to respond with a distinct Arabic symbol for each of the set sizes 1–4 and then a new symbol (*5*) and a new set size (five items) is introduced, the chimp appears to infer that *5* means "any set size that isn't 1–4," rather than one specific numerosity (Matsuzawa 1985). As sets get larger, the encoding of exact number therefore appears to become a last resort for many animals (see Davis & Perusse 1988).

Many experiments provide evidence that prelinguistic human infants represent number as well. Newborn infants discriminate between small sets of dots varying in number (Antell & Keating 1983), as do older infants (Starkey & Cooper 1980; Strauss & Curtis 1981; Starkey, Spelke, & Gelman 1990). Infants respond to the invariant number of elements in a set over changes in spatial patterning (e.g. Starkey & Cooper 1980), over motion and occlusion (van Loosbroek & Smitsman 1990), and over changes in a variety of properties of the elements in a set including size, shape, coloring, and categorical identity (Strauss & Curtis 1981; Starkey *et al.* 1990). Finally, infants pass the same addition and subtraction tasks as monkeys and tamarins (Wynn 1992a; Hauser *et al.* 2000), and they are sensitive to the correctness or incorrectness of simple addition even if the elements move (Koechlin, Dehaene, & Mehler 1996) or change properties (Simon, Hespos, & Rochat 1995). Evidence for exact, small-number representations also has been obtained with adults, who apprehend the exact numerosity of small sets rapidly and track small numbers of objects in parallel (Trick & Pylyshyn 1993). All these findings suggest that the capacity for representing the exact numerosity of small sets is common to humans and other animals and emerges early in human development.

What of the capacity to represent larger sets? Although early studies suggested that infants discriminate between large sets of elements exhibiting large differences in numerosity (Fantz, Fagan, & Miranda 1975), a later reanalysis of this work suggested that infants' discriminations depended on the spatial distribution of contrast in the displays (Banks & Ginsburg 1983). Nevertheless, there is ample evidence for approximate representations of large sets in adults, who can rapidly determine which of two large sets is more numerous if the sets are sufficiently different in number, and who can rapidly estimate the approximate answers to certain large-number arithmetic problems (see Dehaene 1997; Gallistel & Gelman 1992). There is also some evidence for these representations in infants (Xu & Spelke, 2000) and toddlers (Sophian, Harley, & Martin 1995). In particular, 6-month-old infants discriminated between sets of 8 vs. 16 visual elements, even when controls within the experiment insured that discrimination could not be based on the sizes, spacing, or arrangement of the elements, or on continuous variables such as the spatial extent that the set of elements occupies. In contrast, infants failed to discriminate sets with a smaller difference ratio: sets of 4 vs. 6 or 8 vs. 12 elements (Starkey & Cooper 1980; Xu & Spelke 2000). The representations of the approximate numerosity of large sets found in a variety of animals appear to exist in young children.

These findings suggest that both animals and preverbal infants have two systems for representing number. The first system serves to represent small numerosities exactly. It underlies animals' and infants' abilities to keep track of up to four objects in number discrimination and addition/subtraction tasks, and adults' rapid apprehension of small numbers of objects (see also Kahneman, Treisman, & Gibbs 1992; Trick & Pylyshyn 1993). These representations appear to be robust over variations in the properties of their elements such as shape and location. The second system serves to represent large sets. It also persists over development and underlies animals', infants', and adults' rapid apprehensions of the relative numerosity of large sets of entities, and underlies adults' abilities to give approximate answers to large-number arithmetic problems (see Dahaene 1997). These representations do not appear to be limited as to set size beyond the limits of sensory acuity, but their accuracy decreases with increasing set size in accord with Weber's Law.

In contrast to infants and to other animals, human adults have a third system for representing numbers, which typically involves verbal counting. Like the small-number system, this system allows the representation of exact numerosity, independently of other quantitative variables, and the computation of the exact effects of addition and subtraction. Like the large-number system, it has no upper bound on set size, beyond that imposed by limits on time and patience, and it allows comparisons of the relative numerosity of two sets.

It is widely believed that children develop the ability to use this new, uniquely human system of number representation when they learn verbal counting (e.g. Gallistel & Gelman 1992). But how does verbal counting develop, and how does its development give rise to a new system of number representation? Number is arguably our most abstract system of knowledge – why is this system tied to the child's developing understanding of a verbal activity like counting? How can children ever come to understand counting if they do not already understand the entities that counting singles out? Moreover, how does learning to count allow children to form representations, such as "exactly seven," that exceed the limits of their preverbal systems of number representation, circumventing Fodor's (1975) dictum that one can learn words only for concepts that one can already represent?

The most ambitious attempt to answer these questions may be found in Gelman & Gallistel's thesis that the system of knowledge underlying verbal counting is innate and inherent in the large-number representational system, and that children learn to count by gaining access to the principles that define the nonverbal system (Gelman & Gallistel 1978; Gallistel & Gelman 1992; see also Wynn 1990; Sophian, Harley, & Martin 1995, and Dehaene 1997 for further discussion of Gelman & Gallistel's thesis). A different account (Bloom 1994) builds on Chomsky's (1986) thesis that the conception of discrete infinity that grounds knowledge of the natural numbers stems from the generativity of language: children learn to count by gaining access to the grammatical principles giving rise to this generativity. One problem with both these accounts, however, is that it is not clear how the *principles* governing the operation of any modular representational system become accessible to other systems. In many cases, such access plainly does not occur: speaking a language does not give one knowledge of linguistics, and experiencing brightness contrast does not give one knowledge of calculus (Fodor 1983; Gallistel 1990; cf. Rozin 1976). Why and how might the principles underlying nonverbal counting or language generativity become accessible to the child?

The foregoing analysis of spatial orientation prompts a different account of number development. Children may attain the mature system of knowledge of the natural numbers by conjoining together the representations delivered by their two preverbal systems. Language may serve as a medium for this conjunction, moreover, because it is a domain-general, combinatorial system to which the representations delivered by the child's two nonverbal systems can be mapped.

More specifically, consider the precounting child's representation of "two." By our hypothesis, the child has two systems for representing arrays containing two objects: a small-number system of object representation (which represents this array, roughly, as "an object x and an object y, such that y ≠ x"; see Wynn 1992a) and a large-number system of representation (which, as in Gelman & Gallistel's account, represents this array as "a blur on the number line indicating a very small set"). Because of the modularity of initial knowledge systems, these representations are independent. When younger children hear the word *two*, therefore, they have two distinct representations to which the word could map and no expectation that the word will map to both of them.

Because all the number words appear in the same syntactic contexts (see Bloom & Wynn 1997) and occur together in the counting routine, experience with the ambient language may lead children to seek a common representational system for these terms. Thus, children may discover that all the terms map to representations constructed by the large-number system (although specific terms do not map to specific large-number representations, because of the inaccuracy of that system). This possibility is consistent with Wynn's (1992b) evidence that young children pass through a stage when they interpret all number words above *one* as referring to any display containing more than one object. In addition, because the terms *one*, *two*, and *three* map to specific representations delivered by the small-number system, children may discover these mappings. This possibility is consistent with evidence that young children typically learn the meanings of *one*, *two*, and *three* individually, in that order (Wynn 1992b).

With these advances, how do children learn the meanings of words such as *seven*? Unlike *two*, *seven* does not map to any representation in either the small-number system (which can only represent arrays of seven objects as "arrays of too many things to keep track of") or the large-number system (whose inaccuracy will often lead the child to confuse arrays of seven objects with arrays of six or eight objects). Nevertheless, learning the meanings of *one*, *two*, and *three* may provide the seeds of a solution to the problem of learning *seven*. First, because the words for small numbers map to representations in both the small-number system and the large-number system, learning these words may indicate to the child that these two sets of representations pick out a common set of entities, whose properties are the union of those picked out by each system alone. This union of properties may be sufficient to define the set of natural numbers. Second, because the words *two* and *three* are members of a larger set of terms that behave alike in the different linguistic contexts, learning the terms *two* and *three* may open the way to the insight that terms such as *seven* also refer to sets of objects with the same union of properties. Third, because the terms *one*, *two* and *three* form a sequence in the counting routine, children may discover that each of these number words picks out a set with one more individual than the previous word in the sequence, and they may generalize this learning to all the words in the counting sequence.

In summary, knowledge of the natural numbers may be constructed through the conjunction of representations delivered by two distinct, modular systems, with language providing the medium for this conjunction. Once the child learns that a single set of words maps to both these systems, she may gain the ability to conjoin information that the two representational systems deliver, capturing the strengths of both systems and overcoming limitations specific to one or the other system. From the small-number system may come the realization that each number word corresponds to an exact number of objects, that adding or subtracting exactly one object changes number, and that changing the shape or spatial distribution of objects does not change number. From the large-number system may come the realization that these sets of exact numerosity can increase without limit, and that a given symbol represents the set as a unit, not just as an array of distinct objects. The union of these properties may allow not only the learning of words such as *seven* but also the representation of concepts such as "seven", and of beliefs such as "seven plus seven is fourteen." In the domain of number as in the domain of space, the acquisition of language may lead to the development of concepts that were not expressible in the cognitive systems of the preverbal child.

The most direct tests of this possibility would focus on young children and probe for developmental changes in number representations that are consequences of the child's developing understanding of counting. Unfortunately, such studies have not been conducted. Nevertheless, studies of adults with arithmetic impairment and of adults who speak two languages offer suggestive support for this view.

First, studies in cognitive neuropsychology have focused on the number representations and arithmetic abilities of patients with disorders in arithmetic processing, or "acalculia" (see Dehaene 1997 for review). Interestingly, these patients tend to have both impaired language and impaired abilities to give the exact answers to arithmetic problems. In contrast, many are able to provide an approximate answer to a mathematical problem. For example, Warrington (1982) studied a patient who correctly supplied approximate answers to addition problems with fairly large numbers but often was unable to state the exact answer. The patient himself said that he felt that he knew the approximate, but not exact, answers to all the mathematical problems presented. Similarly, Dehaene & Cohen (1991) discussed the case of an aphasic and acalculic patient who failed to identify the errors in mathematical statements such as $7+3=11$ but was able to reject statements such as

$7+3=17$. These findings suggested that the patient's representations of approximate numerosity were intact, but exact calculations were impaired.

Further suggestive evidence in favor of the thesis that language underlies the representation of exact large numerosities comes from research on bilingualism. There are many anecdotal reports of bilinguals who perform arithmetic in their first language, even if that language is hardly used for other purposes: a person who has learned to do arithmetic in one language (L1) during childhood and then moves to a different country and becomes dominant in that country's language (L2) is apt to perform arithmetic calculations in L1.

A number of studies have bolstered these anecdotal reports by showing that bilingual subjects are faster at verifying or producing the correct answer to arithmetic problems in L1 than in L2 (Marsh & Maki 1976; McClain & Huang 1982; Frenck-Mestre & Vaid 1993). These findings do not show, however, that arithmetic is performed in a language-specific format because the L1 advantage may stem from processes that translate between an abstract language of numerical processing and a language of input and output (see McCloskey 1992). Bilinguals may habitually translate from and to L1, leading to their overpractice with L1 numerals and accounting for their faster performance in L1 (although see Gonzalez & Kolers 1982). To determine whether arithmetic truly requires language-specific representations, it would be desirable to conduct training experiments in which the exposure to numerals in L1 and L2 is controlled. If arithmetic calculations are language-specific, they should be performed faster in the language of training, even when practice with encoding and decoding is equalized across the two languages.

An additional limitation of existing bilingual studies is that they have failed to probe the limits of the L1 advantage. Are bilinguals faster and more accurate on all numerical problems when they perform in L1, or only on certain kinds of problems? The thesis that language serves as a medium for representing exact large numerosities leads to the prediction that language-specific effects will be obtained only for tasks requiring representations that are exact and large.

We have recently attempted to test these predictions (Dehaene, Spelke, Pinel, Stanescu, & Tsivkin 1999; Spelke & Tsivkin in press). In our first study, Russian–English bilinguals were trained to solve three sets of mathematical problems in each of their two languages. Half the subjects were trained in Russian on a set of tasks consisting of double-digit addition problems with addend 54, single-digit addition problems in base 6, and estimation of cube roots problems, and they were trained in English on double-digit addition problems with addend 63, single-digit addition problems in base 8, and estimation of logs base 2 problems. For the remaining subjects, the languages were reversed. After two days of training in each language, all the subjects were tested in two sessions, one administered in

Russian and one in English. In each session, the subjects were tested on all combinations of language and tasks, so that their performance on a given task could be compared across the two languages.

This design addresses a number of limitations of previous work on bilingual numerical processing. First, the introduction of training sessions allowed us to simulate the learning situation faced by children and to ensure that subjects spent an equal amount of time on each type of task across languages. Second, in order to test McCloskey's claim that language experience affects encoding and decoding processes but not the representation of numbers and arithmetic facts, we tested subjects on a subset of items requiring the same output across the two languages (e.g. $54+x=z$ in Russian and $63+y=z$ in English). Third, we assessed the language-specificity of training with both tasks calling for exact calculations and tasks calling for estimations.

The results were clear. First, there were improvements with practice across the two training sessions for all the tasks in this experiment. Training benefits of comparable magnitude were observed with both languages and all problem sets. Second, subjects' test performance in the language of training showed a clear advantage over their test performance in the untrained language for large-number addition and for addition in different bases. Language-specific training effects were observed both for the problems trained in L1 (Russian) and for those trained in L2. Importantly, there was no significant difference between performance on the items that shared the same output and those that did not, suggesting that training did not primarily influence the speed of encoding or decoding in a particular language. This finding suggests that there is a genuine language-specificity in the training effects with the tasks requiring exact calculation with large numbers.

The most important result from this study is that there was no advantage for the language of training on the estimation tasks. Subjects trained on cube roots in Russian, for example, were just as fast at estimating the cube roots in English. This finding indicates that language-specificity does not extend to all numerical processing tasks: only tasks requiring exact, large-number calculations appeared to be processed in a language-specific form. This result is especially striking in view of the design of the study. The experiment was set up to allow for identical conditions of training for all of the tasks involved: the number of training problems, repetitions, and other task factors were equated across all the tasks. The different patterns of transfer therefore suggest that there is something about exact, large-number arithmetic that involves language in ways that approximate arithmetic does not.

In a follow-up study, we extended this research in two directions

(Dehaene *et al.*, 1999; Spelke & Tsivkin in press). First, we investigated whether normal bilingual subjects, like acalculic patients, would show a dissociation between exact and approximate arithmetic within a single arithmetic operation. To this end, we trained separate groups of bilingual subjects on the same two sets of large-number addition problems, requiring one group to calculate the exact answer to each problem and the other group to calculate the approximate answer. If exact but not approximate addition requires a language-specific format, then only the subjects trained on the approximate answers should show a transfer of training to their second language.

Second, we attempted to address a conflict in the neuropsychological and animal literatures concerning the nonverbal estimation process. Gallistel (1990; Gallistel & Gelman 1992) has suggested that the large-number approximation system can be used to perform all arithmetic operations, including multiplication. In contrast, studies of at least one patient with impaired language and impaired exact calculation abilities but preserved approximation abilities suggest the patient can perform approximate addition but not approximate multiplication. Although the patient studied by Dehaene & Cohen (1991) correctly identified the errors in addition problems with answers that were distant from the correct answer (e.g. $7+3=17$), he was unable to identify errors in even the simplest multiplication problems (e.g. $3\times3=16$). In view of this conflict, we probed further for an approximate multiplication process by training the same two groups of subjects on a set of exact or approximate multiplication problems.

Each subject in this study was trained on two sets of problems, one in Russian and one in English. Subjects were administered one set of problems involving addition and one involving multiplication, with one task requiring exact calculation and one requiring approximate calculation. The pairings of languages, operations, and tasks (exact vs. approximate) were counterbalanced across the subjects.

Three findings are noteworthy. First, bilinguals showed no language-specific effects when trained on addition problems via estimation. This finding converges with those of Dehaene & Cohen (1991) and provides further evidence for a language-independent process for estimating approximate answers to addition problems. Second, we replicated our studies with bilinguals trained to solve addition problems via exact calculations. For the same items for which estimation-trained bilinguals failed to show any language-specificity, our exact calculation-trained bilinguals showed superior performance in the language of training.

Third, we found no evidence for approximate multiplication. Both the exact calculation-trained and the estimation-trained bilinguals were faster and more accurate when tested in the language of training, and they showed little transfer of training advantages across languages. This finding, like studies of patients with acalculia, suggests that adults can estimate large-number additions but not large-number multiplications in a language-independent manner.

In summary, there is a broad convergence between studies of number representation in children, brain-damaged patients, and bilingual adults. All these studies suggest that number representations and calculations are independent of language when the numbers involved are very small or answers required are imprecise. In contrast, number representations and arithmetic calculations appear to be language-specific when exact answers to large-number problems must be given. These findings are consistent with the thesis that human children and adults have one nonverbal system for representing small numerosities accurately and a second nonverbal system for representing large numerosities inaccurately. Language may serve to conjoin these two systems, allowing for the discovery of a third system of number representation that is both accurate and unbounded.[5]

4 Language, thought, and conceptual change

We have considered two domains in which children's conceptual resources appear to be enriched over development. In each case, we have suggested that this enrichment depends on a process of conjoining representations that were constructed by task-specific, encapsulated systems of knowledge, and that this process in turn depends on language. This puts us in the odd position of advocating a thesis that has been argued cogently to be absurd. Fodor's argument in *The language of thought* (1975) is intricate but its essence is simple: children can only learn the meaning of a term in a language if they can map the term to a preexisting expression in the language of thought. In that case, however, no language that children learn can go beyond the expressive power of the mental language they already possess.

The essential assumption that leads us to question Fodor's (1975) compelling argument comes again from Fodor, this time from *The modularity of mind* (1983). As every student of cognitive science knows, *The modularity of mind* develops the suggestion that an interesting subset of human cognitive systems are task-specific, informationally encapsulated, and autonomous. Recent research on early cognitive development provides considerable evidence, we believe, that the prelinguistic child's concepts and reasoning are subserved by cognitive systems with these central hallmarks of Fodor's modules. If that is true, however, then the infant might not have one language of thought, but many. If children have many languages of thought (or, as we prefer, systems of representation), then there may be thoughts that they cannot entertain, because the elements of which these thoughts are composed reside in different representational systems. These are the thoughts that a natural language (or some other domain-general, combinatorial system of representation) may make available.

We have considered two kinds of thoughts that may become expressible as children put together their domain-specific, modular representations: thoughts such as "left of the truck" and "five plus seven is twelve." In principle, the development of a domain-general, combinatorial system of representation could lead to a wide variety of new thoughts and concepts. If the present suggestions are correct, then human cognition would be limited only by three factors: the nature of the representations formed by initial, modular systems; the nature of the possible mappings from each of these systems to other domain-general representational systems such as a natural language; and the combinatorics of the latter representational systems.

Any domain-general, combinatorial system of representation in principle could serve to conjoin information from modular representations, but our evidence suggests that natural languages are the primary systems that play this role. The shadowing adults in Hermer-Vazquez's studies and the bilingual trainees in our studies did not appear to find a language-independent medium for representing that a toy was left of a blue wall or that 32 + 63 = exactly 95. Although language is not the only domain-general medium of representation that is available to humans, it does appear to provide a particularly powerful system for conjoining domain-specific representations into new concepts.

One research strategy for probing further the role of language in conceptual change is to investigate cognitive development and cognitive performance in speakers of different languages. In particular, the existence of languages lacking terms such as *left* and *right* (e.g. Levinson 1996a) raises the question of how speakers of these languages would perform in Hermer-Vazquez's reorientation tasks, and the existence of languages lacking counting words for all set sizes (e.g. Gordon 1994) raises the question of how children learning these languages develop number concepts and what concepts they develop.

Studies of linguistic variation also could address questions of linguistic relativity: if the acquisition of language allows the expression of new thoughts that conjoin information from domain-specific representations, do speakers of different languages develop different conjoint concepts? Note that, in principle, the present thesis leaves this question completely open. All conjoint, language-dependent concepts could turn out to be universal among speakers of any language, if (a) languages have strongly universal semantic properties, or (b) terms and expressions in a language can only map onto the representations delivered by the modular systems in a restricted set of ways. On the other hand, conjoint, language-dependent concepts could turn out to be highly variable, if the set of potential mappings across distinct languages is very large, relative to the number of mappings that any actual speaker realizes, and if different languages pick out different members of the set. A third, intermediate possibility now tempts us. It is possible that certain conjunctions of information across domains are especially natural for humans and therefore are found in most or all human cultures: the natural numbers are an example. Other conjunctions of information across domains may develop somewhat later and less universally and yet show only limited ranges of variation. Expressions like *left/east of the truck* or *in/loose-fitting against the bowl* may be such cases (see ch. 17 by Brown and ch. 16 by Bowerman & Choi, this volume). Still other conjunctions may be found later in development, be laborious to learn, and appear only in a small number of cultures. For example, when one uses natural numbers to represent the masses of objects, or the location of points in space, one opens the way to advances in physics and mathematics that may lead far beyond species-universal representations. The present thesis therefore does not presuppose linguistic relativity but leaves this possibility as a question for research.

As we have noted, the thesis that language serves as a medium for conjoining domain-specific systems of representation is only a suggestion, but we may close on firmer ground. First, studies of young children provide evidence that children's initial knowledge of space and number derives from cognitive systems that are domain-specific, task-specific, and encapsulated: cognitive modules. These initial knowledge systems appear to be shared by other animals: humans are not the only creatures that represent the shape of the layout, the identity and distinctness of objects, or the approximate numerosity of sets. Second, the studies provide evidence that human cognition extends beyond the limits of these modular knowledge systems, into domains of knowledge that only humans attain, and that this extension depends in part on processes that conjoin together the representations constructed by distinct knowledge systems.

The processes that underlie these conjunctions are still obscure and subject to debate. Nevertheless, the research efforts described in this volume suggest how one can begin to investigate these processes.

Developmental psychologists can study the emergence of these combinatorial processes through a systematic, two-pronged investigation, focusing both on the content and structure of the initial knowledge systems and on the changes that occur over cognitive development. Cognitive anthropologists can study the same processes through a complementary two-pronged investigation, focusing both on the content and structure of universal knowledge systems and on the variability in knowledge across different cultures. Through such studies, we may hope to shed light on what is perhaps the most intriguing aspect of human cognition: our ability to use our special-purpose systems of knowledge in order to extend our knowledge beyond those purposes, into novel territory.

NOTES

We thank Jerry Fodor, Lila Gleitman, and Steve Levinson for insightful comments on a draft of this chapter and Susan Carey and Linda Hermer-Vazquez for clarifying discussion of the issues it raises. Supported by NIH grant HD23103 to Elizabeth S. Spelke.

1 We depart from Fodor in arguing for the isolation of the modular cognitive systems. Fodor (1983) proposed that cognitive modules send their outputs to a single system of representation: the "language of thought" (see Fodor 1975). We propose that each module produces a separate set of representations that serve as inputs to other, specific modules but not to any central system of representation. On our view, therefore, the mind contains multiple, distinct mental languages, none of which is central or domain-general.

2 This pattern makes sense if one assumes that disoriented children have a representation of their geocentric heading, but that it has been rendered wholly inaccurate by the disorientation procedure. For example, a child who is set down facing North might represent her heading erroneously as southwesterly. If disoriented children, like rats, prefer to use the shape of the environment to make small rather than large corrections in perceived heading, then the child will correct her estimated heading from Southwest to South. Remembering that the toy was hidden in the northeastern corner, the child will then turn away from the two corners she is facing and search the geometrically correct corner behind her.

3 It may appear that a central question about conceptual enrichment has been begged: if the prelinguistic child can represent (a) through (c), don't these representations constitute a preexisting translation, in mentalese, of the English expression, *the picture near the window*? We submit that the prelinguistic child has no such translation, because (a)–(c) are not expressions in a single language of thought but expressions in three distinct systems of representation. Because of the modularity of these systems, these expressions cannot be combined directly with one another, although each can be mapped to language. Our account does beg a different question: if (a)–(c) are not expressions in a single mental language, what processes assure that the entity picked out by "x" in (b) is the same individ-