Theory-based Correlations and Their Role in Children's Concepts

Susan E. Barrett
Lehigh University

Hervé Abdi
University of Texas at Dallas and University of Bourgogne at Dijon

Gregory L. Murphy
University of Illinois

Jeanette McCarthy Gallagher
Lehigh University

Barrett, Susan E.; Abdi, Hervé; Murphy, Gregory L.; and Gallagher, Jeanette McCarthy. *Theory-based Correlations and Their Role in Children's Concepts.* Child Development, 1993, **64**, 1595–1616. Recent accounts of conceptual development have emphasized the important role intuitive theories play in concept formation; however, it is still not clear exactly how these theories exert their influence. We present evidence that elementary school age children use theories to link together specific features associated with individual concepts. The results of our first experiment indicate that theory-based correlations play a prominent role in typicality judgments and in decisions about category membership. In a second experiment, we demonstrate that children's theories play an important role in determining which attributes will be considered most central to the concept. The results of these studies suggest that feature correlations can serve to link children's concepts with their intuitive theories of the world.

Young children have rich theories and beliefs about the world, and these must have some impact on the nature of their concepts. But exactly how do these naive theories or intuitive beliefs exert their influence? Traditional accounts of concept formation have for the most part ignored this question. Instead, the focus has been on structural accounts of concept learning, often with an implicit assumption that most, if not all, of the features comprising a concept are immediately given in the perceptual array. In recent years, however, a new framework for considering conceptual development has begun to emerge as investigators come to the realization that even basic level concepts are not "perceptually transparent" but instead must be constructed with reference to a theory or cognitive model (see Keil, 1989; Murphy & Medin, 1985; and chapters in Neisser, 1987).

Although perceptual similarity surely plays a role in decisions about category membership, even toddlers can move beyond surface similarities in their classification of objects (Mervis, 1987). Mervis's son, Ari, had an early vocabulary that included labels for various waterfowl, such as "duck" and "swan." The offspring of these species look fairly similar, and not surprisingly, Ari

This research was supported by a National Academy of Education Spencer Postdoctoral Fellowship and a BRSG grant 2-S07 RR07173 awarded to the first author. We are grateful to Dorrit Billman, Barbara Malt, and several anonymous reviewers for their thoughtful comments on an earlier version of this paper. Thanks are also due to the members of the MLU group for their helpful discussions of the issues raised in this paper. Requests for reprints should be sent to Susan Barrett, Department of Psychology, Lehigh University, 17 Memorial Drive East, Bethlehem, PA 18015-3068.

[*Child Development*, 1993, **64**, 1595–1616. 0009-3920/93/6406-0009$01.00]

initially labeled all the young "ducks." But by 22 months of age, Ari assigned the appropriate names to the young birds based on his knowledge of the parent bird. For example, if he saw a young bird floating beside an adult bird, he concluded that the two birds were members of the same species. Thus, at a relatively young age, Ari seemed to possess an intuitive theory that parents and offspring are members of the same species regardless of their appearances, and this theory played a dominant role in decisions about category membership.

Keil's (1986) studies provide further evidence that theories affect classification. Keil presented elementary school children with objects that were transformed in appearance from one ontological category to another. For example, in one story, a porcupine was changed to look like a cactus. When confronted with these apparent ontological changes, even kindergartners judged that each item retained its original identity despite considerable changes in appearance. Keil argued that this belief is rooted in children's theories about natural kinds. Children believe that natural kinds possess essential properties that cannot be altered, and this belief takes precedence over perceptual information in decisions about category membership.

Both Keil's and Mervis's work demonstrate that children do not form categories solely on the basis of perceptual information. Instead, children interpret the information given in the perceptual array with reference to the theories they hold about the world. A second way in which theories might influence concepts is by restricting the kinds of inferences that can be drawn on the basis of category membership. Even preschoolers selectively generalize properties on the basis of category membership rather than perceptual similarity (Carey, 1985; S. Gelman & Markman, 1986). For example, after learning that a particular bird feeds its young mashed up food, S. Gelman and Markman found that children were willing to credit a perceptually dissimilar bird with the same feeding habit because both animals belong to the bird category. The same children recognized that other attributes, such as weight, do not generalize across category members. Apparently their theories tell them that members of the same species tend to share eating habits but not weight.

Studies focusing on patterns of induction for natural kinds and artifacts have further shown that elementary school children reason differently about objects in these two domains. When informed of a new property, second graders are more likely to draw inferences within natural kind categories than artifact categories (S. Gelman, 1988; S. Gelman & O'Reilly, 1988). For example, with natural kinds, second graders expect a new property to be true of other members of the basic level category but are less willing to make this inference with artifacts. In contrast, preschoolers are equally likely to generalize novel attributes across members of natural kind categories and artifact categories. These results suggest that as children's theories become more elaborate, they begin to recognize that living things and artifacts differ in the richness of the inferences they promote (S. Gelman & Wellman, 1991).

Changes in children's theories can lead to changes in how widely attributes are generalized across members of a category. Carey (1985) believes that 4-year-olds lack an independent theory of biology; instead, their understanding of biological processes is rooted in their intuitive psychology and centers around what they know about people.[1] One consequence of this is that they do not make a clear distinction between essential biological properties (e.g., eating) and nonessential biological properties (e.g., having bones): animals are credited with both kinds of properties based on how similar they are to people. By age 10, however, children make this distinction and attribute essential biological functions to all animals. Carey argues that this shift in attribution patterns reflects a change in theories. Whereas preschoolers' biological knowledge is embedded in a human-centered theory of behavior, older children possess an intuitive theory of bi-

[1] Carey believes that preschoolers lack a theory of biology, but given that preschoolers have some conception of biological growth and inheritance, a number of investigators have argued that preschoolers' understanding of animals is embedded within an intuitive biological theory (e.g., Inagaki & Hatano, 1987; Rosengren et al., 1991; Springer & Keil, 1991), rather than a social theory of human behavior as Carey (1985) claims. Preschoolers' biological theories, however, are primitive at best, and children of preschool age seem unlikely to see the relevance of biological principles in situations where biological functioning is not stressed (Vera & Keil, cited in Keil, 1989).

ology in which the body is likened to a machine that functions to support life. Acquiring this theory produces fundamental changes in the child's concepts of living things and their properties.

One reason the theory-based approach to conceptual development has been received with such enthusiasm is that it offers a new framework for considering whether there are fundamental differences in the concepts of the child and the adult. If children and adults have different intuitive theories, then we would expect these differences to result in qualitatively different concepts. Carey's work on conceptual change suggests that this is true. Using the theory construct as a way to explain conceptual change, however, raises a number of difficult issues. Undoubtedly, the most central of these is what constitutes a theory.

Drawing on work in the philosophy of science, Wellman (1990) makes a distinction between framework theories and specific theories that may help psychologists clarify what they mean by the term "theory" (see also Barrett, Abdi, & Sniffen, 1992). Framework theories are broader than specific theories; they define the ontology of the domain and place limits on the kinds of information and causal mechanisms that specific theories can incorporate. As a result of these constraints, the specific theories within each domain are likely to have their own unique flavor. For example, specific theories within the child's framework theory of mind are likely to incorporate information about an actor's intentions, whereas specific theories within the physical domain would be limited to mechanistic explanations for behavior.

The contrast between framework and specific theories may provide a useful starting point for considering how theories influence concepts. Framework theories focus the learner on the causally relevant properties that structure a domain (see Brown, 1990; R. Gelman, 1990), and consequently bias feature selection and weighting during concept acquisition. In some cases, these causally relevant properties may embody feature correlations. For example, correlations between structural properties and functional attributes are likely to be intimately connected with the set of principles that organize the domain. But it is also true that many feature correlations are more specific in nature. Perceived correlations between education and income, dizziness and earaches, and rate of speech and intelligence, to name a few (see Murphy & Medin, 1985; Pazzani, 1991, for other examples), do not really typify the core causal principles of a domain, but instead, seem to embody the more specific theories that are formulated within the constraints of the framework theory.

In sum, there is a complex relation between the child's theories and his or her concepts. Children's theories enable them to move beyond surface properties in decisions about category membership and also help determine which properties can be extended to other category members. An intriguing proposal that emerges from these findings is that the emergence of new framework theories can lead to fundamental changes in the child's concepts. In addition, subtle differences between the concepts of the child and the adult might be a function of age-related changes in specific theories. In the present experiments, we explore how both framework theories and specific theories affect children's concepts, and, more specifically, whether children take special note of feature pairs that can be easily connected through theory-based explanations.

Several influential accounts of concept learning have recognized the importance of feature correlations. For example, Rosch (1978) pointed out that features tend to cluster together in the environment, and that basic level concepts have evolved to take advantage of this correlational structure (see Malt & Smith, 1984, for evidence supporting this view). Rosch's insight has given rise to an important question: How is the correlational structure of the environment represented in our concepts? One possibility suggested by Rosch's work is that a concept is represented by its central tendency or prototype. Prototypes are often characterized as a listing of the typical features associated with a concept (see Smith & Medin, 1981), and, in a limited sense, prototypes represent correlations by listing features together. For example, the prototypical features of a bird are all correlated to some extent. Most things with wings and feathers also fly and live in nests. But prototypes by themselves do not represent all that we know about feature correlations. For example, all features listed in the prototype are not equally correlated (Malt & Smith, 1984). Possessing wings seems to be more strongly correlated with flying than with living in a nest, but this information is not conveyed in the feature list. One way to remedy this situation would be to add information about the strength of each feature correlation (Smith & Medin, 1981), but this would quickly result in an explosion of feature links, and the efficiency of the prototype representation would be lost.

Another more fundamental problem with this solution is that a simple probabilistic linking between features does not seem to capture the essence of feature correlations (Murphy & Medin, 1985). Statistical properties do not predict which feature correlations people will notice. Adults are not very good at detecting arbitrary feature correlations in a concept learning task even when the features are perfectly correlated (Murphy & Wisniewski, 1989). Instead, the extent to which feature correlations are noticed seems to be largely a function of prior expectations that are likely to reflect an individual's intuitive theories about the world (Alloy & Tabachnik, 1984; Wright & Murphy, 1984). An adequate account of feature correlations must show how such explanations can be incorporated into the conceptual representation.

Theories can provide a means for linking together the features within a concept because most theories include information about the interdependence or causal connection between features (Murphy & Medin, 1985). But these correlations need not be explicitly stored within a theory—instead, theories may provide a context for generating specific explanations. Feature correlations, then, might be linked together through pre-existing or newly constructed specific theories. Regardless of when they were created, however, these explanations may help unite the features of a concept. It is also worth pointing out that specific theories need not always be self-generated. In many situations, explanations for feature correlations are introduced by parents and teachers (Callanan & Oakes, 1989), and we expect that when these explanations are consistent with the child's framework theories, they also play a central role in shaping the child's concepts.

In the present set of experiments, we explore whether children use their theories to link together feature pairs and whether these theory-based correlations are an integral part of the concept. We test these claims using artificial concepts that contain features that can be linked together through the child's intuitive theories. Our prediction is that objects possessing theory-based correlations will be considered among the most representative examples of a particular concept, whereas exemplars in which the theory-based correlation is explicitly broken will be judged poor examples of the concept. This prediction is based on the view that theory-based explanations will link together specific feature pairs and become an integral part of the concept. As past studies (e.g., Anglin, 1977; Mervis, 1987) have demonstrated that young children are often reluctant to extend category labels to less typical members of the category, we also expected to find that children are more willing to extend category membership to exemplars that preserve theory-based correlations.

Experiment 1

One of the issues the first experiment addresses is whether feature pairs that can be linked together through theory-based explanations play a central role in defining concepts. In this context, we will also consider whether the extent to which children have developed elaborate specific theories in a domain may affect the likelihood that they will notice theory-based correlations. In past work, fundamental differences in the concepts of the child and adult have generally been attributed to changes in the framework theories that encompass the targeted domain. For example, Carey (1985) has argued that children's concepts of "living thing," "person," and "animal" change over time as these early concepts are encompassed by a newly differentiated biological theory. However, major changes in framework theories such as this one are probably not frequent occurrences in development, and when they do occur, they are likely to be accompanied by fundamental changes in how the child deals with a variety of problems. For example, in the "theory of mind" literature, it has been argued that a shift from a copy-based theory of knowledge to a constructivist theory underlies the child's newfound success on false belief tasks, appearance-reality problems, and level 2 perspective-taking tasks (see, e.g., Flavell, Green, & Flavell, 1990; Gopnik & Astington, 1988; Wellman, 1990).

Clearly, fundamental changes in framework theories play an important role in conceptual change. Refinements in a framework theory are likely to have less pervasive effects but may still be an important source of developmental differences. When children and adults subscribe to the same framework theory, there may be differences in the kinds of specific theories they generate. Children's theories might differ from adults' in terms of the number of variables they incorporate or in the salience of specific attributes. Developmental differences in specific theories may also reflect differences in the amount of domain-specific knowledge children possess. Chi and her colleagues' (Chi & Koeske, 1983; Gobbo & Chi, 1986) studies on child dinosaur experts suggest that as children learn more about dinosaurs, their knowledge structures become more integrated and cohesive. One consequence of this is that children begin to notice the different correlations that exist among various attributes, for example, that plant-eating dinosaurs tend to have either armor or horns to protect themselves against predators. Presumably, the same framework theory encompasses the dinosaur knowledge of the novice and expert, and yet, only the expert is able to generate specific explanations for these feature correlations.

In our first experiment, we focused on three interrelated issues. The first issue concerns possible developmental differences in the specific theories children possess about the relation between the brain and cognitive functioning. The second issue concerns whether children spontaneously notice theory-based correlations, even when these features do not co-occur more frequently than other attribute pairs. The third issue concerns the relative importance of theory-based correlations, that is, the extent to which children see these features as central to the concept.

Like their older counterparts, young elementary school children associate the brain with higher cognitive functions and believe the brain plays an active role in manipulating information (see, e.g., Wellman, 1990). These similarities between children's and adults' conceptions of the brain suggest that this knowledge is encompassed by the same framework theory. However, even though younger and older children appear to subscribe to the same framework theory, there are likely to be differences in the specific theories they possess. In general, it seems that younger children are less likely to focus on physical aspects of the brain and the relation between physiological attributes and cognitive functioning. Whereas older children consider the brain an essential component of the nervous system, younger children simply associate the brain with specific cognitive functions (Carey, 1985; Johnson & Wellman, 1982). Young elementary school children recognize that the brain is needed for remembering (Johnson & Wellman, 1982), but in contrast to older elementary school children, they are not likely to consider how the structure of the organ makes it possible for it to accomplish this function (Crider, 1981). Yussen and Kane's (1985) work on children's conceptions of intelligence also suggests that memory skills may be more closely tied to intellectual capacity and possibly "brain power" in the older child. In sum, the available literature suggests that elementary school children's understanding of the relation between the brain and cognitive processes is in a state of flux and that older children may be more inclined to posit a connection between physical aspects of the brain and cognitive functioning. These differences can be conceptualized at the level of specific theories. Once children begin to focus more on the structure of specific organs, they may generate specific theories that relate structure and function, and these theories would result in the relevant knowledge base becoming more integrated and cohesive.

The first experiment explores how these expected differences in the specific theories of younger and older children might influence the nature of their concepts. In our experiment, children were introduced to two types of birds: One had a large and complicated brain and an excellent memory, whereas the other had a small and simple brain and very limited memory abilities. Given older children's more sophisticated understanding of the relation between physiological structures and functional attributes, we expected them to be more sensitive to the connection between brain structure and memory capacity.

Each of the bird categories that was used in Experiment 1 was defined by four features, and only two of these features were theory-correlated, which means that these features could be linked together through an intuitive theory. The theory-correlated features did not always co-occur in the familiarization items—they co-occurred exactly as often as the "theory-neutral" features did. Thus the theory-correlated features were not statistically correlated to a greater extent than any other feature pair. Moreover, children were not cued in any way as to the presence of these theory-based correlations. The first experiment examines whether children spontaneously notice such feature pairs.

The third question addressed in this study concerned whether theory-correlated

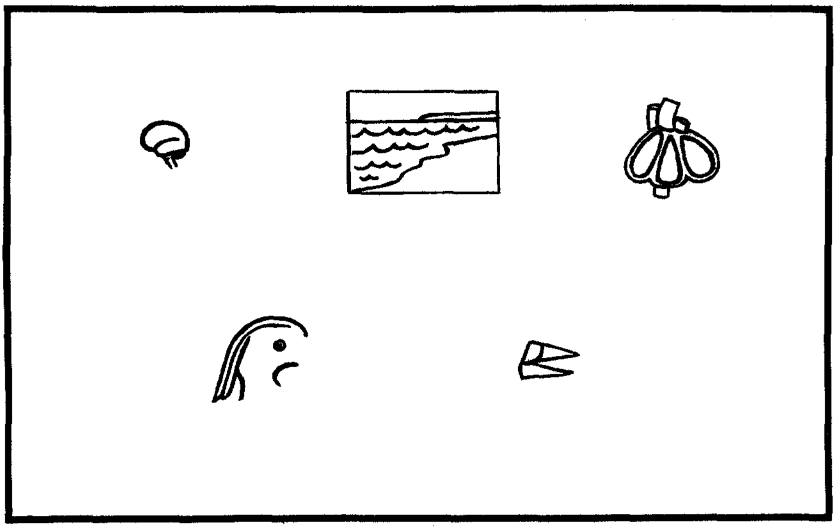

FIG. 1.—One of the familiarization exemplars used in the first experiment. The experimenter described the five features associated with the bird while the child looked at the card. This bird was introduced as a Loppet that has a small and simple brain, lives near the ocean, has a three-part heart, head feathers that point down, and a pointed beak.

features are especially central to a concept. Two experimental tasks were used to evaluate this possibility, a classification task and a typicality task. In both tasks, children were presented with two types of test items. Half of the test items preserved the theory-based correlation by including both theory-correlated features, whereas in the remaining items, the theory-based correlation was broken: A theory-correlated feature from one category was paired with the accompanying feature from the contrasting category. For example, in one broken correlation item, the feature "has a small and simple brain" was paired with "can remember all the different places it has found food." We will call these two types of items CORR and BCORR (correlated and broken correlation), respectively. Both the CORR and BCORR items contained the same number of typical and atypical features; the only difference between them was whether the theory-based correlation remained intact or was broken. Thus, the typicality of the items, as measured by their family resemblance (Rosch & Mervis, 1975), was identical for the preserved-correlation and broken-correlation conditions.

Before detailing our predictions and procedures, it may be helpful to explain some of the factors that influenced our choice of stimulus materials. The familiarization and test items consisted of lists of features. As Callanan's (1985) work demonstrates, parents frequently use verbal statements to draw children's attention to properties. In our study, we presented each feature visually with an oral description. Each feature was depicted in a single picture, and the five pictures associated with each exemplar were presented on a single stimulus card (see Fig. 1). This procedure made it possible for us to include information about internal attributes that are not typically visible and also helped avoid a problem that arises when all the features are represented in a single picture. As Kemler (1981) has pointed out, when exemplars are presented as unitary figures composed of multiple features, there is no guarantee that the features the experimenter used in constructing the stimuli are the same features subjects extract from the figure. By presenting and labeling each feature separately, we hoped to avoid this problem.

In this first experiment, the structure of children's concepts was probed by comparing their performance with different test items. If children use their theories about the world to link together the features of a concept, then exemplars that violate these theory-based correlations should prove harder to classify and should be perceived as less typical category members. If older children are sensitive to the connection between brain structure and memory capacity,

TABLE 1

THE CONSISTENT AND RANDOM FEATURES USED IN EXPERIMENT 1

	Bird Categories	
	Jassler	Loppet
Consistent features:		
Feature 1.[a]	Has a big and complicated brain	Has a small and simple brain
Feature 2.[a]	Can remember all the different places it has found food	Can only remember the very last place it found food
Feature 3.	Has a two-part heart	Has a three-part heart
Feature 4.	Has a rounded beak	Has a pointed beak
Random features for both categories:		
	Has head feathers that point down	Has head feathers that point up
	Has straight tail feathers	Has curly tail feathers
	Lives in the mountains	Lives near the ocean

[a] Theory-based correlations.

but younger children are not, then only the older children should differentiate between the CORR items that preserve the theory-based correlation and BCORR items in which the theory-based correlation is broken.

METHOD

Subjects

Twelve first graders (M = 6-10; range 6-4 to 7-5; four males and eight females) and 12 fourth graders (M = 9-11; range 9-7 to 10-4; nine males and three females) from a middle-class suburban elementary school participated.

Materials

Two contrasting bird categories, labeled "jasslers" and "loppets," were constructed. Four different consistent features were associated with each category. These features included information about the bird's brain, the bird's memory capacity, the bird's heart, and the shape of the bird's beak. The exact features used are given in Table 1; the first pair of features (brain size and memory capacity) represents the theory-based correlation. Features involving the feathers on the bird's head, its tail feathers, and its habitat were also included and appeared equally often with both categories. (These are labeled random features in the table.)

Familiarization items.—During the familiarization phase of the experiment, children were introduced to 16 exemplars of each category. Each exemplar contained a mixture of consistent and random features. Only three of the four consistent features were present in any one item, making these features sufficient but not individually necessary for category membership. Each consistent feature appeared equally often across the familiarization items for the relevant category. Hence, no special status was given to the theory-correlated features in the familiarization phase: The theory-correlated features co-occurred exactly as often as the "theory-neutral" features did. As a result, each theory-correlated feature frequently appeared without its complement during the familiarization phase. To make the categories "fuzzier," each familiarization exemplar also contained two random features. These random features appeared equally often with both categories (e.g., "has straight tail feathers" occurred with both jassler and loppet exemplars). Thus, in learning the categories, the child needed to distinguish between the consistent and random features and had to keep track of which consistent features were associated with each concept.

Test items.—The test phase of the experiment consisted of two parts, a classification task and a typicality task. For the *classification task,* the focus was on whether the presence of a theory-based correlation facilitates categorization. Children were presented with 16 new exemplars, half of which contained the theory-based correlation (CORR) and half of which explicitly broke or violated the theory-based correlation. The broken correlation (BCORR) items were created by pairing a theory-correlated feature from one category with a companion feature from the contrasting category—for example, by combining "has a small and simple

TABLE 2

THE TEST ITEMS FOR THE CLASSIFICATION TASK IN EXPERIMENT 1

	Jasslers					Loppets			
Item	F1	F2	F3	F4	Item	F1	F2	F3	F4
High family resemblance items:									
CORR	J	J	J	L	CORR	L	L	L	J
CORR	J	J	L	J	CORR	L	L	J	L
BCORR	J	L	J	J	BCORR	L	J	L	L
BCORR	L	J	J	J	BCORR	J	L	L	L
Low family resemblance items:									
CORR	J	J	L	...	CORR	L	L	J	...
CORR	J	J	...	L	CORR	L	L	...	J
BCORR	J	L	J	...	BCORR	L	J	L	...
BCORR	L	J	...	J	BCORR	J	L	...	L

NOTE.—Features that were characteristic of jasslers are labeled J in the table and features that were characteristic of loppets are labeled L. Features 1 and 2 are theory-correlated. Only the consistent features are included in this table.

brain" with "can remember all the different places it has found food." The test items for the classification task are presented in Table 2. Note that with the CORR items, the theory-based correlation remains intact, but the theory-neutral features are split. In contrast, with the BCORR items, the correlation between the theory-neutral features is preserved, whereas the correlation between the theory-relevant features is broken. Thus the crucial difference between the CORR and BCORR items is whether the feature correlation that is preserved is theory-relevant or theory-neutral; family resemblance is held constant across the two conditions. As the table makes clear, children's classification decisions under two conditions of family resemblance were also examined: The high family resemblance items contained three features consistent with one category, one feature from the contrasting category, and one random feature, while the low family resemblance items were composed of two consistent features, one contrasting feature, and two random features.

A *typicality task* was included to assess whether exemplars that preserve a theory-based correlation are perceived as representative of a particular category. In this task, children were presented with eight pairs of items (see Table 3). Each pair consisted of two exemplars from the same category, and the child had to indicate which item was more typical of the category. Each exemplar contained three consistent features, one inconsistent feature, and one random feature. Thus family resemblance was held constant across the pair. For one member of the pair, the consistent features included the theory-based correlation (CORR), and the other member contained one theory-correlated feature from each category, thereby breaking the correlation (BCORR). A similar forced-choice technique has been used by Malt and Smith (1984) to demonstrate sensitivity to feature correlations.

Procedure

At the beginning of the session, children were told they would be learning about two new kinds of birds and that they would see pictures to help give them an idea of what each bird was like. Each exemplar was depicted by five feature pictures placed in a 3-2 arrangement on a 5 × 8 index card. The feature pictures were randomly assigned to the five positions on their respective cards. During the familiarization phase, children were alternately shown an exemplar from one category followed by an exemplar from the contrasting category. The starting category was counterbalanced across subjects. For each exemplar, the features were pointed out and then the child was told which category the item belonged in. The child then placed the card in the box designated for that category.

The familiarization phase of the experiment was followed by the classification task. Children were presented with a single exemplar and asked which category they thought it belonged in, for example, whether

TABLE 3

THE TEST ITEMS FOR THE TYPICALITY TASK IN EXPERIMENT 1

CORR EXEMPLAR					BCORR EXEMPLAR			
F1	F2	F3	F4		F1	F2	F3	F4
"Which is more like the Jasslers?"								
J	J	J	L	vs.	J	L	J	J
J	J	J	L	vs.	L	J	J	J
J	J	L	J	vs.	J	L	J	J
J	J	L	J	vs.	L	J	J	J
"Which is more like the Loppets?"								
L	L	L	J	vs.	L	J	L	L
L	L	L	J	vs.	J	L	L	L
L	L	J	L	vs.	L	J	L	L
L	L	J	L	vs.	J	L	L	L

NOTE.—The characteristic features for the jasslers are labeled J and the characteristic features for the loppets are labeled L. Only the consistent features are listed in the table. Each exemplar consisted of four consistent features and one random feature (not shown). Thirty-two unique exemplars were created by varying the identity of the random features.

they would call it a "jassler" or a "loppet." After children completed the classification task, they were presented with the typicality task. Here, two exemplars were placed side by side and the child was asked to point to the one that seemed most like a given category (e.g., "Which is more like the jasslers?"). The items within both the classification and typicality tasks were presented in a different random order for each subject. Children were tested individually in an empty classroom in a 25–30-min session.

RESULTS

The classification task measured the accuracy of children's judgments about category membership. Recall that the test items for the classification task consisted of either two or three features that were typical of one category and one feature that was typical of the contrasting category (see Table 2). The number of typical features determined the correct answer for each item. The number of correct responses was analyzed using an ANOVA for a mixed factorial design; both type of correlation (CORR vs. BCORR) and family resemblance (high vs. low) were within-subject factors and age (first graders vs. fourth graders) was a between-subjects factor. The age effect was significant, $F(1, 22) = 8.58$, $p < .01$; the fourth graders were correct, on the average, on 69% (SD = 8.8) of the items, whereas the first graders classified only 54% (SD = 17.0) of the items correctly. Only the fourth graders' overall performance was significantly above chance, $t(11) = 7.82$, $p < .001$. As predicted, items containing a theory-based correlation were identified more accurately than items in which the correlation was broken, $F(1, 22) = 40.93$, $p < .001$. Children were correct, on the average, on 75% (SD = 13.6) of the CORR items but performed at chance with the BCORR items (48% correct; SD = 18.5). Contrary to our expectation, however, this effect did not vary with the age of our subjects, $F(1, 22) = 0.14$, and the family resemblance structure of the items did not have an effect, $F(1, 22) = 0.79$.

Although the data provide strong support for our view that children are more likely to extend category membership to novel exemplars that preserve theory-based correlations, an alternative explanation needs to be ruled out. Ward and his colleagues (Ward, Becker, Hass, & Vela, 1991; Ward, Vela, Perry, Lewis, Bauer, & Klint, 1989) have shown that under certain conditions children tend to rely on single attribute rules when learning ill-defined concepts. If children based their decisions on either one of the theory-correlated attributes, their overall performance would be better with the CORR items than the BCORR items since they would be using an attribute present in more of the CORR items. Looking at the individual patterns of results, it is indeed possible to find two first graders and three fourth graders whose behavior is consistent with this interpretation; all of their decisions were consistent with one of the single attribute rules. Single attribute rules, however, do not adequately account for the data from the remaining subjects. When each child's decisions are evaluated in relation to the single attribute rules encompassing the theory-correlated features and the proportion of responses consistent with the optimal decision rule for each subject are calculated, these optimal single attribute rules account for only an average of 62% (SD = 22.2) of children's decisions. To be sure that the five single attribute responders were not causing the significant type of correlation effect that was observed, these children were eliminated and an analysis was performed on the data from the 19 remaining subjects. The most important finding to emerge from this second analysis was that the effect of type of correlation remained significant, $F(1, 17) = 30.99$, $p < .001$. Thus, the difference between the CORR and BCORR items is not simply due to some children's having adopted a single attribute rule. In this second analysis, the age effect was no longer significant, but because the three fourth graders and two first graders that were dropped from the analysis were using essentially the same strategies, the lack of significance can be attributed to a smaller sample size. Once again, the family resemblance manipulation did not have an effect.

The data from the typicality task were analyzed using a similar approach. The responses from the five single attribute responders were eliminated from the data set. Each of the remaining children's responses was coded as a series of 1's and 0's, 1 when the CORR member of the pair was considered to be more typical, and 0 when the BCORR member was chosen as more typical. The data were analyzed using an ANOVA for a mixed design, with type of correlation and question ("more like the jasslers" vs. "more like the loppets") as within-subject factors and age as a between-subjects factor. The type of correlation effect proved significant, $F(1, 17) = 4.70, p = .05$; on average, children picked the CORR item as more typical on 61% (SD = 4.8) of the trials. No other main effects or interactions reached significance.

Taken together, the results from these two tasks suggest that children perceive theory-based connections between feature pairs, and that these features do not have to be statistically correlated for this effect to emerge. Our findings with the first graders suggest that it is also not necessary that the relevant theoretical domain be highly elaborated. So whereas prior research (Crider, 1981) suggests that first graders tend not to focus on the physical aspects of an organ when they consider its functioning, if they are presented with information about a simple relation between an increase in brain size and memory capacity, they link together these attributes.

Discussion

The results of the first experiment suggest that children are sensitive to theory-based correlations. Even though the theory-correlated features did not co-occur with greater frequency than any other feature pair, children were better at classifying exemplars that preserved the theory-based correlation than they were at classifying exemplars that did not preserve this correlation. Children also considered exemplars that preserved the theory-based correlation to be more typical of the category. Although the first experiment provided critical support for our view that children use their intuitive theories to link together specific feature pairs, the expected interaction between age and type of correlation was not observed. We had initially thought that fourth graders' more sophisticated understanding of how the structure of an internal organ is related to its functioning would result in fourth graders evidencing greater sensitivity to the connection between brain structure and memory abilities than first graders. However, in retrospect, we would suggest that a simpler specific theory is sufficient for linking these features. If, as Wellman (1990) argues, children conceive of the brain as a storehouse for ideas and memories, then they might think a larger brain affords more storage space. Hence, the size information alone may have been sufficient for connecting memory capacity with brain type, and information about brain complexity may have seemed redundant or irrelevant. Moreover, even if the relation between brain size and memory ability was not explicitly represented in younger children's theories, it seems likely that first graders could easily explain it once they encounter it.

Experiment 2

In the second experiment, we again looked for evidence that children's concepts vary with their theories, but this time we manipulated the framework theories that children would be likely to draw on when learning novel concepts. We believed that the results would be especially powerful if we could show that the same set of features is represented differently depending on whether or not the feature pair embodies a theory-based correlation.

Experiment 2 tested this possibility by focusing on two domains that children know a great deal about: animals and artifacts. A wide range of studies has shown that infants and preschoolers distinguish between animate and inanimate objects on the basis of appearance and activity (R. Gelman, 1990; R. Gelman, Spelke, & Meck, 1983; Golinkoff, Harding, Carlson-Luden, & Sexton, 1984; Jones, Smith, & Landau, 1991) and that they recognize that only animals are capable of self-generated movement (Massey & R. Gelman, 1988). Moreover, even preschoolers expect animals and artifacts to have different internal structures (R. Gelman, 1990; S. Gelman & O'Reilly, 1988), and they recognize that people can make artifacts but not animals or other natural kinds (S. Gelman, 1988; S. Gelman & Kremer, 1991; Keil, 1989). Preschoolers also realize that people can impart certain properties to artifacts, for example, by painting a can blue, but that the color of an animal depends on something the mother contributes (Springer & Keil, 1991).

Although preschoolers recognize that biological and nonbiological kinds are fundamentally distinct, the theories that encompass these domains are not highly developed. For example, although preschoolers understand that animals, but not artifacts, can become larger with time (Rosengren, S. Gelman, Kalish, & McCormick, 1991), they fail to recognize that changes in surface characteristics have different effects on the identity of animals and artifacts (Keil, 1989). Second graders, however, are sensitive to these differences, and this tendency is even more pronounced in fourth graders.

In the second experiment, groups of third graders were presented with either two artifact or two animal categories. All the categories were defined by lists of properties that could be sensibly applied to either tools or animals (e.g., "found in the desert"). Half of the children were told that they were learning about animals, whereas the remaining children were told the items were tools. Both groups of children were presented with exactly the same stimuli throughout the experiment; the only difference was the domain children were told was relevant. For children who were told the items were animals, the relevant theory-based correlation involved the animal's habitat and a property that would be adaptive in this environment. For example, one animal lived in the desert and could hold a lot of water, whereas the other lived in the mountains and was covered with thick wool. When the items were introduced as tools, a different pair of features was theory-relevant. The tool that was found in the mountains could crush rocks, whereas the tool that was found in the desert could cut cactus plants. For the tools, information about the item's water capacity or covering was considered to be theory-neutral, and for the animals, the neutral information included whether it could crush rocks or cut cactus plants.

In this experiment, each category was defined by five consistent features, and only features that could be sensibly applied to both animals and tools were included. One consequence of this is that the five consistent features used to define the categories in the second experiment did not cohere as well as the features associated with the bird categories in Experiment 1 and so were expected to be more difficult to learn. The features that were used in Experiment 2 were pretested to ensure that young elementary school children would be sensitive to the relevant theory-based correlations. The details of this procedure are described below. In the experimental task, children were presented with a wide range of items so that we could evaluate a number of decision rules children might use when classifying novel exemplars. Most of the test items pit a decision based on the theory-correlated features for the animal category against a decision based on the theory-correlated features for the tool category. Recall that in the first experiment, five children adopted single attribute decision rules. In the present experiment, we included a larger set of test items so that we could more thoroughly test whether individual children adopted this strategy. To test the relative importance of the theory-based features we also included a subset of test items that pitted the theory-based correlation against three consistent, but theory-neutral, features from the contrasting category.

Method

Subjects

Twenty-four third graders (M = 8-8; range 8-0 to 9-8; 16 males and 8 females) from a middle-class suburban elementary school participated in the main experiment.

TABLE 4

NUMBERS OF CHILDREN SELECTING EACH FEATURE IN THE PRETEST

LOCATION	ASSOCIATED FEATURE				
	Wool	Water	Rocks	Cactus	Undecided
Animals:					
Mountains	11	4	5	1	3
Desert	0	18	0	3	3
Tools:					
Mountains	0	7	10	1	6
Desert	0	7	2	13	2

Materials

Pretest.—Twenty-four undergraduates and 24 third graders who did not take part in the main experiment helped pretest possible stimuli. The purpose of the pretest was to find a cluster of three features such that one pair of features would embody a theory-based correlation when the items were introduced as animals and a different pair would be theory-relevant when the items were described as tools. An additional constraint was that all three features needed to be plausible attributes for both animals and tools. Sets of possible feature pairs were generated that appeared to meet these constraints, and these served as the stimuli for the initial pretest with undergraduates. The items that were judged as most promising based on the undergraduate data were then pretested with the third graders. The pretest questions were presented in a multiple choice format and administered in a group setting. Each question introduced the item as either an animal or machine,[2] and one property of the item was described. Subjects were then presented with four additional attributes and asked to select the one they thought was most likely to be true of the item; they could also indicate that none of the answers seemed appropriate. Sample items are presented in the Appendix, and the first item includes the features we based the experimental stimuli on. The pretest data for the features we used in the main experiment are presented in Table 4. When an object was introduced as an animal, children associated different properties with animals found in the mountains and in the desert, $\chi^2(4) = 25.91$, $p < .001$. With the machines, children's judgments also depended on where the machine was found, $\chi^2(4) = 17.62$, $p < .05$.[3] More important, for both the animals and tools, the most frequent response within each condition was the property that preserved the hypothesized theory-based correlation.

Familiarization items.—Each category was defined by a set of five features that were uniquely associated with it during the familiarization phase. The categories were ill-defined, meaning that no single feature was necessary and sufficient for category membership. The categories were labeled "numsters" and "zibbots," and the features associated with each category are presented in Table 5. When the items were introduced as animals, features 1 and 2 were theory-correlated: Animals living in the mountains would benefit from a wool coat, and it would be advantageous for animals living in the desert to be able to store a lot of water. When the items were introduced as tools, features 1 and 3 were theory-correlated. A tool that cuts cactus plants would be especially useful in the desert, and a tool for crushing rocks would be more useful in the mountains. During the familiarization phase, children were presented with five examples of numsters and five examples of zibbots. Four of the five consistent features were present in each exemplar, and each feature appeared four times in the familiarization phase. Once again, it is important to note that the theory-

[2] These items were labeled "machines" in the pretest, but we switched to the more common term "tools" in the main experiment.

[3] Given the relatively small expected frequencies in some cells, it could be argued that a Fisher's exact test is more appropriate than a chi-square (see Kendall & Stuart, 1979). In our case, when the data are examined using Fisher's exact test, the estimated probability values are very close to the values given by the standard chi-square analysis.

TABLE 5

THE FIVE CONSISTENT FEATURES USED IN EXPERIMENT 2

Numster	Zibbot
Feature 1. Is found in the mountains	Is found in the desert
Feature 2. Has thick wool	Can store a lot of water
Feature 3. Can crush rocks	Can cut cactus plants
Feature 4. Is long and straight like an arrow	Is round like a plate
Feature 5. Catches snakes	Catches spiders

correlated features co-occurred exactly as often as the theory-neutral features did. To keep the design more manageable and make the categories easier to learn, random features were not used in this experiment.

Test items.—All children completed a classification test consisting of 22 items. These test items were designed to make it possible to determine if children's decisions depended on whether the stimuli were animals or tools and whether children considered the relevant theory-based correlation to be an important part of the concept. In this experiment, a wide range of items was used so that we could more clearly distinguish between children who were sensitive to the theory-based correlation and children who were relying on either a single attribute rule or simply counting the number of features present in the concept. The items that were used in the classification test are presented in Table 6. Within this table, items are grouped together based on how many features were present in the test items.

The first set of items all contain five features. Each item includes three features from one category and two features from the other. Information about whether the stimulus is a CORR or BCORR item is also included in this table. Note that for the five-

TABLE 6

TEST ITEMS FOR THE CLASSIFICATION TASK FOR EXPERIMENT 2

						ITEM TYPE	
ITEM NO.	F1	F2	F3	F4	F5	Animal	Tool
Five-feature items:							
1	N	N	Z	Z	Z	CORR	BCORR
2	Z	Z	N	N	N	CORR	BCORR
3	N	Z	N	Z	Z	BCORR	CORR
4	Z	N	Z	N	N	BCORR	CORR
Four-feature items:							
5	N	N	Z	Z	...	CORR	BCORR
6	N	N	Z	...	Z	CORR	BCORR
7	Z	Z	N	N	...	CORR	BCORR
8	Z	Z	N	...	N	CORR	BCORR
9	N	Z	N	Z	...	BCORR	CORR
10	N	Z	N	...	Z	BCORR	CORR
11	Z	N	Z	N	...	BCORR	CORR
12	Z	N	Z	...	N	BCORR	CORR
13	N	N	...	Z	Z	CORR	...
14	Z	Z	...	N	N	CORR	...
15	N	...	N	Z	Z	...	CORR
16	Z	...	Z	N	N	...	CORR
Three-feature items:							
17	N	N	Z	...	...	CORR	BCORR
18	Z	Z	N	...	...	CORR	BCORR
19	N	Z	N	...	...	BCORR	CORR
20	Z	N	Z	...	...	BCORR	CORR
21	N	Z	Z	...	...	BCORR	BCORR
22	Z	N	N	...	...	BCORR	BCORR

feature stimuli, every item that preserves the theory-based correlation for one domain breaks the correlation for the other. So, for example, Item 1 would be a CORR item for children in the animal condition and a BCORR item for children in the tool condition. Another aspect of the design to note is that within each CORR item, the theory-correlated features are the only two features that are typical of that category; the three theory-neutral features are all associated with the contrasting category. The five-feature CORR items, then, pit the theory-based correlation against the number of typical features.

The next group of items is the four-feature items. Each exemplar in this group contains two features associated with the numsters and two features associated with the zibbots. For Items 5 through 12, an exemplar that is a CORR item for children in the animal group would be a BCORR item for children in the tool group and vice versa. Items 13–16 preserve the theory-based correlation for one domain without violating the theory-based correlation for the other domain. This occurs because only one of the theory-relevant features for the contrasting domain is present in each item. The CORR items were included so we could evaluate whether children based their classification decisions on the theory-correlated features. The corresponding BCORR items were included in the larger set of items that were used to test whether children were relying on a single attribute rule.

The last group of items is the three-feature items. Although we have noted in the table whether each exemplar represents a CORR or BCORR item for the category, because the same decision would be made regardless of whether children relied on the two theory-correlated features or the two consistent features, these items cannot be used to evaluate the relative importance of the theory-based correlation. These items were designed to test whether children are basing their decisions on a single attribute rather than the co-occurrence of the theory-relevant features.

Procedure

The experimental task was modified so it could be administered in a group setting, at the request of the school. Although subjects in group settings are probably not as attentive and motivated as in individual testing situations, we considered the possible consequences of this in our analysis. At the beginning of the group session, children silently read a short story about a boy named Danny who was watching as either animals or tools were sorted into two trucks: one for the numsters and one for the zibbots. To help new workers learn what these categories were like, Danny was asked to write in his notebook four things he knew about each item as it was loaded in the truck. For half of the class, the items were described as animals, and for the remaining children, the items were described as tools.

Each child was given a copy of Danny's notebook. The notebook contained the 10 familiarization items; each item was presented on a separate page. The word "NUMSTER" or "ZIBBOT" was alternately printed at the top of each page, and below this word, four consistent features were listed in a random order. The notebooks for the children in the animal and tool conditions were identical. The experimenter read aloud each page of the notebook as the children followed along in their individual copies. Illustrations were not included in the notebooks as we did not want to influence children's conceptions of the objects as animals or tools.

Immediately following the familiarization period, children completed the classification task. Each child was given a packet consisting of 22 randomly ordered test items. For each item, children followed along as the experimenter read the attributes associated with each test item and then responded to the question "What is it?" by underlining either the word "numster" or "zibbot." Children were tested in a group session lasting approximately 30 min.

RESULTS

Children seemed to have more difficulty with this task than with the task in Experiment 1. The children in this experiment were only briefly familiarized with the two contrasting categories, and familiarization took place in a group setting in which some children may have been less likely to pay attention. In addition, the tool and animal categories in the present experiment were less coherent than the bird categories in Experiment 1, and children were not presented with pictures of any of the objects in the present experiment. As a result, we anticipated that children would be likely to have more problems with the experimental procedures than did the children in Experiment 1. Thus, the first question we asked was whether children were able to give princi-

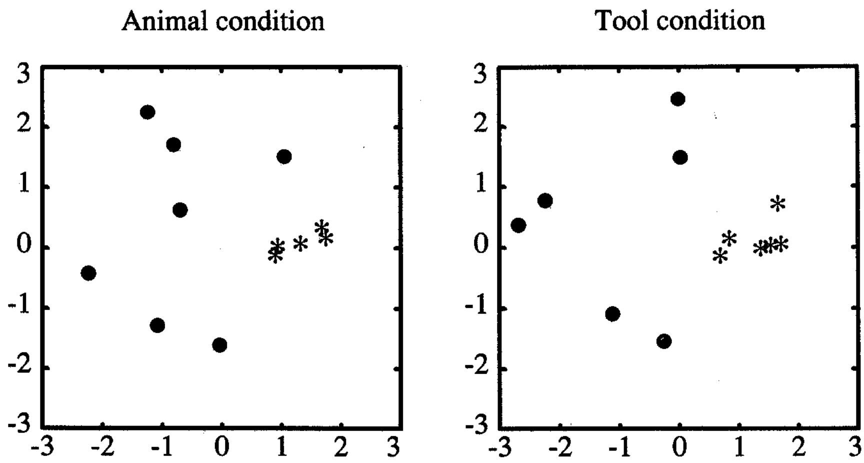

FIG. 2.—The two-dimensional scaling solutions for the children in the animal and tool conditions of Experiment 2. The two-dimensional solution accounts for 56% of the variance in the animal condition and 63% of the variance in the tool condition. The "consistent" subjects in each condition are designated by asterisks.

pled responses in this situation. If the task was indeed too difficult, we expected to find a great deal of variability both within and between subjects. In our first analysis, we focused on the extent to which children within a condition gave similar patterns of responses. To assess this, two intersubject agreement matrices were derived, one for the animal condition and one for the tool condition. For each matrix, intersubject agreement was computed based on the number of times each pair of subjects gave the same answer for a given test item. Separate multidimensional scaling solutions were computed for children in the animal and tool conditions.[4] As can be seen in Figure 2, there were individual differences in response patterns. In the animal condition, five subjects were tightly clustered, whereas the seven remaining subjects were scattered throughout the space. In the tool condition, a similar pattern emerged; six subjects clustered together and six subjects were scattered. These results suggested that it would be helpful to focus on individual patterns of responding.

Our next analysis focused on the response patterns that encompassed the largest number of children in each condition. More specifically, we focused on those five children in the animal condition and six children in the tool condition who agreed on their decisions (i.e., the "consistent" subjects). Recall that in designing our test items, we selected stimuli that would allow us to evaluate the relative importance of the theory-based correlations. We first performed a simple ANOVA to test whether children's responses varied as a function of the test items and domain. Planned contrasts were used to assess the relative importance of the theory-correlated features. For these analyses, children's classification decisions were coded as a series of 0's and 1's (0 if they labeled the item a numster and 1 if they labeled it a zibbot). A two-way ANOVA was performed with domain (animal vs. tool) as a between-subjects variable and stimuli (the 22 test items) as a within-subject variable.[5] There was a significant main effect for stimuli, $F(21, 189) = 10.97$, $p < .001$, meaning that different test items were assigned to the

[4] Both multidimensional scaling (Torgerson, 1952) and additive tree representations (Abdi, 1990; Sattath & Tversky, 1977) were used because these approaches sometimes yield different results (see Abdi, Barthelemy, & Luong, 1984), although in this case the results were similar.

[5] Hsu and Feldt (1969) and Lunney (1970) have investigated the appropriateness of the analysis of variance for random samples from binomial distributions (see Edwards, 1985). Using Monte Carlo simulations, these authors found that when the degrees of freedom are greater than 20, an analysis of variance can be used with a dichotomous variable. Although these conclusions were based on balanced designs, the slight difference in the number of subjects in the animal and tool conditions should not pose problems.

TABLE 7

CLASSIFICATION DECISIONS FOR THE CRITICAL ITEMS IN EXPERIMENT 2

Condition and Item No.	Theory-Correlation Prediction	Family Resemblance Prediction[a]	% of Subjects Responding "Zibbot"
Five-feature items:			
Animal:			
1	Numster	Zibbot	20
2	Zibbot	Numster	100
Tool:			
3	Numster	Zibbot	17
4	Zibbot	Numster	83
Four-feature items:			
Animal:			
5	Numster		40
6	Numster		0
13	Numster		20
7	Zibbot		60
8	Zibbot		100
14	Zibbot		80
Tool:			
9	Numster		17
10	Numster		0
15	Numster		17
11	Zibbot		83
12	Zibbot		83
16	Zibbot		100

[a] The family resemblance model does not offer any predictions for the four-feature items because each exemplar consisted of two features that were associated with the numsters and two features that were associated with the zibbots.

two categories. There was no effect for domain, which indicates that both groups assigned the same proportion of test items to the two categories. The critical domain × stimuli interaction proved to be significant, $F(21, 189) = 2.21$, $p < .001$. Thus, as predicted, subjects in the animal and tool conditions differed in which test items they assigned to the two categories.

A series of planned contrasts was performed to determine the nature of these differences. The first analysis focused on the five-feature items. Recall that these items were included because they make it possible to evaluate whether subjects rely on theory-correlated features when the correlated pair of features predicts membership in one category and the three remaining features are typical of the contrasting category. The predictions outlined in Table 6 are presented again in Table 7 along with the classification decisions made by the consistent subjects, expressed in terms of the percentage of subjects who considered that item to be a zibbot. For children in the animal condition, Items 1 and 2 were considered critical because the theory-correlated features predicted membership in one category whereas the family resemblance structure of the items predicted membership in the other category. For children in the tool category, Items 3 and 4 were considered critical for the same reasons. As can be seen in Table 7, children based their decisions on the theory-correlated features. The contrast analysis confirmed what is readily apparent; children in the animal and tool conditions relied on the theory-correlated features when classifying these items, $F(1, 189) = 23.36$, $p < .001$. (For this analysis, positive weights were assigned to items where a numster response was expected based on the theory-correlated features, and negative weights were assigned to items where a zibbot response was expected.) Additional contrasts revealed that there was no overall effect for domain, nor was there an interaction between the domain and the extent to which children relied on the theory-correlated features.

A second set of contrast analyses was performed on the four-feature items. For children in the animal condition, six items were considered critical because they preserved the theory-correlation, and a different set of six items were considered critical for children in the tool condition for the same reason. These predictions and the data from the consistent subjects are presented in Table 7. The contrast analysis comparing numster and zibbot responses confirmed what is readily apparent; children used the theory-correlated features to classify the items, $F(1, 21) = 61.83$, $p < .001$. Additional analyses revealed that there was no overall effect of domain or interaction between domain and fit. Thus, the results of the contrast analyses suggest that for children in both the tool and animal conditions, the dominant response pattern was consistent with our prediction that children would rely heavily on the theory-correlated features when classifying new items.

Our next task was to determine whether the judgments of any of the remaining children conformed to an identifiable response pattern. First we identified five possible decision rules. Three of these were single attribute decision rules, each of which focused on a single theory-relevant feature (i.e., Feature 1, 2, or 3). In addition to evaluating these three single attribute decision rules, we also looked at how well children's predictions fit a decision rule based on the two theory-relevant features (i.e., a CORR decision rule) and whether they were consistent with a simple family resemblance decision rule. For each of these five decision rules, we computed the proportion of judgments on the relevant items that were consistent with the rule. For the single attribute rule associated with Feature 1, all 22 items were relevant, and for the single attribute rule associated with Feature 2 and the single attribute rule associated with Feature 3, 20 items were relevant. (As Table 6 makes clear, Items 15 and 16 could not be used to evaluate Feature 2 responses and Items 13 and 14 were not relevant for Feature 3 responses.) Ten items were used to evaluate the CORR decision rule for each condition; the relevant items are labeled CORR items in Table 6. Ten items were also used to evaluate the family resemblance decision rule; all the five feature and three feature items in Table 6 were relevant for this comparison.

First consider the data from the animal condition. The decisions of the five consistent subjects conformed to the CORR decision rule. More specifically, 86% ($SD = 5.5$) of their judgments were consistent with this rule. In contrast, the best-fitting single attribute rule for each individual subject accounts for an average of only 73% ($SD = 6.4$) of the judgments. Moreover, for each of these five subjects the CORR decision rule accounted for a higher percentage of their judgments. Clearly, then, these five children were not single attribute responders. Of the remaining children, however, two may have been single attribute responders; for both of these children, the optimal decision rule was a single attribute rule (86% of one child's responses were based on Feature 1 and 75% of the other child's responses were based on Feature 2). One child appeared to have mixed up the category labels; after the labels are interchanged, 100% of his judgments fit the CORR rule and 95% fit a single attribute rule. The data from the remaining four subjects were not on the surface consistent with any of these rules.

Next consider the data from the tool condition. For the six consistent subjects, the degrees of fit to the CORR and Feature 1 decision rules were strikingly similar, 85% ($SD = 15.2$) and 88% ($SD = 10.7$) agreement, respectively. A seventh child showed more variability in his responses than the consistent subjects, but the CORR decision rule provided the best fit to his data; 70% of his judgments fit this rule but for him the next best fit was provided by the Feature 2 rule (60%). Of the remaining five children, two relied on either the CORR features or Feature 1 but mixed up the category labels; for both of these children, if the labels are reversed, all of their judgments are consistent with the CORR rule, and the fit to Feature 1 is 95% for one child and 88% for the other. The responses of the remaining three children did not correspond to any of the decision rules.

Discussion

Experiment 2 provided additional support for our view that children's theories influence the nature of their concepts. In this study, two groups of children were familiarized with the same sets of features, but the way in which they structured this information depended on whether these items were introduced as animate or inanimate objects. Children who were told the objects were animals posited a link between the animal's habitat and a property that would be adaptive in this environment. These findings complement S. Gelman's (1988) work; children expect the properties of natural kinds to be highly correlated, and they draw on their intuitive theories to help them discover the relevant connections. The strong domain effects that emerged in our study suggest that the specific theories children use are constrained by the framework theory that encompasses the domain.

Although the data from the animal condition fit nicely with our predictions, theory-based correlations proved less central in the tool condition, and we were not able to rule out the possibility that children based their decisions entirely on where the tool was found. However, even if this were the case, it is still important to keep in mind that regardless of whether children in the tool condition used a single attribute rule or based their decisions on the two theory-correlated features, it is clear that their judgments about the test items differed from those made by children in the animal condition, even though both groups had been presented with the same stimuli in the familiarization phase. Clearly children's expectations about the features likely to be relevant for tools influenced what they learned about the tool categories.

It should be noted that our experiments did not take advantage of the structure-function relations which are likely to be the best candidates for theory-based correlations in this domain. Even young children appear sensitive to correlations between the structure of an artifact and its function. For example, before extending category membership to new tools and instruments, they consider whether the structure of the objects makes it possible for them to serve the relevant function (Kemler-Nelson et al., 1991). Preschoolers also tend to focus on those attributes or parts that are likely to have functional significance for particular types of artifacts (see Ward, 1990; Ward et al., 1989).

With our stimuli, there was no single functional attribute that distinguished between our two categories. In fact, a single exemplar typically had more than one function associated with it. Moreover, we did not provide structural information to support these functional attributes. If we had provided such information, it is likely that children would have focused on a theory-based correlation involving the structure of the tool and its function.

General Discussion

Concept learning involves more than simply keeping a running tally of which features are associated with which concept. Instead, both children and adults selectively process information based on the kinds of concepts they are learning. The data we have presented suggest that children's theories play an important role in determining which attributes will be selected and how this information will be structured. More specifically, we have shown that children focus on feature correlations that fit with their intuitive theories.

Several researchers have suggested that children tend to focus on some subset of attributes that are either especially relevant for the task at hand (Brown, 1990; Kemler Nelson et al., 1991; Macario, Shipley, & Billman, 1990) or that have proven useful for learning similar concepts in the past (Jones et al., 1991; Ward et al., 1991). These accounts suggest that different attributes may be relevant in different domains. For example, Macario et al. posit that children may focus on color when learning about different types of fruit, which may prove a useful strategy until the child encounters such anomalies as Granny Smith apples. Jones et al. have further shown that these biases differ for animate and inanimate objects. Preschoolers attributed animacy to objects that possess eyes, and they attend more to the texture of these objects than they do to the same objects without eyes. Jones et al. believe this reflects children's sensitivity to general "co-relations" that exist between properties, here, the presence of eyes and the importance of texture variations. But given that eyes signal animacy, their findings might also be interpreted as suggesting domain-specific influences on how the feature space is sampled (Keil, 1987).

The correlations we focused on in our studies involved relations among specific attributes. Although past research has considered whether children are sensitive to multiple attributes (e.g., Ward et al., 1989), these studies have generally not considered why children posit links between some attributes and not others. In the present experiments, we looked more directly at whether theory-correlated features have any special status in children's concepts. In the first experiment, we focused on the connection between physical aspects of the brain and cognitive functioning as past work suggested that elementary school children are likely to possess specific theories relating these attributes. Despite the fact that these theory-correlated features were no more likely to co-occur than other feature pairs during the familiarization phase, children were more likely to extend category membership to exemplars that preserved the theory-based correlation while rejecting exemplars that broke the theory-based correlation. Items preserving the theory-based correlation were also judged to be more representative of the category. In the second experiment, we exposed two groups of children to exactly the same stimuli, but we told one group that they would be learning about animals, and the other that they would be learning about tools. Consistent with our prediction, children were sensitive to those theory-based correlations that fit with the framework theory we had activated; children included the theory-correlated items in the category but excluded items that broke this correlation. In both experiments, children chose to link together feature pairs because their theories suggested a connection between these attributes; children were not bothered by and probably did not notice that all possible feature pairs co-occurred equally often in the familiarization phase. Thus, a statistical correlation between attributes is not critical for linking attribute pairs, but a causal link or theoretical basis for the connection may be necessary. The work on illusory correlations (e.g., Chapman & Chapman, 1967, 1969; Crocker, 1981; Wright & Murphy, 1984) suggests that theory-based correlations remain important even when the relation is contradicted by the statistical evidence.

Our findings are not easily explained by traditional models of concept learning. In particular, most simple prototype theories posit abstraction processes that are insensitive to people's naive theories of the world. The influential theory of Rosch and Mervis (1975), for example, suggests that people process a concept's features independently, keeping a count of how often each feature occurs with a category. The related models of Homa, Rhoads, and Chambliss (1979) and Posner and Keele (1970), which use a wholistic prototype representation (rather than a featural representation—see Smith & Medin, 1981), also fail to predict these effects. In these models, a prototype is formed by a global similarity judgment of the category members—no influence of world knowledge is mentioned, and it is hard to see how it could be incorporated. Finally, "exemplar" theories of concepts (Brooks, 1978, 1987; Kossan, 1981; Medin & Schaffer, 1978) have a similar problem. These theories propose that concepts are represented merely as remembered exemplars, rather than as a more abstract representation. Although exemplars may indeed be important, they do not include the world knowledge necessary to detect theory-based correlations. Furthermore, in our experiments, the features that the children perceived as correlated were not actually correlated in the familiarization exemplars, which suggests that the children did not represent the categories only as exemplars.

In short, major theories of concept learning do not include the information necessary to explain how subjects represent correlations. (Murphy & Medin, 1985, discuss this problem further.) However, we do not wish to suggest that these theories are completely wrong or misguided. Most likely, a complete theory of concept learning will have to include both a prototype-extraction device, perhaps one that learns correlation rules in the absence of external feedback (see Billman & Heit, 1988), and processes that draw on an individual's intuitive theories about the world. Thus, our experiments argue for the insufficiency of traditional accounts, rather than their invalidity (see Murphy, 1992, in press). For example, if one conceives of concepts as lists of features (see Smith & Medin, 1981), our results suggest that knowledge outside the feature list influences which features will be selected and how this list will be structured. Our understanding of how this occurs might be advanced by considering the different ways in which framework theories and specific theories might affect this process. Framework theories, undoubtedly, place some general constraints on this process; they suggest which classes of features should be afforded special status within a domain and also limit the kinds of mechanisms that can be used to explain feature correlations. Many of the explanations for why features are likely to co-occur, however, are probably better conceptualized as specific theories. They seem more restrictive in nature and more malleable than framework theories.

Our findings are also relevant for researchers in semantic and conceptual development who are attempting to specify the constraints that guide the child in acquiring new concepts. We believe that children's theories constrain their concepts by helping them rule out impossible feature combinations. In making this claim, we should point out that we are using a weak version of the term "constraint." Children's theories seem to bias feature selection and storage, but it does not appear that theories completely rule out contradictory feature correlations (e.g., a winged bird that is unable to fly).

Imagine for a moment a child who is faced with the task of learning about an object that embodies a contradictory feature correlation. In the laboratory, the child is likely to consider the object nonsensical and may simply attempt to memorize lists of features. Under these conditions, adults can learn to distinguish categories consisting of contradictory features, though not as well as more reasonable categories (Murphy & Wisniewski, 1989). However, we have found in unpublished experiments that elementary school children simply do not learn contradictory feature lists when taught using the procedures outlined in the first experiment. One tempting conclusion to make from this difference is that children may be more dependent on their world knowledge in learning new concepts than adults are. However, even if this were true, it could be due, in part, to children's memory limitations (which make it difficult to memorize arbitrary features) or other cognitive differences. Clearly, more explicit comparisons between adult and child concept learning are called for (see, e.g., Kemler Nelson, 1984). But what happens when the child is confronted with a real object that embodies a feature correlation that contradicts his or her current theories? Confronted with this apparent contradiction, the child may deny or ignore the feature correlation. Or alternatively, the child may be motivated to seek out new information to explain this correlation, and this may lead the child to modify or qualify his or her present theories. An important question for future research concerns the relationship between seemingly contradictory feature correlations and conceptual change.

In exploring the importance of theory-based correlations, we have found that such correlations are a central part of concepts even in newly emerging theoretical domains. This suggests that theory-based correlations may serve a similar function throughout development. What may change, however, is the relative importance of specific feature correlations. It would be interesting to explore this question in domains such as human reproduction or religion, where children are likely to have strongly held, but unique, theories. It may be that the concepts within these domains change as the encompassing theories change (Carey, 1985, discusses this possibility). If this is true, an important question would be whether it is possible to specify this restructuring of individual concepts by tracing changes in which features are considered to be theory-correlated.

Appendix

This animal is found in the mountains. Do you think it is more likely to:

a. be able to crush rocks
b. be able to cut cactus plants
c. be covered with thick wool
d. be able to store lots of water
e. none of these answers seems right

This machine needs a lot of nuts. Do you think it is more likely to:

a. be helpful for making ice pops
b. be helpful for making peanut butter
c. stay inside all winter
d. get wet and sweaty on the outside
e. none of these answers seems right

References

Abdi, H. (1990). Additive tree representations. *Lecture Notes in Biomathematics*, **84**, 43–56.

Abdi, H., Barthelemy, J. P., & Luong, X. (1984). Tree representations of associative structures in semantic and episodic memory research. In E. Degreef & J. VanBuggenhaut (Eds.), *Trends in mathematical psychology* (pp. 3–31). New York: Elsevier.

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans: The joint influence of prior expectations and current situational information. *Psychological Review*, **91**, 112–149.

Anglin, J. M. (1977). *Word, object, and conceptual development*. New York: Norton.

Barrett, S. E., Abdi, H., & Sniffen, J. M. (1992). Reflecting on representation and process: Children's understanding of cognition. In B. Burns (Ed.), *Percepts, concepts and categories: The representation and processing of information* (pp. 275–322). Amsterdam: Elsevier.

Billman, D., & Heit, E. (1988). Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science*, **12**, 587–625.

Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

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

Brown, A. L. (1990). Domain-specific principles affect learning and transfer in children. *Cognitive Science*, **14**, 107–133.

Callanan, M. (1985). How parents label objects for young children: The role of input in the acquisition of category hierarchies. *Child Development*, **56**, 508–523.

Callanan, M., & Oakes, L. M. (1989, April). *Parent-child collaboration in the development of category inferences and causal theories*. Paper presented at the biennial meeting of the Society for Research in Child Development, Kansas City, MO.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Chapman, L. J., & Chapman, J. P. (1967). Genesis of popular but erroneous diagnostic observations. *Journal of Abnormal Psychology*, **72**, 193–204.

Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, **74**, 272–280.

Chi, M. T. H., & Koeske, R. (1983). Network representation of a child's dinosaur knowledge. *Developmental Psychology*, **19**, 29–39.

Crider, C. (1981). Children's conceptions of the body interior. In R. Bibace & M. E. Walsh (Eds.), *New directions for child development: Children's conceptions of health, illness, and bodily functions* (pp. 49–65). San Francisco: Jossey-Bass.

Crocker, J. (1981). Judgments of covariation by social perceivers. *Psychological Bulletin*, **90**, 272–292.

Edwards, A. L. (1985). *Experimental design in psychological research* (5th ed.). Cambridge, MA: Harper & Row.

Flavell, J. H., Green, F. L., & Flavell, E. R. (1990). Developmental changes in young children's knowledge about the mind. *Cognitive Development*, **5**, 1–27.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and animate-inanimate distinction as examples. *Cognitive Science*, **14**, 79–106.

Gelman, R., Spelke, E. S., & Meck, E. (1983). What preschoolers know about animate and inanimate objects. In D. Rogers & J. A. Sloboda (Eds.), *The acquisition of symbolic skills* (pp. 297–324). London: Plenum.

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

Gelman, S. A., & Kremer, K. E. (1991). Understanding natural cause: Children's explanations of how objects and their properties originate. *Child Development*, **62**, 396–414.

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, **23**, 183–208.

Gelman, S. A., & O'Reilly, A. W. (1988). Children's inductive inferences within superordinate categories: The role of language and category structure. *Child Development*, **59**, 876–887.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understanding of the non-obvious. *Cognition*, **38**, 213–244.

Gobbo, C., & Chi, M. (1986). How knowledge is structured and used by expert and novice children. *Cognitive Development*, **1**, 221–237.

Golinkoff, R. M., Harding, C. G., Carlson-Luden, V., & Sexton, M. E. (1984). The infant's perception of causal events: The distinction between animate and inanimate objects. In L. P. Lipsitt (Ed.), *Advances in infancy research* (Vol. 3, pp. 145–165). Norwood, NJ: Ablex.

Gopnik, A., & Astington, J. W. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development*, **59**, 26–37.

Homa, D., Rhoads, D., & Chambliss, D. (1979). The evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory*, **5**, 11–23.

Hsu, T. C., & Feldt, L. S. (1969). The effect of limitations on the number of criterion scores on the significance level of the *F* test. *American Educational Research Journal*, **6**, 515–527.

Inagaki, K., & Hatano, G. (1987). Young children's spontaneous personification as analogy. *Child Development*, **58**, 1013–1020.

Johnson, C. N., & Wellman, H. M. (1982). Children's developing conceptions of the mind and brain. *Child Development*, **53**, 222–234.

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

Keil, F. C. (1986). The acquisition of natural kind and artifact terms. In W. Demopoulos & A. Marras (Eds.), *Language learning and concept acquisition* (pp. 133–153). Norwood, NJ: Ablex.

Keil, F. C. (1987). Conceptual development and category structure. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 175–200). New York: Cambridge University Press.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Kendall, M., & Stuart, A. (1979). *The advanced theory of statistics: Vol. 2. Inference and relationship*. London: Griffin.

Kemler, D. G. (1981). New issues in the study of infant categorization: A reply to Husaim and Cohen. *Merrill-Palmer Quarterly*, **27**, 457–463.

Kemler Nelson, D. G. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior,* **23,** 734–759.

Kemler Nelson, D. G., Almasy, L., Crowley, K., Duke, N., Gardner, J. A., Kiggins, V., Lasher, K., McQuilken, A., O'Connell, M., Russell, R., Sterner, D., & Tirk, E. (1991, April). *Principle-based inferences in preschoolers' categorization of novel artifacts.* Paper presented at the biennial meeting of the Society for Research in Child Development, Seattle, WA.

Kossan, N. E. (1981). Developmental differences in concept acquisition strategies. *Child Development,* **52,** 290–298.

Lunney, G. H. (1970). Using analysis of variance with a dichotomous variable: An empirical study. *Journal of Educational Measurement,* **7,** 263–269.

Macario, J. F., Shipley, E. F., & Billman, D. O. (1990). Induction from a single instance: Formation of a novel category. *Journal of Experimental Child Psychology,* **50,** 179–199.

Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior,* **23,** 250–269.

Massey, C., & Gelman, R. (1988). Preschoolers' ability to decide whether pictured unfamiliar objects can move themselves. *Developmental Psychology,* **24,** 307–317.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review,* **85,** 207–238.

Mervis, C. B. (1987). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 201–233). Cambridge: Cambridge University Press.

Murphy, G. L. (1992). Theories and concept formation. In I. Van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 173–200). New York: Academic Press.

Murphy, G. L. (in press). A rational theory of concepts. In G. V. Nakamura, R. M. Taraban, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol.* **29.** *Acquisition, representation and processing of categories and concepts.* New York: Academic Press.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review,* **92,** 289–316.

Murphy, G. L., & Wisniewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol.* **2.** *Theory and applications* (pp. 23–45). Chichester, England: Ellis Horwood.

Neisser, U. (1987). *Concepts and conceptual development: Ecological and intellectual factors in categorization.* Cambridge: Cambridge University Press.

Pazzani, M. J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning and Memory and Cognition,* **17,** 416–432.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology,* **77,** 353–363.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology,* **7,** 573–605.

Rosengren, K. S., Gelman, S. A., Kalish, C. W., & McCormick, M. (1991). As time goes by: Children's early understanding of growth in animals. *Child Development,* **62,** 1302–1320.

Sattath, S., & Tversky, A. (1977). Additive similarity trees. *Psychometrika,* **42,** 319–345.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Springer, K., & Keil, F. C. (1991). Early differentiation of causal mechanisms appropriate to biological and nonbiological kinds. *Child Development,* **62,** 767–781.

Torgerson, W. S. (1952). Multidimensional scaling: Theory and method. *Psychometrika,* **17,** 401–419.

Ward, T. B. (1990). The role of labels in directing children's attention. In J. T. Enns (Ed.), *The development of attention: Research and theory* (pp. 321–342). Amsterdam: Elsevier.

Ward, T. B., Becker, A. H., Hass, S. D., & Vela, E. (1991). Attribute availability and the shape bias in children's category generalization. *Cognitive Development,* **6,** 143–167.

Ward, T. B., Vela, E., Peery, M. L., Lewis, S. N., Bauer, N. K., & Klint, K. A. (1989). What makes a vibble a vibble? A developmental study of category generalization. *Child Development,* **60,** 214–224.

Wellman, H. M. (1990). *The child's theory of mind.* Cambridge, MA: MIT.

Wright, J. C., & Murphy, G. L. (1984). The utility of theories in intuitive statistics: The robustness of theory-based judgments. *Journal of Experimental Psychology: General,* **113,** 301–322.

Yussen, S. R., & Kane, P. T. (1985). Children's conception of intelligence. In S. R. Yussen (Eds.), *The growth of reflection in children* (pp. 207–241). New York: Academic Press.