Early Category and Concept Development
Making Sense of the Blooming. Buzzing Confusion
Edited by Devid H. Rakison and Lisa M. Oakes   OXFORD University Press 2003

CHAPTER ELEVEN

Making an Ontology
Cross-linguistic Evidence

Linda B. Smith, Eliana Colunga, and Hanako Yoshida

And every language is a vast pattern-system, different from others, in which are culturally ordained the forms and categories, by which the personality not only communicates, but also analyzes, notices, or neglects types of relationship and phenomena.

> Benjamin L. Whorf, Language, Thought, and Reality

For the vocabulary of the language, in and of its self, to be a molder of thought, lexical dissections and categorizations of nature would have to be almost accidently formed, rather as though some Johnny Appleseed had scattered named categories capriciously over the earth.

> E. Rosch, "Linguistic Relativity"

HUMAN CULTURES AND LANGUAGES are diverse. To some, these differences imply incommensurate ways of being human. To others, these differences only serve to underscore our profound sameness. Most cross-linguistic studies of categorization offer up their evidence on one side or the other of this philosophical divide. In this chapter, we summarize recent results from our cross-linguistic studies of early noun learning by English-speaking and Japanese-speaking children. The findings are clearly relevant to issues of linguistic and conceptual diversity. However, these issues were not the proximal impetus for our studies. Instead, our questions were pitched at a different level, to a mechanistic understanding of the development of categories and early noun learning. Still, by pursuing mechanisms of developmental change, we arrive at a deeper understanding of the processes that create both universal and linguistically specific ways of knowing.

Universal Ontological Distinctions?

The things we encounter in our everyday lives seem to fall naturally into different kinds. There are animate things that react and intentionally move; there are discrete things with stable forms that we move; and there are substances, masses with less regular forms, that also do not move on their own. This partition of things into animals, objects, and substances is sometimes considered an ontological partition in two senses: in the Aristotelian sense, that these are three different kinds of existence, and in the psychological sense, that these are distinct psychological kinds that provide a foundation for human category learning. There is empirical support for the second idea from children's judgments in novel noun generalization tasks.

Kind-Specific Generalizations of Newly Learned Nouns

The novel noun generalization task measures children's expectations about the category organization of different kinds. In this task, the experimenter presents the child with a novel entity and names it with a novel name, saying, for example, "this is the mel." The experimenter then presents choice items and asks the child which of these can be called by the same name, saying, for example, "show me the mel." This is an interesting task because the naming event itself provides the child with few constraints on the class to which the name applies. Thus, children's generalizations from this minimal task input provide insights into children's expectations about how nouns map to categories. And the evidence indicates that children's generalizations honor an organization of kinds into animates, inanimate objects, and substances.

In particular, when 2½- to 3-year-old children are presented with novel solid and rigidly shaped things, they consistently generalize the name only to new instances that match the exemplar in shape but not to instances that match in other ways (Imai, Gentner, & Uchida, 1994; Landau, Smith, & Jones, 1988, 1992, 1998; Soja, Carey, & Spelke, 1991). However, when the named entity is a nonsolid substance such as hair gel or lotion molded into a shape, same-aged children are more likely to generalize the name by its material and color (Soja et al., 1991; Soja, 1992). Finally, when the named entity has properties typical of animate things—eyes or feet or limbs—children generalize the name narrowly to objects that match the named example in both shape and texture (Jones, Smith & Landau, 1991; Jones & Smith, 1998; Yoshida & Smith, in press; see also Gelman & Coley, 1991; Keil, 1994; Markman, 1989). Further, increasing evidence suggests that children learning a variety of languages such as Korean, Japanese, English, and Spanish make similar distinctions, naming rigidly shaped things by shape, nonsolid substances by material, and depictions of animate things by shape and texture (e.g., Gathercole & Min, 1997; Imai & Gentner, 1997; Lucy, 1996; Yoshida & Smith, in press).

Where Does This Knowledge Come From?

Evidence That Language Learning Plays a Role

Four facts suggest that language learning contributes to children's developing understanding of different kinds, as follows:

- Categorization taking place in naming and non-naming tasks
- Emergence of kind-specific name generalizations with vocabulary growth
- Modulation of kind-specific name generalizations
- Cross-linguistic differences

First, children's attention to the different properties of different kinds is evident most robustly in naming tasks. Many of the experiments showing that children systematically extend novel names in different ways for different kinds have included non-naming control tasks (e.g., Imai, et al., 1994; Jones et al., 1991, 1998; Landau et al., 1988, 1992, 1998; Soja et al., 1991). These control tasks are identical to the novel noun generalization task, except the object is not named. Instead, children are shown the exemplar and then are asked what other objects are "like" or "go with" the exemplar. In these non-naming tasks, children do not systematically attend to the different properties of different kinds. This fact suggests a mechanistic link between naming and knowledge about the category organizations of different kinds.

Second, kind-specific name generalizations emerge with vocabulary growth (Jones & Smith, 1997; Jones et al., 1991; Landau et al., 1988; Samuelson & Smith, 1999, 2000; Smith, 1999; Soja et al., 1991). The evidence indicates that the tendency to attend to shape in the context of naming emerges only after children already know some nouns. Moreover, this so-called shape bias in naming becomes stronger with development and more specific to solid and rigidly shaped objects. A bias to extend names for animates by similarity in shape and texture and a bias to extend names for substances by similarity in material emerge later (see, especially, Jones et al., 1991; Samuelson & Smith, 2000). Thus, biases to attend to different properties when extending names for different kinds codevelop with increasing vocabulary, a fact consistent with the idea that children's word learning helps create their category knowledge.

Third, kind-specific name generalizations are modulated by syntactic cues. One area of relevant research concerns the influence of count and mass syntactic frames on English-speaking children's interpretations of novel object and substance names. Count nouns are nouns that take the plural and can be preceded by words such as *a*, *another*, *several*, and *few*, as well as numerals. Count nouns thus label things we think of as discrete—chairs, trucks, shirts, studies, and hopes. Mass nouns, in contrast, cannot be pluralized but instead are preceded by words such as *some*, *much*, and *little*. Mass nouns thus label things that are conceptualized as unbounded continuous masses—water, sand, applesauce, research, and justice. Past research shows that count syntactic frames (e.g., *a mel*, *another mel*) push children's attention to the shape of the named thing, whereas mass syntactic frames (e.g., *some mel*, *more mel*) push attention to material (e.g., Gathercole, Cramer, Somerville, & Jansen, 1995; McPherson, 1991, Soja, 1992). In brief, language exerts an on-line influence on children's category formation.

Fourth, although there are clear universals in the name generalizations of children learning different languages—solid rigid things tend to be named by shape, nonsolid things by material, and things with features suggesting animacy by joint similarity in shape and texture—there are differences as well, differences that we

believe provide a potentially rich window on the role of language in creating knowledge about kinds. In the next section we present background evidence on differences between English and Japanese.

Language Differences

Individuation

Lucy (1992) proposed an animacy continuum that is intimately related to how languages individuate kinds. As illustrated in figure 11.1, this continuum orders kinds by the degree to which instances are marked as individuals by devices such as the plural and indefinite articles. On one extreme of Lucy's proposed continuum are animate entities, the kinds most likely to be treated as discrete entities by a language. On the other extreme are substances, the kinds least likely to be individualized by languages. In the middle are objects, entities that are treated as individuals by some languages but not by others. The key point is this: different languages emphasize different boundary points along a continuum of kinds from animate to substance.

English, with its count/mass distinction, is said to partition the continuum between objects and substances. Both common animal and object names—*cow* and *cup*—are count nouns. Both are thus kinds that English treats as discrete entities. Common substance names such as *milk, sand,* and *wood,* in contrast, are mass nouns in English. These are treated by the language as unbounded continuous entities. Thus, through devices such as the indefinite article, pluralization, and quantification,

The animacy continuum

animates	objects	substances

←——————————————————————————————
likelihood individuated

English

animates	objects	substances

individuals

Japanese

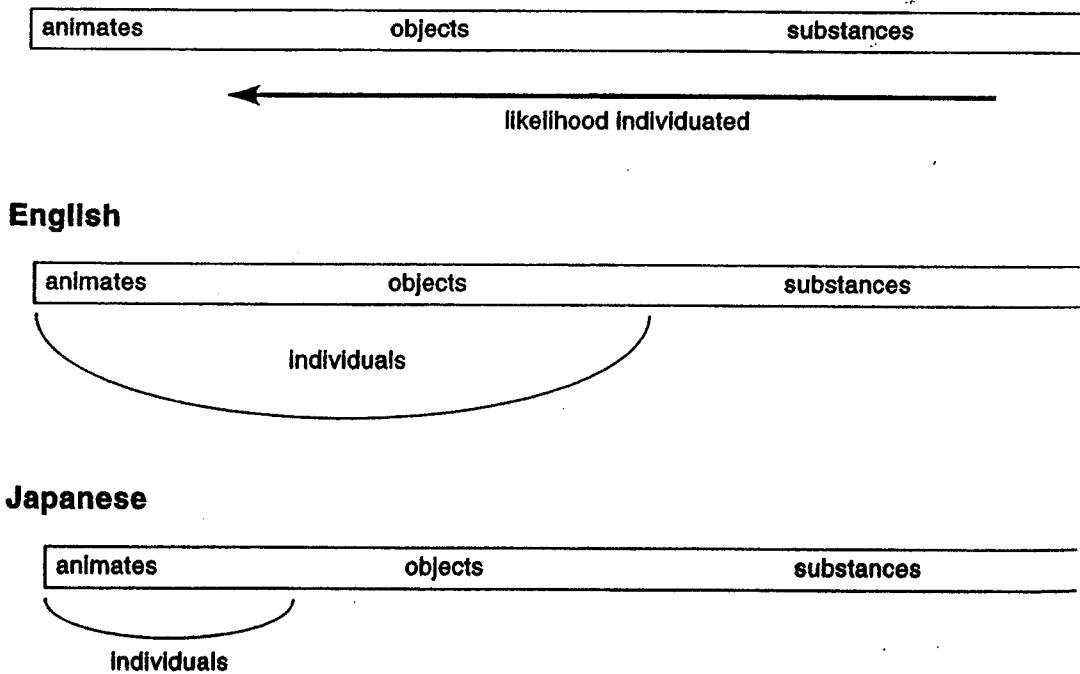| animates	objects	substances

individuals

Figure 11.1. The animacy continuum and individuation in English and Japanese.

English treats animate and object names in the same way and differently from substance names.

The Japanese language, in contrast to English, appears to partition the continuum between animates and inanimates, treating only animates as discrete individuals. First, Japanese nouns that refer to multiple entities are not obligatorily pluralized. Thus *inu ga ita* could mean either "there was a dog" or "there were dogs." However, nouns referring to multiple humans or young animals are optionally pluralized with the suffix *tachi.* Thus, *koinu tachi ga ita* is "there were some puppies." The plural suffix appears not to be used on inanimate nouns. Second, when Japanese speakers do need to count discrete entities, they use a system of classifiers that often depend on the kind of thing being counted, much as English speakers count *loaves* of bread or *panes* of glass. The Japanese classifiers used for animates tend not to overlap with those used for inanimates. Finally, a distinction between animates and inanimates is also supported by other aspects of Japanese than plurals and quantification. Although not traditionally viewed as markers of individuation, there are additional aspects of Japanese that are closely linked to individuation and animacy (see Yoshida & Smith [in press], for further discussion). One of these is the distinction between *aru* and *iru.* For the very fundamental notion of existence ("there is") and spatial location ("be located"), Japanese has separate verbs for animates and inanimates: *aru* is "inanimate object exists/is located" and *iru* is "animate object exists/is located." Thus Japanese, through pluralization, its classifier system, and the *iru/aru* distinction in locative constructions, imposes a boundary between people and animals on the one hand and objects and substances on the other.

These are systematic language differences of the kind likely to matter in children's developing conceptualizations of kinds (Gumperz & Levinson, 1996; Lucy, 1996): noun categories in English are systematically partitioned into object names versus substance names, whereas noun categories in Japanese are systematically partitioned into names for animates versus names for inanimates. Both Quine (1969) and Lucy (1992) suggested that the partitions made by a language's system for marking individuals determines the ontological partitions made by speakers of that language.

Complete linguistic determination, however, seems unlikely, as there is relevant perceptual information about category structures that is available to speakers of all languages. Indeed, prelinguistic infants distinguish animate categories from objects that hold their shape and form, and also distinguish rigid forms from nonrigid ones (e.g., Spelke, Vishton, & Van-Hofsten, 1995).

Imai and Gentner's Results

Imai and Gentner's (1997) cross-linguistic study of the object–substance boundary provides clear evidence that both linguistic and perceptual information contribute to an object–substance distinction. In their study, they compared Japanese-speaking and English-speaking children's generalizations of names for novel solid and nonsolid forms. They used three kinds of stimulus sets: solid and complexly shaped things, solid but simply shaped things, and nonsolid and thus simply shaped substances. They did this because solids and nonsolids differ in the kinds of shapes they usually take. Solid things can be quite complex—with many angles and multiple parts.

Nonsolid substances, however, cannot take angular multipart shapes and over time relax toward rounded and accidental-appearing forms like splatters and drops.

In the experiment, Imai and Gentner presented children with an exemplar and named it with a novel noun. They used a syntactic frame in English that was neutral, consistent with either a count or mass noun. In this way, any language effects would be off-line effects, effects of a history of making distinctions between count and mass nouns in English and not making such a distinction in Japanese. After the exemplar was named, the child was shown two choice objects, one that matched the exemplar in shape and one that matched the exemplar in material. The child was asked to indicate the one called by the same name as the exemplar.

Imai and Gentner found that Japanese speakers and English speakers formed similar categories for solid complexly shaped things, generalizing a newly learned object name to new instances by shape. And speakers of both languages increased attention to material when the named entity was nonsolid. Imai and Gentner concluded from these similarities that the partition of objects from substances does not depend on linguistic individuation, since both English-speaking and Japanese-speaking participants treated solids and nonsolids differently, even though Japanese does not mark objects and substances differently.

However, Imai and Gentner also found differences between the novel noun generalizations of English and Japanese speakers. Most notably, English and Japanese speakers differed in their generalizations of names for simply shaped solids. Simply shaped solid things are like objects in the rigidity of their shapes but are like substances in the simplicity of their shapes. English speakers treated the simply shaped solid things as objects and generalized their names by shape, whereas Japanese speakers were more likely to generalize the name by material. The results suggest that as a consequence of different systems of individuation, Japanese and English speakers place the boundary between objects and substances in slightly different places. For speakers of English, solid things—both complexly and simply shaped—are categorized as objects, that is, by shape. For speakers of Japanese, simply shaped things—both solid and nonsolid—are more likely to be categorized as substances, that is, by material.

Ontologies as Statistical Regularities

Imai and Gentner's results show both universal and language-specific influences on children's "ontological" distinctions. We propose that both the universals and the differences are the product of the same statistical learning mechanism, arising from correlations among the perceptual properties of different kinds, lexical category structures, and linguistic devices concerned with individuation. This proposal is based on the following five core ideas.

1. *There are regularities that distinguish kinds of things in the world and our perceptual systems are sensitive to these regularities.* Solids, nonsolids, and animates present correlated bundles of perceptual properties.
2. *The nominal categories of languages honor these correlational bundles.* Languages evolved to fit the perceptual system and the world. Thus it

makes sense that lexical categories across-languages respect and make use of the same correlated perceptual properties that distinguish animates from solid objects and from nonsolid substances.

3. *Word learning enables higher-order generalizations.* Word learning may be mechanistically crucial to going beyond specific knowledge about specific kinds to developing the higher-order correlations that constitute kind-specific noun generalizations and ultimately abstract knowledge, knowledge we might rightly call an "ontology."

4. *The mechanism is associative learning.* The statistical regularities characteristic of early noun categories may be sufficient in and of themselves to create a partition of things into animals, objects, and substances. Ontologies in their psychological sense could be the generalizations that arise naturally from the statistical regularities across lexical categories.

5. *Linguistic regularities are part of the associative mix and thus bend knowledge in language-specific ways.* Linguistic forms that are regularly associated with correlated bundles of perceptual cues may reinforce the connections between those perceptual cues. In this way, systematic linguistic contrasts, such as those that compose a language's system of quantifying individuals, may differentially bolster and weaken perceptual correlations, changing how things are perceived and conceived.

We present preliminary support for these ideas in the remainder of this chapter. We do so by first concentrating on Imai and Gentner's finding of differences in the object-substance boundary for children learning English and Japanese. We then turn to a parallel phenomenon at the animal-object boundary. Finally, we propose how abstract ideas about even abstract kinds might emerge from these correlations across categories of concrete things.

Creating an Object-Substance Boundary

Early noun categories are highly structured. They present the kinds of regularities that could yield a partition of kinds into objects and substances. Specifically, early learned categories of solid things are well organized by shape, and early learned categories of nonsolids are typically well organized by material. This is so in both English and Japanese.

Regularities in the Early English Lexicon

Samuelson and Smith (1999) asked: What kinds of nouns do young children learning English know? Do they learn names for solid things in shape-based categories and names for nonsolid things in material-based categories? To answer these questions, Samuelson and Smith examined the structure of a set of nouns that are typically known by children at 30 months. More specifically, they examined the list of nouns that compose the MacArthur Communicative Developmental Inventory (MCDI). The MCDI is a parent checklist that is used by many researchers to mea-

sure the vocabulary of children from 16 to 30 months. The list of words on the MCDI was developed from extensive studies of parental diaries, in-laboratory testing of early vocabularies, and large normative studies (Fenson et al., 1993). The nouns contained on the MCDI are known by 50 percent of children at 30 months. Samuelson and Smith specifically examined the category structures of 312 nouns—all the nouns in the animals, vehicles, toys, food and drink, clothing, body parts, small household items, and furniture and rooms sections of the MCDI.

The method used to examine the category structure of these 312 early learned nouns was borrowed from the pioneering work of Rosch (1973). Adults were presented with each noun on the list of 312 and asked to think of the instances named by each noun. For example, they might be told: "Think of apples that you commonly experience." Then, while thinking about these instances, the adults were then asked a series of yes/no questions: "Are these similar in shape? Are these similar in color? Are these similar in material? Are these solid? Are these nonsolid?" A separate group of adults was presented with the criteria for distinguishing count and mass nouns and asked to judge whether each noun on the MCDI was a count or a mass noun or could be used both syntactic frames (e.g., *cake*). To classify a nominal category as possessing any of these properties, Samuelson and Smith required that 85 percent of the adults agreed with that characteristic. This conservative criterion was used to ensure that the regularities attributed to the early lexicon were likely to be ones that are manifest in the experiences of most young learners. In this way, each noun was categorized as shape-based, material-based, color-based, based on a combination (or all) of these properties, or based on none of these properties. Each noun was also classified as referring to solid or nonsolid things or ambiguous insolidity, and each noun was classified as a count noun, a mass noun, or as ambiguous in its syntactic category.

Figure 11.2 summarizes the key regularities in terms of Venn diagrams. In these diagrams, the relative size of each circle represents the relative numbers of nouns of that kind, and the size of the overlap between intersecting circles represents the relative number of nouns of both kinds. The circles on the left depict the relative number of count nouns, names for solid things, and names for categories organized by shape. The circles on the right represent the relative numbers of mass nouns, names for nonsolid substances, and names for things in categories organized by material. (Color is not shown because so few categories were judged to be similar in color independently of similarity in material.) What the figure shows is that many early nouns are count nouns, many refer to solid objects, and many name objects in shape-based categories. Moreover, count nouns, solid things, and shape similarity go together. The right side of figure 11.2 shows that there are many fewer nouns in this corpus that are mass nouns, name nonsolid things, and name categories organized by material. However, nonsolidity, mass-noun syntax, and material-based categories are correlated. Thus, the early English lexicon presents correlations among category structures, the perceptible properties of solid and nonsolid things, and count-mass syntactic cues. The regularities are clearly lopsided—much stronger on the solid, shape, count side than on the nonsolid, material, mass side.

One might ask: Why does the early noun corpus have the structure it does? Sandhofer, Smith, and Luo (2000) examined transcripts of parent speech to young
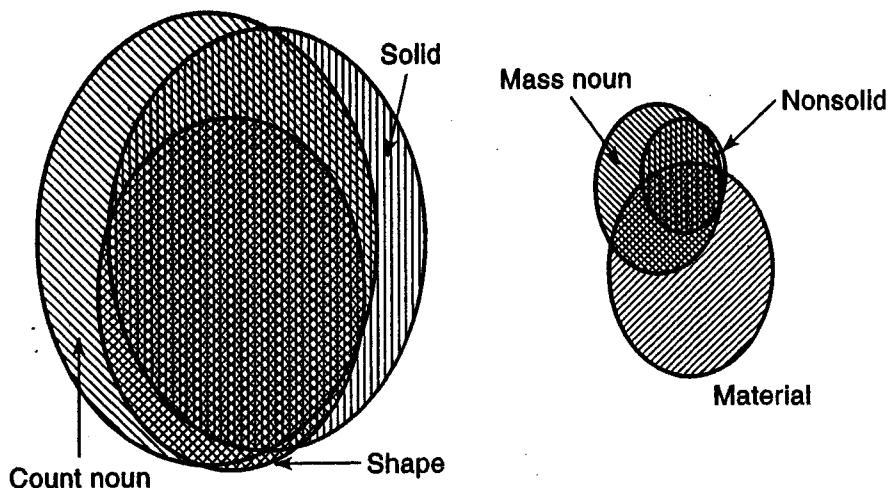
Figure 11.2. Venn diagrams illustrating the overlap among shape-based categories, solidity, and count syntax and material-based categories, nonsolidity, and mass syntax among the 312 early-learned English nouns.

children. They selected the most common 100 nouns and asked adults to judge the category structure, using the same method used by Samuelson and Smith. They found evidence for the same correlational structure as had Samuelson and Smith and the same emphasis on naming solid things in shape-based categories. We suspect that the structure of the common nouns children hear and use reflects deep truths about the perceptual regularities in the world and their functionality from a human perspective.

Children, however, must individually learn these deep truths. The evidence indicates that learning names for things is a crucial part of this. Children's kind-specific name generalizations become organized *as* they learn more and more names for different kinds (for review, see Smith, 1999). In line with previous results, Samuelson and Smith (1999) found that when children knew few nouns, they did not honor a distinction between solid and nonsolid things. Instead, they generalized novel names for solid things by shape only after they had already learned a substantial number of names for solid things, a fact that fits the idea that children's novel noun generalizations are themselves generalizations over the structure of already learned nouns. Further, children generalized names for solid things by shape long before they generalized names for nonsolid things by material—a fact that also aligns with the statistical regularities across early English noun categories.

Regularities in the Early Japanese Lexicon

What are early learned nouns in other languages like? Do they name the same kinds of categories as do the early English nouns? Colunga and Smith (2000) addressed this question by examining the nouns on the Japanese MCDI. The Japanese MCDI, like the English one, is a parent checklist of early-learned words and phrases. The Japanese MCDI was independently constructed and normalized across large samples of children learning Japanese as their first language (Ogura & Watamaki, 1997; Ogura,

Yamashita, Murase, & Dale, 1993). The Japanese MCDI, like its English counterpart, contains the words and phrases that 50 percent of children in the normalized samples know at 30 months. The list of nouns on the Japanese MCDI thus are a good measure of the first nouns learned by children learning Japanese.

When one compares the list of early English nouns and the list of early Japanese nouns, some differences are immediately apparent. For example, the Japanese checklist has more animal terms (52 vs. 43), more food terms (77 vs. 68), more people terms (34 vs. 26), and more body parts (33 vs. 27). In contrast, the English checklist has more names for artifacts. And the specific lexical categories differ. There is little overlap among early food categories, the dominant segment of names for nonsolid substances in both vocabularies. In addition, there are many differences in animal names. The Japanese list of animal terms includes *shrimp, crab, hippopotamus, kangaroo, koala, rhinoceros,* and *swallow*—none of which are on the English list of early known animal names. But, importantly, the early Japanese corpus, like the early English one, presents clear evidence of different category organizations for solid and nonsolid things.

In an effort to understand whether early English and Japanese nouns lexicalize categories of solid and nonsolid things similarly, Colunga & Smith (2000) examined the category structures of all food and concrete object terms on the Japanese and English lists. (That is, unlike Samuelson and Smith, they excluded animal terms and abstract terms such as "friend.") In total, 167 nouns on the Japanese MCDI are food or concrete object terms, and 150 nouns on the English MCDI are food or concrete object terms. Colunga and Smith asked native speakers to judge the category structure of each noun category using the same method as did Samuelson and Smith.

The results are presented in figure 11.3 as Venn diagrams. The larger outline area represents all the nouns that were judged in the language—including those that did not reach the strict agreement criteria. The smaller rectangles inside the larger area indicate by size the numbers of lexical items that did reach the strict agreement criteria. Black areas represent the numbers of nouns judged to refer to solid things, and white areas represent the numbers of nouns judged to refer to nonsolid things. Horizontal stripes indicate the numbers of nouns judged to refer to objects of similar shape, and vertical stripes the numbers of nouns judged to be similar in material and/or color.

As can be seen, in both languages about half of these early learned nouns refer to solid objects (42 percent in English, 48 percent in Japanese) and there are fewer (24 in English, 21 in Japanese) that name nonsolids. Further, in both languages more nouns were judged to refer to things similar in shape (38 percent in English, 49 percent in Japanese) than to things similar in material and/or color (31 percent in English, 20 percent in Japanese). And, crucially, solidity and category organization are correlated. Again, the correlation is very strong for solidity and within-category similarity in shape, with most of the words that were classified as referring to solid things also judged to refer to things that were similar in shape (79 percent in English, 93 percent in Japanese), and most of the words that were classified as referring to things similar in shape were also classified as referring to solid things (88 percent in English, 90 percent in Japanese). Again, the correlation was weaker for nonsolids and material-based category organizations. Whereas words that were classified as refer-

Japanese   English

■ solid   ▤ shape
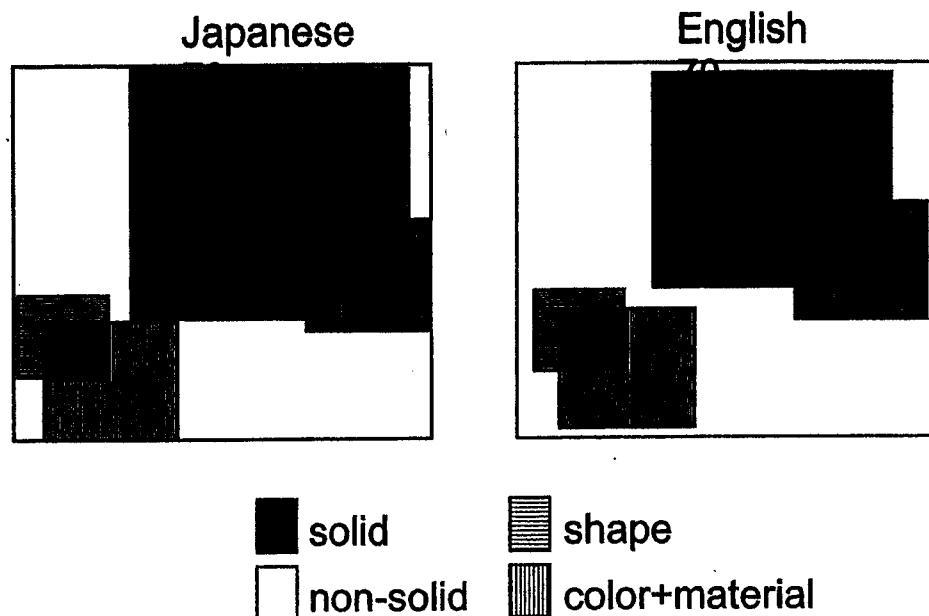□ non-solid   ▥ color+material

Figure 11.3. Venn diagrams of the overlap among shape-based and material based categories and the solidity and nonsolidity of instances for early learned food and object terms in English and Japanese.

ring to nonsolids were judged to refer to things that were similar in material (96 percent in English, 81 percent in Japanese), the correlation did not hold in the opposite direction (49 percent in English, 52 percent in Japanese).

The key result, then, is this: the same regularities characterize object and substance terms in the two languages.

Network Simulations

Are these regularities enough in and of themselves to create the universals in children's kind-specific generalizations? If children's knowledge about solid objects and nonsolid substances are the direct product of the statistical regularities among the nouns children know, then a simple learner of statistical regularities, a connectionist net, should develop similar knowledge if trained on a "vocabulary" similar to that of young children. Thus, we tested the idea of "ontology" as statistical regularities by feeding these regularities to a simple statistical learner.

Importantly, although connectionist networks are simple associative learners, the generalization the network needs to make to reproduce children's kind-specific noun generalizations is not simple. It requires going from simple associations to abstract, rule-like generalizations. For example, in the training phase of the simulations, we taught networks names for specific instances of specific categories—for example, the word "ball" associated with round things of variable color and material and the word "sand" associated with things of a particular material and range of colors. These kinds of associations are easy for networks to learn; and it is easy for networks to generalize from some specific instances of a category (e.g., from specific balls) to new in-

stances of the same category (e.g., to never-before-encountered balls). The theoretical question, however, concerns not the learning and generalization of these specific categories but the emergence of the higher-level abstraction: that solidity signals the relevance of shape and that nonsolidity signals the relevance of material—for objects and substances never encountered before and shapes and materials never experienced before. Thus, prior to the simulations, it was an open question: Are the correlational structures manifest in early English and Japanese noun vocabularies enough to yield kind-specific category organizations when given novel things?

To address this question, Colunga and Smith (2000) taught the early English vocabulary to one set of networks and the early Japanese vocabulary to another. Two specific issues were at stake: (1) Would both sets of networks learn the same distinction, naming complexly shaped solid things by shape and nonsolid substances by material? and (2) Could the small differences in the statistical structures of the early noun lexicon in the two languages possibly be sufficient to create the differences in how English-speaking and Japanese-speaking children generalize names for simply shaped solids?

The Network   We used a Hopfield network, which is a simple recurrent network. The networks were trained using contrastive Hebbian learning, an algorithm that adjusts weights on the basis of the correlations between unit activations. Figure 11.4 shows the architecture of the network. It has a word layer, in which each unit corresponds to one word in the training vocabulary. Individual objects are represented on what we call the object layer. Activation patterns on this layer represent the shape and material of each individual object or substance presented to the network. More specifically, the shape and material of an object (say the roundness of a particular ball and its yellow rubbery material) are represented by an activation pattern along the whole layer, in a distributed fashion. In the solidity layer, one unit stands for solid and another for nonsolid. Finally, there is a hidden layer that is connected to all the other layers and recurrently with itself. Note that the word layer and the object layers are only connected through the hidden layer; there are no direct connections among them.

Training   We trained networks on the "English" or "Japanese" nouns. The goal was to mimic the vocabulary learning that a child brings into a novel noun generalization experiment. The statistical regularities characteristic of the early vocabularies were built into the network's training set in the following way. First, for each word that the network was to be taught, a pattern was generated to represent its value along the relevant dimension—the dimension that the English-speaking and Japanese-speaking adults said characterized the similarities of objects named by the noun. Second, at each presentation of the word, the value along the irrelevant dimension for that lexical category was varied randomly. For example, the word "ball" was judged by the English-speaking adults in the Samuelson and Smith (1999) study to refer to things that were similar in shape; thus, a particular pattern of activation was randomly chosen and then assigned to represent ball-shape. All balls presented to the network were defined as having this shape, although each ball presented to the network also consisted of a unique and randomly generated pattern defining the ma-
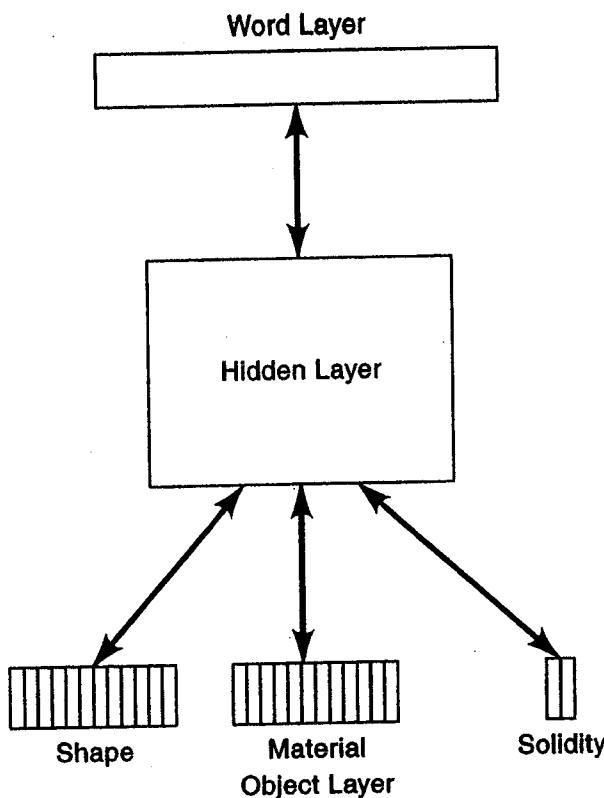
Word Layer

Hidden Layer

Shape   Material   Solidity
            Object Layer

Figure 11.4. The architecture of the network used by Colunga and Smith. See text for further clarification.

terial and color. So whenever the network got the unit representing the word "ball," it also got the pattern representing ball-shape along the shape dimension and a different pattern along the material dimension.

We also built into the training the shape regularities that distinguish solid and nonsolid things. Specifically, in the simulations, although instances of most solid categories were the same shape (in proportion to the adult judgments), instances of different solid categories differed greatly in shape, instantiating the full range of possible shapes. In contrast, instances of the same nonsolid category typically differed in shape (in the same proportions as the adult judgments), but overall, nonsolid instances for all categories of nonsolid things were drawn from a relatively restricted range of possible shapes.

Noun Generalization Test   After teaching a network the "English" or "Japanese" vocabulary, we tested the network's expectations about how novel solid and nonsolid things should be named. Our approach to testing the networks is based on our conceptualization of the novel noun generalization task. In that task, the child sees an exemplar and hears its name and then is presented with two choice items—one matching the exemplar in shape and one in material. We propose that the child generalizes the name to the choice item that is *perceived* as most similar to the exemplar. If, for example, the child *attends* exclusively to the shape of the named exemplar,

then a test object that matches the exemplar in shape (although different from the exemplar in material) should be perceived as highly similar to the exemplar, and the child should generalize the name to that item. Thus, to measure selective attention, we asked if the network's internal representations of a named exemplar and a test object were similar. More specifically, we asked if the patterns of activations on the hidden layer for the named exemplar and shape-matching test item were more or less similar than the patterns of activation on the hidden layer for the named exemplar and the material-matching choice item. Thus on each test trial, a novel exemplar object was generated by randomly creating an activation pattern along the shape and material dimensions. Then a novel shape-matching test object was generated by combining the exemplar's shape pattern with a novel randomly generated material pattern. A similarity measure of the exemplar and the shape match was computed in terms of the distance between the activation patterns in the hidden layer after the exemplar and its shape match were presented. Similarly, a novel material-matching test object was generated by combining the exemplar's material pattern with a new randomly generated shape pattern and the similarity between exemplar and material match was computed. Finally, we used these similarity measures between the emergent patterns of activation on the hidden layer to calculate the probability of choosing the shape and the material match using Luce's forced choice rule (Luce, 2000).

In this way, we trained 10 networks (with 10 different randomly generated initial connection weights) with categories structured like the object and substance terms young English-speaking children know. During training, we presented multiple instances of each trained noun until the network stably produced the right noun when presented an instance of each kind. We taught nouns with different category organizations in the same proportions that are found in young English-speaking children's lexicons. We then tested each of these English networks in the novel noun generalization task—with 30 novel exemplars. These 30 test trials were divided evenly into three kinds: the exemplars were defined by patterns of activation that represented (1) solid and complexly shaped things, (2) solid and simply shaped things, and (3) nonsolid and simply shaped things. In the same way, we trained 10 networks with all the words in the Japanese corpus and, at the end of this training, tested those 10 Japanese networks with the same 30 novel noun generalization trials. If the statistical regularities in the two vocabularies are sufficient to create a common solidity–nonsolidity distinction as well as the cross-language differences, then the performances of these networks should look like the performances of the children in the Imai and Gentner's study.

Results   In figure 11.5A, we compare the performances of the networks to the patterns reported by Imai and Gentner (1997) for 2-year-olds—the relevant age for the training corpus. The solid bars show the 2-year-old children's performances from the Imai and Gentner study—the proportion of times children extended the name of the object to the test object matching in shape. Since children always chose between a shape-matching and material-matching test object, chance is .50, and systematic extensions by material are indicated by below-chance performance in the figure. The striped bars in the figure show the mean of the networks' performances.
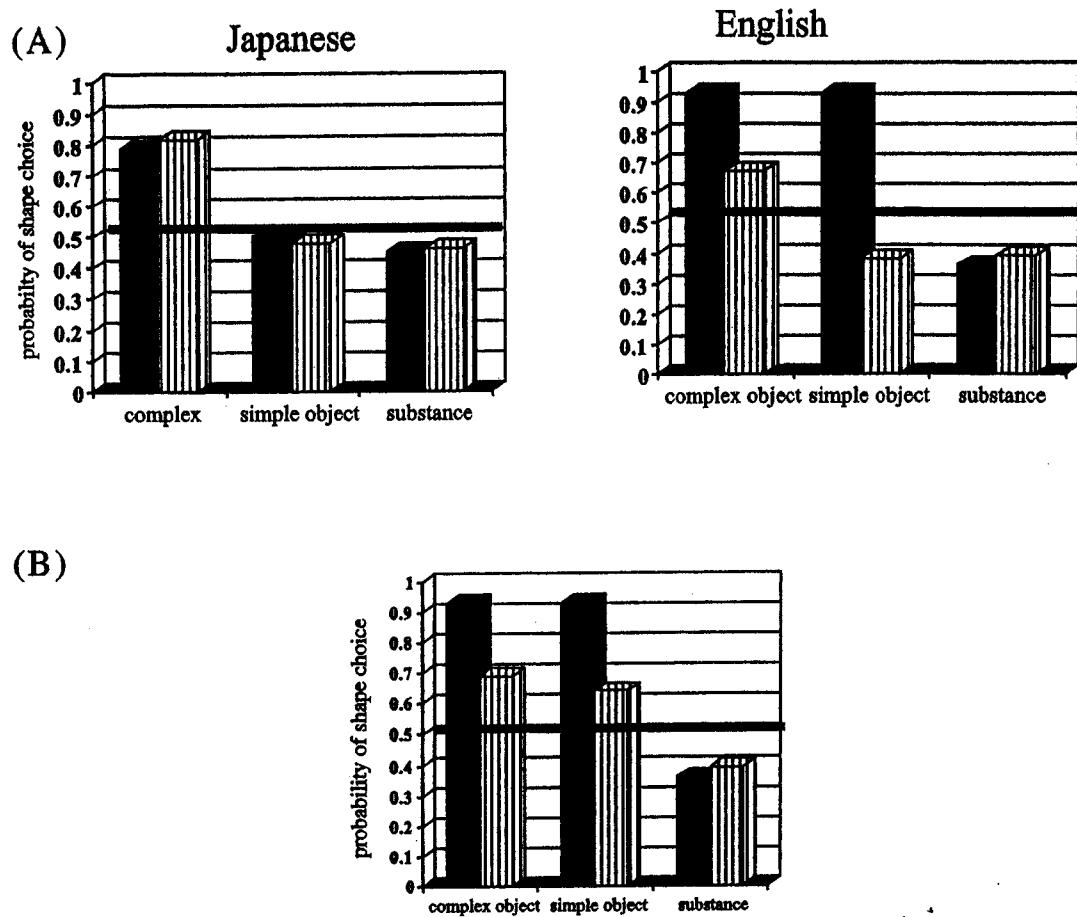
(A)

Japanese   English

(B)

Figure 11.5. (A) *The mean proportion of shape choices by English-speaking and Japanese-speaking 2-year-olds in Imai and Gentner (1997) experiment, as a function of the solidity of the exemplar and the mean proportion of shape choices predicted by the networks trained on early English or Japanese nouns. (B) The mean proportion of shape choices predicted by network trained on English nouns with correlated count-mass syntax, and for comparison, the mean proportion of shape choices by-2-year-old English-speaking children as a function of solidity, as reported by Imai & Gentner (1997).*

Consider first the performances of the Japanese-speaking children and the networks trained on the Japanese noun categories. Names for complexly shaped objects are generalized by shape. Names for simple solids and for nonsolid substances are much less likely to be generalized by shape and often (more than half the time) are extended to new instances that match the named exemplar in material. The networks taught noun vocabularies with the same statistical structure as the noun vocabularies known by 2-year-old Japanese children generalize names for novel entities in the same way as the Japanese-speaking children. Complexly shaped things are named by shape, but simply shaped things—solid or nonsolid—are not. The fact that the networks mimic the performances of Japanese-speaking children tells us that the structure of the early noun lexicon is itself enough to create a distinction between objects and substances—with the boundary between object and substance being determined by the complexity of the shape.

Now consider the performances of the English-speaking children and the networks trained on the English noun categories. The children show a much stronger bias to extend names for solid things by shape than do the networks. This is particularly so for the solid simply shaped things. The children extend names for all solid things—simple or complex in shape—by shape; but they are more likely to extend names for nonsolid substances by material. Thus for English-speaking children, the partition between object categories organized by shape and substance categories organized by material is defined by solidity. The networks trained on the English nouns, in contrast, extend names for solid complexly shaped things by shape (reliably more often than expected by chance), but extend names for simply shaped things—solid or nonsolid—by material. The boundary between object and substance imposed by the English trained networks is based only on the statistical regularities in the early English noun categories and, like Japanese children but not English children, the boundary appears to be defined by complexity of shape rather than solidity. These results tell us that the structure of the early English noun lexicon is not enough in and of itself to explain English-speaking children's novel noun generalizations.

Adding Syntax

What is missing from the simulations of the English-speaking children? The obvious additional factor relevant to English-speaking children's learning is count–mass syntax. Therefore, in the next simulation, we added the count–mass syntax correlations to the English-trained networks.

For this simulation we added an additional input layer to those illustrated in figure 11.4: the syntax layer. The syntax layer had two units, one to represent count syntax and one for mass syntax. The networks were trained on the same English vocabulary, but now each noun was associated with count/mass syntax information, according to adults' judgments as collected by Samuelson and Smith (1999). Nouns that adults judged to be both count and mass nouns (e.g., "cake" and "muffin") were associated equally often with both the count and mass units "on."

The results of the network simulations are shown compared to children's performance in figure 11.5B. Although the connectionist networks trained on English with the correlated count–mass syntactic cues show a quantitatively weaker shape bias than do children, they were successful in simulating the qualitative pattern. The networks, like the children, now generalize names for solid things—simple and complex—by shape, and names for nonsolid things by material. Learning names for concrete objects and substances in both languages appears to create knowledge that objects and substances are named by different properties. But language-specific syntactic cues in English shift this "ontological" boundary relative to that of Japanese speakers.

Conclusion

The kinds of nouns known early by children learning English and by children learning Japanese present an organized structure. Most name solid things, and solid things with the same name tend to be similar in shape. A coherent subset of nouns name nonsolid substances, and substances with the same name tend to be similar in mate-

rial. The simulations show that these regularities are sufficient to create the similarities in English-speaking and Japanese-speaking children's novel noun generalizations; that is, a shape bias when naming complexly shaped solids and (to a lesser degree) a material bias when naming nonsolids. Similarly structured lexical categories create similar knowledge about object and substance categories. The results also suggest that language-specific syntactic cues are part of the correlational mix, modulating the object–substance partition in language-specific ways.

The Animate-Object Boundary

If language-specific cues correlate with perceptible differences among kinds and influence ontological boundaries, then there should be cross-linguistic differences at the animate-object boundary for Japanese- and English-speaking children. This should be so because Japanese adds linguistic cues to the statistical mix that are correlated with an animal versus object-substance partition, just as English adds cues to the associative mix that are correlated with a animal-object versus substance partition.

Iru/Aru

Of all the distinctions in Japanese that focus on animacy, the *iru/aru* distinction seems a likely powerful force on the way Japanese children think about objects. This distinction involves fundamental notions of existence ("there is") and spatial location ("be located"). In English we use the same verb "be" for a dog, a cup, and water, saying: *there is a dog, there is a cup*, and *there is water*. However, the Japanese verb *iru* is used for a dog, and *aru* is used for a cup or water. *Iru* implies being in a place by one's own will. *Aru*, on the other hand, implies "having been left" at a place. Importantly, *iru* is used whenever one refers to entities that behave intentionally, for example, people and animals. Critically, *iru* is also used by adult speakers (and children) when inanimates are conceptualized as animates. For example, *iru* is used by adults when referring to dolls and toys as the animates they depict in play and conversations with children. Thus every time a Japanese-speaker refers to the location of an object, the speaker must decide if the object is to be conceptualized as animate or inanimate.

Yoshida and Smith (2001, in press) demonstrated that 2- to 3-year-old Japanese-speaking children understand the implications of *iru* and *aru* in a novel noun generalization task. The children in this study were monolingual and were tested in Japan. The children were presented with three-dimensional objects that were ambiguous and could be seen as depictions of animates or artifacts. As illustrated in figure 11.6, each object had four pipe-cleaner appendages. The objects could be conceptualized as animal depictions if the appendages were construed as limbs, but they also could be easily viewed (as least by our intuitions) as artifacts and not animal-like at all. The exemplar objects were named either using a sentence frame with *aru* (suggesting an artifact) or with *iru* (suggesting an animate entity). In a yes/no version of the novel name generalization task, the child was shown the exemplar and told its name—"This
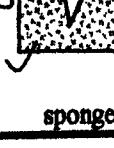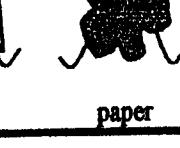
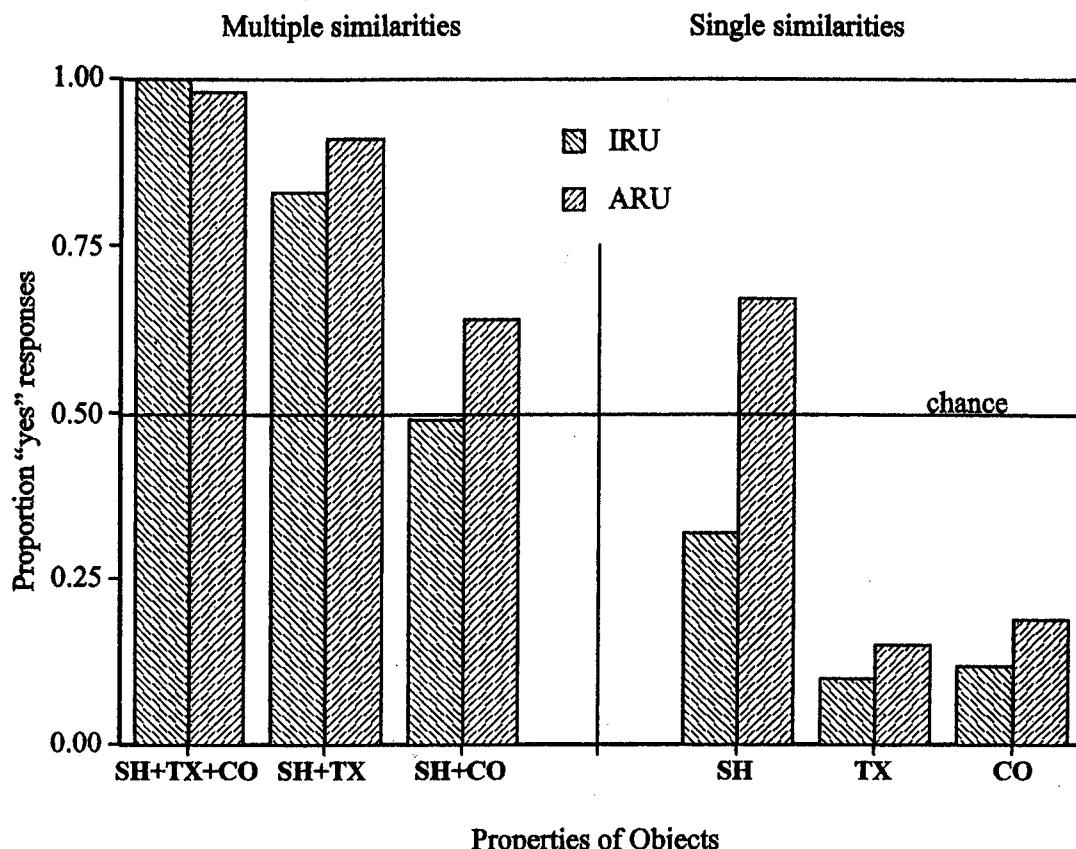| Exemplar	Test objects					
	Multiple similarities			Single similarity		
Keppuru	sh+tx+co	sh+tx	sh+co	sh	tx	co
smooth clay	smooth clay	smooth clay	sponge	shiny hair	smooth clay	perforated clay
Tema	sh+tx+co	sh+tx	sh+co	sh	tx	co
sponge	sponge	sponge	paper	tinsel	sponge	paper

Figure 11.6. The ambiguous objects used by Yoshida and Smith.

is a mobit"—and then each test object was presented individually. The child was asked about each test object "Is this a mobit?" The results are shown in figure 11.7. When the novel name was presented in a sentence frame containing *aru*, the Japanese-speaking children generalized the name to all test objects matching the exemplar in shape, regardless of whether these test objects matched or mismatched the exemplar on the other properties. In contrast, when the exemplar's name was presented in the context of *iru*, children generalized the name conservatively, only to objects that matched the exemplar on multiple properties, and particularly in shape and texture. The pattern in the *iru* condition fits past findings on children's extensions of names for animal-like things; for animals, shape alone is not enough, and multiple similarities are required to extend the name (e.g., Jones et al., 1991; Jones & Smith, 1997).

These results provide three important pieces of information. First, the linguistic cues of *iru/aru* alter the way Japanese-speaking children categorize novel objects. This tells us that young Japanese-speaking children do have knowledge of at least one linguistic device that privileges animate kinds. Second, young Japanese-speaking children generalize names for implied artifacts to new instances more broadly than they generalize names for implied animals. Third, the results tell us that linguistic cues, at least explicitly present ones, can alter how the same perceptual entity is conceptualized—as a depiction of an animate or artifact kind.

Variation at the Animate-Object Boundary

In the world, animate and inanimate things differ in many ways. They have different properties, such as eyes and limbs versus angular parts. They move differently. And people talk about them differently. In brief, the world presents the learner with

Multiple similarities   Single similarities

Figure 11.7. Mean "yes" responses—the name of the exemplar extends to the test object— by Japanese-speaking children, when the name was presented in locative contructions using iru or aru, as a function of the properties of the test object that match those of the named exemplar: shape (SH), texture (TX), color (CO).

a richly structured set of associations. *Iru* and *aru* and a host of other linguistic distinctions centered on animacy are part of this associative mix for children learning Japanese. Does this alter the way Japanese children perceive animate and inanimate things? We hypothesized that Japanese children, relative to their English counterparts, might be hypersensitive as to whether some object should be construed as an animate versus an artifact. That is, given ambiguous objects with features merely suggestive of limbs, Japanese-speaking children should be more likely than English-speaking children to see the appendages as limblike and to construe the objects as depictions of animate things, even when the linguistic context is neutral and offers no suggestion as to how the object should be construed. This should be so if the linguistic distinction in the language heightens attention to cues relevant to making perceptual distinctions (see Lucy, 1996).

Yoshida and Smith (in press) tested this prediction by comparing 2- to 3-year old Japanese- and English-speaking children's name generalizations using the same stimuli as in figure 11.6. The sentence frames used in Japanese were nonlocative constructions that did not require *iru/aru*, rather, the same sentence frame could be used with both animates and inanimates.

Figure 11.8 shows the mean proportion of "yes" responses as a function of language and individual test objects. As is apparent, when presented with ambiguous objects named with novel names in a neutral sentence frame, Japanese-speaking children generalized the names in the same way they did when the name was presented in the context of *iru*, a context that unambiguously implies animacy. That is, Japanese-speaking children generalized the exemplar's name only to items that matched the exemplar in both shape and texture and rejected all other test objects as instances of the lexical category. In contrast, the English-speaking children generalized the novel names in the same way that Japanese-speaking children had when the name had been presented in the context of *aru*, a context that unambiguously implies an inanimate thing. English-speaking children, like Japanese-speaking children in the *aru* condition, generalized the name to all objects that matched the exemplar in shape—both when that object matched in other properties and when it did not.

Here, again, we see the effect of the language one is learning on the ontological boundary. *Iru* and *aru* are correlated with things that present different perceptible properties—those that distinguish a real and unambiguous animate, like a living dog, from an inanimate thing, like a cup. *Iru* and *aru* are also correlated with lexical category structure—categories organized by joint similarity in shape and texture versus
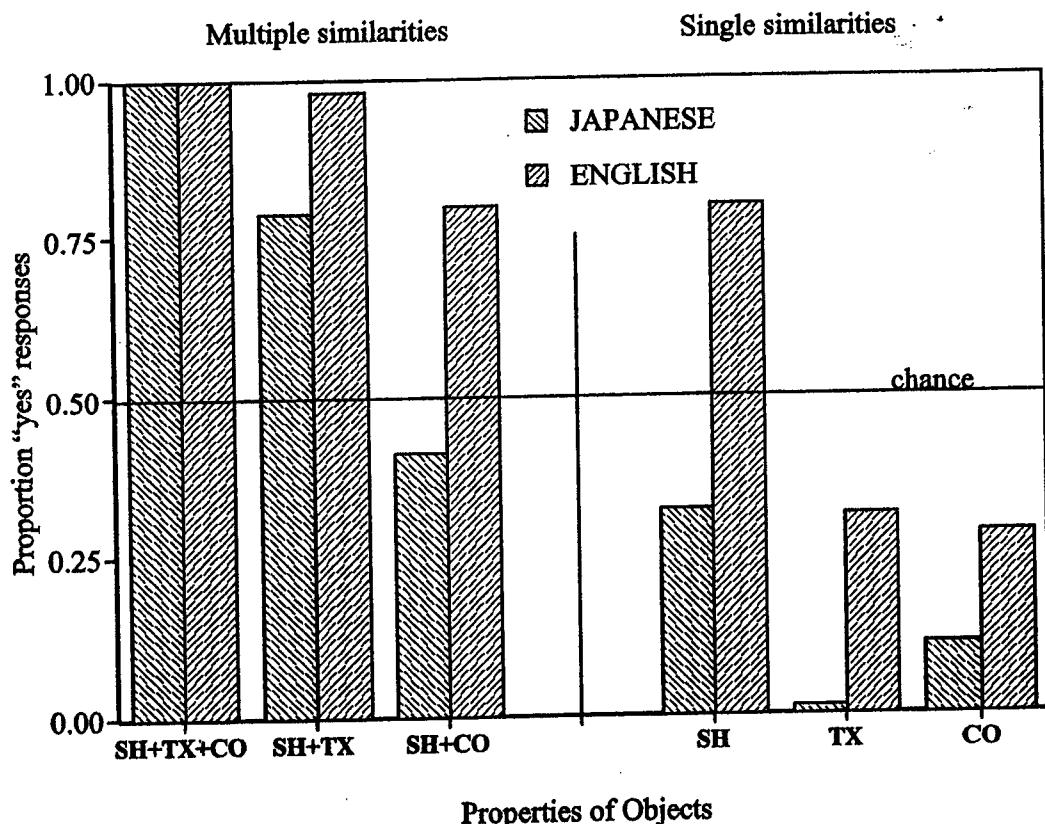

Figure 11.8. Mean "yes" responses—the name of the exemplar extends to the test object—by Japanese-speaking and English-speaking children, when the name was presented in a neutral syntactic frame, as a function of the properties of the test object that match those of the named exemplar: shape (SH), texture (TX), color (CO).

categories organized by shape alone. The result of these added linguistic cues to the correlational structure is that Japanese-speaking children are more likely than English-speaking children to see the appendages as limbs and the objects as depictions of an animate kind.

Cross-language Differences Only at the Boundary

Yoshida and Smith (in press, 2001) also compared Japanese- and English-speaking children's generalizations of names for unambiguous depictions of animates (rounded bodylike forms with eyes) and unambiguous depictions of artifacts (angular, complex, multipart shapes). With unambiguous exemplars, Japanese- and English-speaking children extended the exemplars' names in the same way. Names for unambiguous depictions of animates were extended to new instances narrowly, by shape and texture. Names for unambiguous artifacts were extended broadly by shape. Thus, the cross-linguistic effects at the animate-object boundary, like those at the object-substance boundary, appear evident only for ambiguous entities that lie near the boundary.

These findings make sense if linguistic cues are *one* influence in a correlational soup that also includes perceptual cues and learned lexical category structures. If perceptual cues strongly predict (and perhaps determine) lexical category structure, as seems to be the case in the early noun lexicon, linguistic cues may push conceptualizations one way or the other only in perceptually ambiguous cases. This is an important idea for thinking about how language-specific structures might influence the formation of even more abstract ideas.

Summary

Children learning all languages are presented with three sources of information relevant to forming "ontological" distinctions. These are illustrated in figure 11.9. First, there are the different kinds of things in world—from formless liquids to deformable substances to simple wood to complex artifacts and natural kinds to animate things. These different kinds present—*statistically and in a graded way*—different perceptible properties. Second, there are also the lexical categories that children are learning. The similarity structure of these categories is statistically correlated with the perceptible properties of different kinds. Liquids that have no shape of their own may tend to be named by material (and color), artifacts with rigid and stable shapes may tend to be named by shape, and animates with their rich correlational structures may be named by multiple properties that include texture and shape. Third, there are linguistic devices specific to specific languages that correlate with these regularities in perceptual properties and category organizations. If children are associative learners, and if all these sources of information are blended together in a learned ontology, then one would predict both universals and a coherent bending of those universals in language and culturally specific ways. This is what the pattern of results here suggests. Children learning English and children learning Japanese learn to carve up the world in the same way because of deep regularities in that world and because the two languages organize lexical categories of concrete kinds in pretty much the same way.

Perceptual properties	water sand chunk of wood cup car tree doll cat people

Lexical categories	mostly material-based mostly shape based multiple similarities

Linguistic contrasts	--- mass syntax-------- /-------------count syntax---- -----------aru------------------ /---iru-----------

Figure 11.9. Three sources of information about animates, objects, and substances.

How These Correlations May Build Abstract Ideas

The early lexicon is small and is not representative of the full range of nouns that children ultimately learn. For the most part, the early noun lexicon is filled with names for the concrete, palpable things that dominate domestic life. The adult lexicon includes, in addition, names for abstract ideas, ideas that sometimes also seem to be divided into abstract categories of animate, object, and substance. For example, speakers of English pluralize and count "hopes" as if hopes were bounded and discrete kinds. Speakers of English, however, do not pluralize "justice" but speak of meting it out in portions, as if justice were a continuous and unbounded substance. These abstract ideas may be built on or be metaphoric extensions of the perceptual structures of concrete objects and substances (see Lakoff & Johnson, 1980; Levinson, 1996).

Figure 11.10 illustrates hypothesized correlations among perceptual properties and from perceptual properties to lexical category structure. Although not illustrated, it seems likely that these various connections vary in strength, depending on the strength of relations in the world. For example, objects with angles and multiple parts are highly likely to be solid (since complex angular shapes cannot be readily formed from nonsolid substances). Thus angularity strongly predicts solidity and multiple parts, and each of these cues and the whole cluster predicts categorization by shape. Analogously, nonsolid objects tend to be rounded and simply shaped, although many simply shaped things can also be solid. Thus, simple shape and roundedness weakly predict nonsolidity and categorization by material, but simple shape, roundedness, and nonsolidity would jointly predict more strongly categorization by material. Finally, a strong cluster of interrelated cues would seem to characterize animate things, and all these cues predict categorization by multiple similarities. The correlations in figure 11.10 derive from the perceptual regularities in the world, regularities that appear to be honored in the category structures of the common concrete nouns of both English and Japanese.

What do the differences between English and Japanese languages add to these perceptual correlations? As illustrated in figure 11.11, perceptual properties and category structures characteristic of animates are also associated with particular linguistic forms in Japanese and perceptual properties and category structures characteristic of inanimates are associated with contrasting forms. Figure 11.12 illustrates how perceptual properties and category structures characteristic of animates and objects are
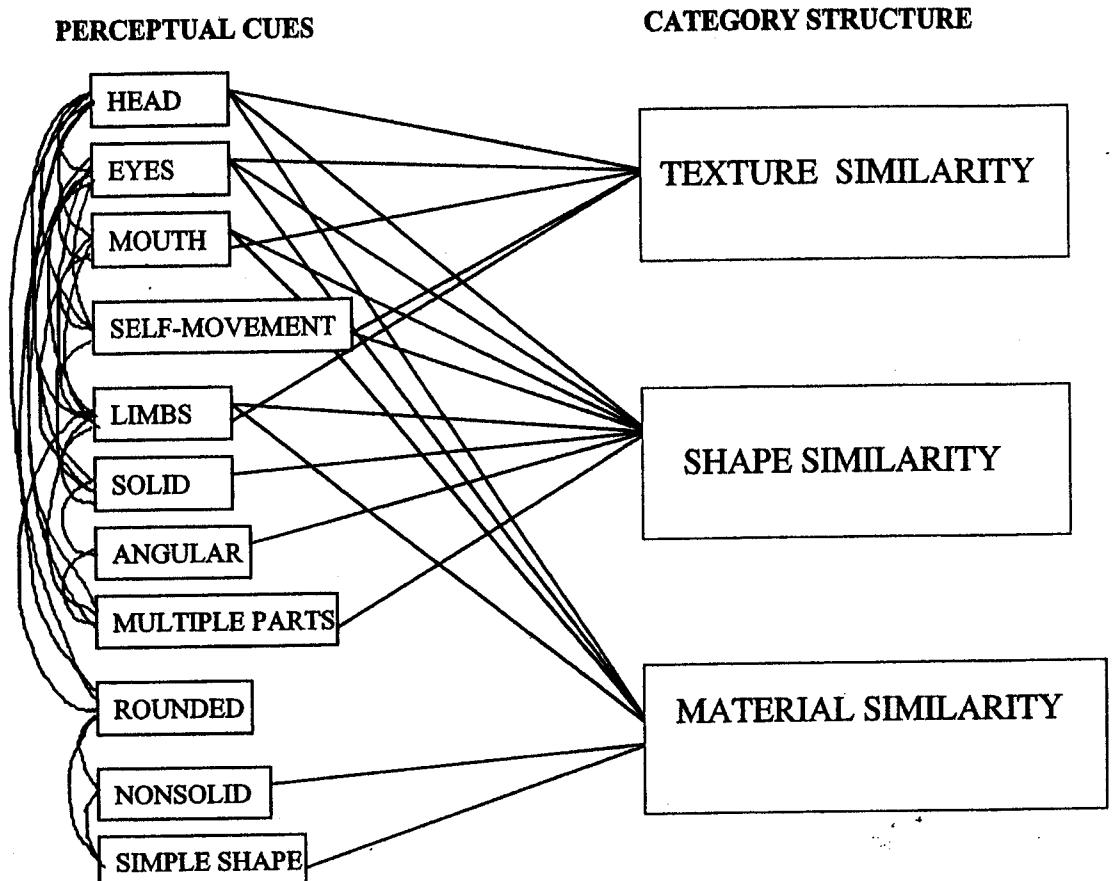
PERCEPTUAL CUES

CATEGORY STRUCTURE

HEAD
EYES
MOUTH
SELF-MOVEMENT
LIMBS
SOLID
ANGULAR
MULTIPLE PARTS
ROUNDED
NONSOLID
SIMPLE SHAPE

TEXTURE SIMILARITY

SHAPE SIMILARITY

MATERIAL SIMILARITY

Figure 11.10. Correlations between perceptual properties and lexical category structure.

also associated with particular linguistic forms in English and how perceptual properties and category structures characteristic of nonsolids are associated with contrasting forms. One can also see in these illustrations how the addition of linguistic cues to a name generalization task can influence children's name generalizations; how in Soja's (1992) study, saying *a mel* increased English-speaking children's generalizations by shape, whereas saying *some mel* increased their generalizations by material, and how in Yoshida and Smith's (2000) study, saying *iru* increased Japanese-speaking children's generalizations by shape and texture, whereas saying *aru* increased their generalization by shape alone.

Importantly, however, systematic linguistic contrasts do more than just shift attention on-line. The evidence suggests that they also differentially bolster and weaken perceptual correlations, changing, in a sense, how things are perceived. Figures 11.11 and 11.12 illustrate how this may be so in an associative learner. The interconnections among perceptible cues associated with animacy—head, eyes, limbs, self-movement—may be strengthened by their joint association with linguistic forms in Japanese. Because of their connections to the same cluster of linguistic cues, the feature "limblike appendages" may be more strongly linked to self-movement and to eyes for Japanese speakers than for English speakers. The implication is that for Japanese speakers, vaguely suggestive limbs—because of reinforcing connections pro-
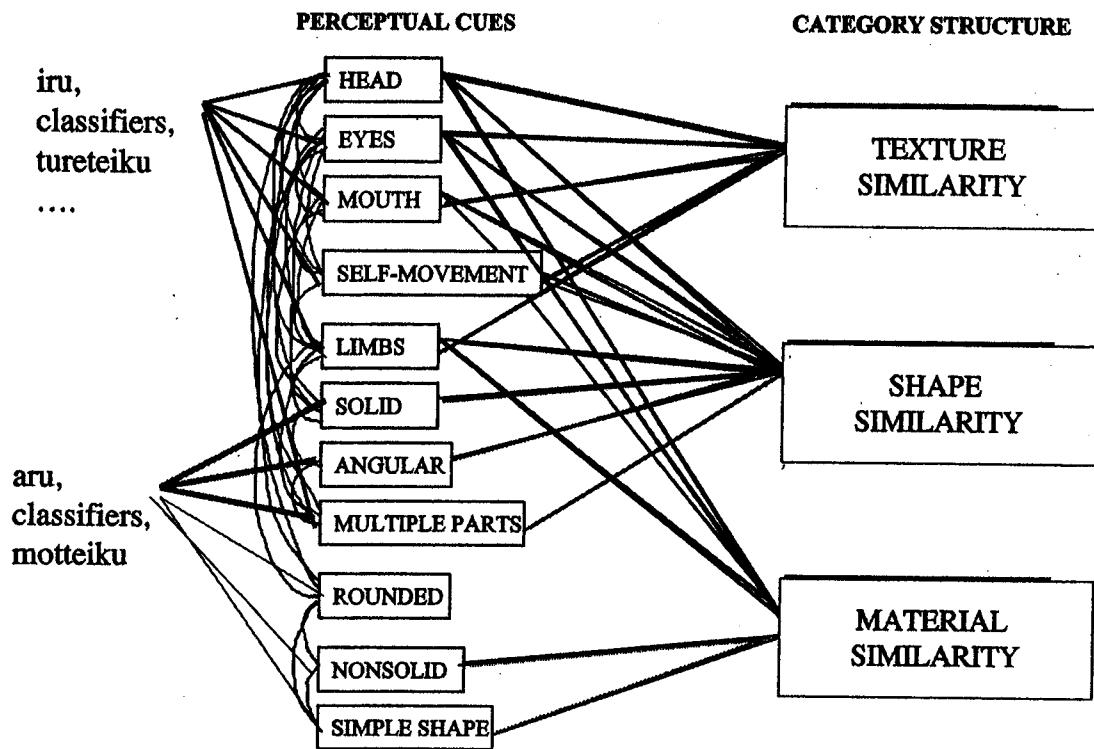
PERCEPTUAL CUES   CATEGORY STRUCTURE

iru,
classifiers,
tureteiku
....

aru,
classifiers,
motteiku

HEAD
EYES
MOUTH
SELF-MOVEMENT
LIMBS
SOLID
ANGULAR
MULTIPLE PARTS
ROUNDED
NONSOLID
SIMPLE SHAPE

TEXTURE
SIMILARITY

SHAPE
SIMILARITY

MATERIAL
SIMILARITY

Figure 11.11. Correlations among Japanese linguistic cues, perceptual properties, and lexical category structure.

vided by the Japanese language—may be more likely to bring forth ideas associated with animate things, including categorization by multiple properties. Thus, vaguely limblike appendages may be a stronger cue suggestive of animacy for Japanese than English speakers. Analogously, the linguistic forms in English that signal discrete countable things may reinforce the connections between cues that are characteristic of objects and between those cues and categorization by shape. Thus, even in tasks in which those linguistic cues are not present, solidity—even in the context of a simple shape—may robustly lead to categorization by shape. Although speculative, these ideas fit the general workings of interactive-activation models of associative learning (Billman & Heit 1989; Colunga & Gasser 1998; Kersten & Billman 1997; McClelland & Rumelhart 1981): overlapping connections reinforce each other such that one cue alone can bring forth activation of a whole correlated cluster.

Intriguingly, the strengthened connections that are the consequences of these so-called gang effects in associative learning may play an important formative role in abstract ideas. Ideas of animacy or objectness that do not depend on perceptual cues may emerge through links from linguistic cues to category structures. If the relations illustrated in figures 11.11 and 11.12 capture the regularities that actually exist, then the most basic assumptions of associative learning predict that linguistic cues like the indefinite article should give rise to ideas of boundedness and that *iru* should give rise to ideas of self-movement. In this way, hopes may be abstract objects and spirits may have intention.
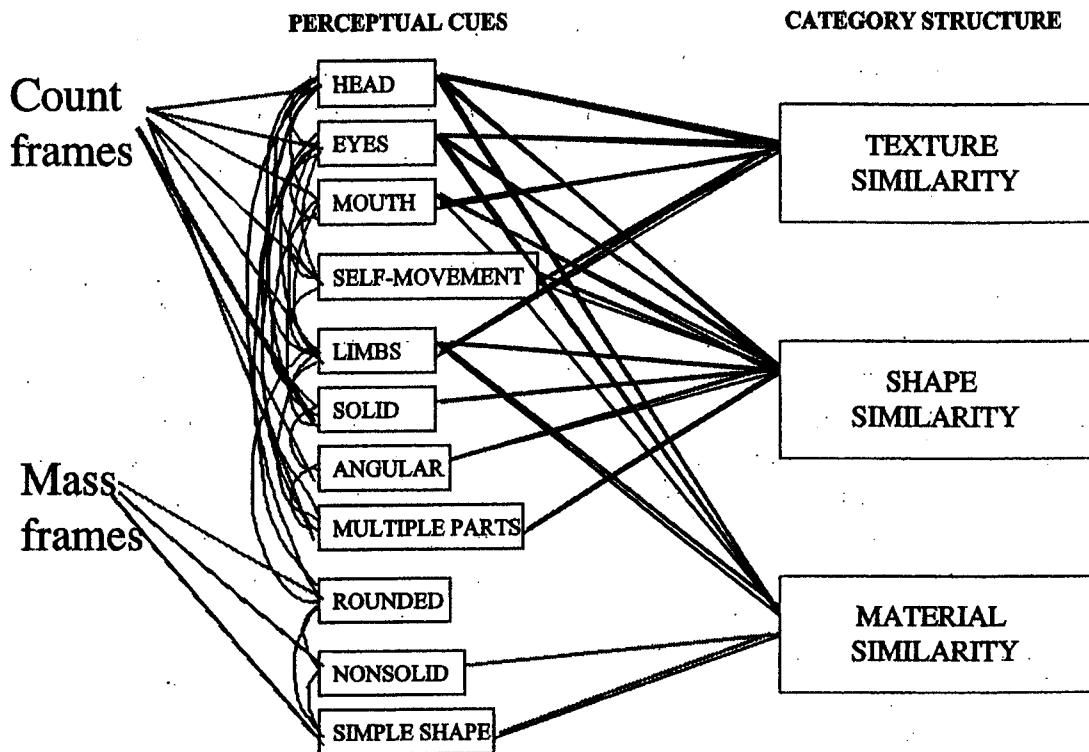
PERCEPTUAL CUES   CATEGORY STRUCTURE

Count frames

HEAD
EYES
MOUTH
SELF-MOVEMENT
LIMBS
SOLID
ANGULAR
MULTIPLE PARTS
ROUNDED
NONSOLID
SIMPLE SHAPE

Mass frames

TEXTURE SIMILARITY
SHAPE SIMILARITY
MATERIAL SIMILARITY

Figure 11.12. Correlations among English linguistic cues, perceptual properties, and lexical category structure.

Conclusions

The evidence presented in this chapter provides empirical support for five core ideas, as follows:

1. *There are regularities that distinguish kinds of categories.* Solid things can be complexly shaped, nonsolid things cannot, and animate things are characterized by bundles of correlated properties.
2. *The nominal categories of languages honor these correlational bundles.* Concrete nouns in both English and Japanese—the nouns learned early—name complexly shaped solid things by shape, nonsolid things by material, and animate things by multiple similarities, including similarities in shape and texture.
3. *Learning names for things enables higher-order generalizations.* Simple associative devices, when taught pairings between names and individual object categories, learn more than just how those trained names map to categories. They also learn the correlations that characterize different kinds, for example, how object categories are structured differently from substance categories.
4. *The mechanism is associative learning.* The simulation studies clearly demonstrate how ontologies in the psychological sense could arise

naturally from the correlational bundles in the world and the regularities across lexical categories that are mapped to those bundles.

5. *Linguistic regularities are part of the correlational mix that creates ontologies, and thus language-specific properties will bend psychological ontologies in language-specific ways.*

These ideas of ontologies as statistical regularities suggest a profound sameness in all human knowledge. They also suggest the genuine possibility that there are culturally distinct ways of knowing. Universality will be found amid the correlations and statistical regularities that are grounded in perception, the structure of the world, and in concrete lexical categories. Diversity, unique ways of knowing specific to specific cultures, will arise from variations in how the systematic contrasts in a language correlate with early-learned statistical regularities and will show itself most dramatically in ideas about abstract kinds. Both universality and diversity are the natural products of the statistical regularities among properties of concrete things, their category structures, and the exquisite variations in how languages reflect and extend deep truths about concrete kinds.

Acknowledgments

References

Billman, D., & Heit, E. (1989). Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science, 12,* 587–625.

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

Colunga, E., & Smith, L. B. (2000) Learning to learn words: a cross-linguistic study of the shape and material biases. *Boston University Child Language Development Proceedings, 24,* 197–207.

Fenson, L., Dale, P., Reznick, J. S., Bates, E., Hartung, J., Pethick, S., & Reilly, J. (1993). *MacArthur Communicative Development Inventories.* San Diego: Singular.

Gathercole, V., & Min, H. (1997). Word meaning biases or language-specific effects? Evidence from English, Spanish and Korean. *First Language, 17,* 31–56.

Gathercole, V., Cramer, L., Somerville, S., & Jansen, M. (1995). Ontological categories and function: Acquisition of new names. *Cognitive Development, 10,* 225–251.

Gelman, S. A., & Coley, J. D. (1991). Language and categorization: The acquisition of natural kind terms. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 146–196). Cambridge: Cambridge University Press.

Gentner, D. & Boroditsky, L. (in press). Individuation, relativity, and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development.* England: Cambridge University Press.

Gumperz, J. C., & Levinson, S. C. (1996). Introduction to part 1. In J. C. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 21–36). Cambridge: Cambridge University Press.

Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. *Cognition, 62,* 169–200.

Imai, M., Gentner, D., & Uchida, N. (1994). Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive Development, 9,* 45–76.

Jones, S., & Smith, L. B. (1998). How children name objects with shoes. *Cognitive Development, 13,* 323–334.

Jones, S. S., Smith, L. B., & Landau, B. (1991). Object properties and knowledge in early lexical learning. *Child Development, 62,* 499–516.

Keil, F. (1994). Explanation, association, and the acquisition of word meaning. In L. R. Gleitman & B. Landau (Eds.), *Lexical acquisition* (pp. 169–197). Cambridge, MA: MIT Press.

Kersten, A., & Billman, D. (1997). Event category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 638–658.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago: University of Chicago Press, 1980.

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

Landau, B., Smith, L. B., & Jones, S. S. (1992). Syntactic context and the shape bias in children's and adult's lexical learning. *Journal of Memory and Language, 31,* 807–825.

Landau, B., Smith, L. B., & Jones, S. (1998) Object shape, object function, and object name. *Journal of Memory and Language, 38,* 1–27.

Levinson, S. C. (1996). Relativity in spatial conception and description. In J. J. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 177–202). Cambridge: Cambridge University Press.

Luce, R. D. (2000). *Utility of gains and losses: Measurement-theoretical and experimental approaches.* Mahwah, NJ: Erlbaum.

Lucy, J. A. (1992). *Language diversity and thought: A reformulation of the linguistic relativity hypothesis.* Cambridge: Cambridge University Press.

Lucy, J. A. (1996). The scope of linguistic relativity: An analysis and review of empirical research. In J. C. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 37–69). Cambridge: Cambridge University Press.

McClelland, J., & Rumelhart, D. E. (1981). An interactive-activation model of context effects in letter perception: An account of basic findings. *Psychological Review, 88,* 373–407.

McPherson, L. (1991). A little goes a long way: Evidence for a perceptual basis of learning for the noun categories COUNT and MASS. *Journal of Child Language, 18,* 315–338.

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

Ogura, T., Yamashita, Y., Murase, T., & Dale, P. (1993, April). *Some preliminary findings from the Japanese Early Communicative Inventory.* Paper presented at the Sixth International Conference for Child Language, Toronto.

Ogura, T., & Watamaki, T. (1997). *Japanese communicative developmental inventories: User's guide and technical manual.* San Diego: Singular.

Quine, W. V. (1969). *Ontological relativity and other essays.* New York: Columbia University Press.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 111–144). San Diego: Academic Press.

Sandhofer, C. M., Smith, L. B., & Luo, J. (2000). Counting nouns and verbs in the input: Differential frequencies, different kinds of learning? *Journal of Child Language, 27,* 561–585.

Samuelson, L., & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category structure, and syntax correspond? *Cognition, 73,* 1–33.

Samuelson, L., & Smith, L. B. (in press). Attention to rigid and deformable shape in naming and non-naming tasks. *Child Development.*

Smith, L. B. (1999). Children's noun learning: How general learning processes make specialized learning mechanisms. In B. MacWhinney (Ed.), *The emergence of language* (pp. 277–303). Hillsdale, NJ: LEA.

Soja, N. (1992). Inferences about the meanings of nouns: The relationship between perception and syntax. *Cognitive Development, 7,* 29–46.

Soja, N., Carey, S., & Spelke, E. (1991). Ontological categories guide young children's inductions of word meanings: Object terms and substance terms. *Cognition, 38,* 179–211.

Spelke, E. S., Vishton, P., & Von-Hofsten, C. (1995). Object perception, object-directed action, and physical knowledge in infancy. In M. S. Gazzaniga (Eds.), *The cognitive neurosciences* (pp. 165–179). Cambridge, MA: MIT Press.

Yoshida, H., & Smith, L. B. (in press). Shifting ontological boundaries: How Japanese- and English-speaking children generalize names for animals and artifacts. *Developmental Science.*

Yoshida, H., & Smith, L. B. (2001). Early noun learning depends on the language being learned, *Cognition, 82*(2), B63–B74.