Perceptual Similarity and Conceptual Structure

LINDA B. SMITH
DIANA HEISE
Indiana University

I. INTRODUCTION

Many recent discussions of categories and concepts contrast the perceptual and conceptual bases for categories. These discussions pick up a persistent theme in developmental theory. This theme posits a trend from

perception to conception and it has been played throughout the works of Piaget (1929), Vygotsky (1934/1962) and Werner (1948). It is echoed by Flavell (1970), Wohlwill (1962), and Bruner and Olver (1963) and more recently by Gentner (1989) and Keil (1989). The idea is that children shift from perceptually-bound representations of objects that are global and holistic to ones that are principled and articulated along abstract dimensions. So, for example, Flavell (1970) wrote that conceptual development grows from "equivalences based on the more concrete and immediately given perceptual, situational, and functional attributes of objects to equivalences of a more abstract, verbal-conceptual sort." Since development is directional, this theme implies that conception is in some way "better than" perception.

Many developmentalists argue against a shift from perceptual categories to conceptual ones (e.g., Brown, 1990; Mandler & Bauer, 1988; Wellman & Gelman, 1988). They argue that categories are conceptually structured from the very beginning. But even for these theorists, the research agenda is defined in terms of categories structured by "mere appearance" versus categories with a "rich theoretical structure that goes beyond superficial similarity." In these discussions, perceptual and conceptual categorization are presented as mutually exclusive processes. Moreover, conceptual categorization is presented as smart and mature and perceptual categorization is presented as deficient, and to be abandoned.

We believe that this perception-*versus*-conception organization of research works to the detriment of empirical and theoretical progress. This definition of the research agenda dismisses the contribution of perceptual structure to conceptual structure and does not study it seriously. Our purpose in this chapter is to defend perception. We counter the perception-*versus*-conception approach by arguing that the contribution of perceptual categories to conceptual categories is considerable and continuous throughout development. Conceptual structure does not replace or even override perceptual categories. Instead, conceptual structure is based on perceptual structure.

This view that perception is the grounding force for conception has been argued recently by Johnson (1983) and Lakoff (1987). This view is also gaining attention in the cognitive development literature (see Gelman, 1990; Gentner & Rattermann, 1990; Mandler, 1990). In this chapter, we specifically consider the role of perceptual similarity in category development. Our defense of perceptual similarity rests on three points:

1. Perceptual similarity is dynamic. It varies with the attributes attended to.

2. Experience in perceiving the relations in the world influences attention and perceptual similarity. Because of this, perceptual similarity moves about in meaningful ways.

3. Perceptual categories can be abstract and can convey -- indeed be the source of -- conceptually relevant knowledge.

Before presenting the defense, we assess the case against perceptual similarity.

II. THE CASE AGAINST PERCEPTUAL CATEGORIZATION

Keil (1989) building on the writings of Quine (1970) contrasts an "animal similarity space" and a "conceptual similarity space." Animal similarity according to Keil, is original similarity. It is uninterpreted feature counts and lists of correlations between perceptual features. According to Keil, this animal or original similarity is the *atheoretical* tabulation of information we get when we look at the world. It is constant and unchanging and therefore stupid. Fish and whales look alike regardless of what we know. But what objects look like is not always what they really are. Thus, development is away from "the immediate, the subjective, animal sense of similarity to the remoter sense of similarity determined by scientific hypotheses and posits and constructs" (Quine, 1977). In this view, the one we argue against, perceptual similarity is replaced by belief systems composed of causally related features and nonperceptual properties.

We offer three examples of how ideas such as these organize the current research on categories and concepts.

A. From Basic to Superordinate Categories

Medin and Ortony (1989; see also, Medin et al., 1990) suggest that perceptual similarity although not a very intelligent mechanism may precisely fit the needs of the infant in discovering her first categories. Research on perceptual development (e.g., Kemler Nelson, 1990; Smith, 1989) indicates that very young children often compare multidimensional objects wholistically across all dimensions at once. Medin and Ortony argue that this original similarity space fits the child's task of making a first partitioning of the world's objects into categories.

Medin and Ortony point to the fit between overall perceptual similarity and the basic level categories described by Rosch & Mervis (1975). The categories dog, cat, car, house, and bird are basic level categories and these categories do seem to be organized by perceptual similarity or in Rosch's terms by family resemblance. Dogs *look* alike, or at least most dogs seem to be mostly alike in most of their perceptual properties (see, Rosch, et al. 1976; Biederman, 1985).

In contrast to basic level categories, superordinate categories seem to be structured by a few general abstract properties. Animals, for example, differ widely and are not as a group perceptually like each other. Moreover, by the consensus view, there are no specifiable perceptual properties that all (or even most) animals possess. The consensus view is that superordinate categories are structured by nonperceptual properties (Carey, 1985; Gelman, 1988; Mandler & Bauer, 1988; Markman, 1989).

There is much developmental data consistent with a trend from basic to superordinate categories. In classification tasks, young children form spatial groupings of basic level categories such as shoes versus dogs before they form superordinate groupings (Rosch et al., 1976). Names for basic categories are learned fast and considerably before names for superordinate categories (MacNamara, 1982). These data and arguments suggest a developmental trend that proceeds from Quine's "animal" similarity which gives us basic categories to conceptual similarity which gives us superordinate categories.

However, in their full complexity, the developmental data do not squarely fit a *unidirectional* trend from basic to superordinate categories (see, Mervis, 1987). An early sensitivity to superordinate category structure can be seen in young children's overgeneralizations in word learning. Young children's overgeneralizations honor superordinate category distinctions in that children sometimes mistakenly call cows "doggie" but do not mistakenly call cars "doggie" (see Waxman, 1980, for relevant data). Mandler and Bauer (1988; see also, Mandler, Bauer & McDonough, in press) have shown that even prelinguistic children are sensitive to superordinate categories. They observed 12-month-olds' "categorizations" in a free play situation. Their measure of categorization consisted of the sequence of touching objects. Same category objects were touched in rapid succession more frequently than different category objects (see Sugarman, 1983). By this measure, Mandler and Bauer found that infants readily made superordinate classifications of the sort "dog and horse versus car." Indeed, in their task, 16-month-olds made superordinate classifications more readily than they made basic level ones (e.g., poodle and collie versus horse).

Mandler and Bauer take for granted that perceptual properties alone are not enough to organize superordinate categories. Under this assumption, their data provide support for early conceptually based categorization. Mandler and Bauer point to the categories formed by the children -- chicken, fish, cow, and turtle versus motorcycle, airplane, van, and train engine -- as proof that perceptual similarity was not involved. They point to the way the children played with the toys -- making the animals walk and talk and making the vehicles speed about with "vroom-vroom" sounds as evidence for nascent conceptual structure. Perceptual similarity, they conclude, does not control categorization even at its beginning.

We are not sure that this conclusion is warranted. Perceptual structure may well be sufficient to explain Mandler and Bauer's data. Specifically, there may be heretofore undiscovered perceptual properties that distinguish animals and vehicles, and babies might well be sensitive to these properties. Of course, to find these properties, we have to look. Later, in this chapter, we will present some evidence of such properties that we found when we did look.

B. Perceptual Categorization versus Essences

Keil (1989) recently argued that perceptual similarity is not sufficient to explain the psychological structure of even basic level categories. A thought experiment makes the point. Consider a cow. Is it still a cow if some psychologist covers it in sheepskin, paints it purple, cuts off a leg, and adds moose horns? Keil (1989) has shown that people strongly believe that naturally occurring objects do not change their identity with changes in their perceptual appearance -- no matter how severe. People possess beliefs about naturally occurring objects (or natural kinds) that attribute their identity to the processes of their origin or that imbue them with an undefined essence.

In this research, Keil affirms the distinction between the perceptual processes we use to identify objects in the world and our concepts of them (see Smith & Medin, 1981, for more on this important distinction). The idea is that when we come across some unknown object, we may use perceptual similarity to identify it and assign it to a category. But our categories -- our concepts of what it *means* to be a particular kind of object -- are much more than lists of perceptual or even functional properties. People have organized sets of causal beliefs, or theories, that distinguish kinds of objects according to whether they are naturally occurring or manmade, alive or not, terrestrial or water living and so on (see Carey, 1985; S. Gelman, 1988; & Keil, 1989). According to Keil, it is these

theories about the causal connections between perceptual properties and their origins that form the conceptual core of a category. And thus, alterations in the "mere appearance" of an object in ways (e.g., painting) that do not violate central theoretical beliefs (e.g., an object is a cow if its mother was a cow), do not alter the object's identity. This theoretical core is sometimes referred to as the category's "essence" in recognition of the relation between these views and nominal essentialism in philosophy (see Keil, 1989; and also Murphy & Medin, 1985 and Medin & Ortony, 1989).

Causal theories increase in complexity and scientific accuracy with age (e.g., Carey, 1985; Gelman, 1988). Consistent with this age-related growth in conceptual structure, Keil (1989) has shown that the belief in "category essences" that transcend appearance also increases with age. Three- and 4-year-olds often maintain that an object's identity does change with changes in appearance. Thus, again, the developmental trend is described as moving from perceptual similarity to conceptual structure.

There are a number of ongoing disputes in this literature. They all concern whether perceptual appearance is enough to explain performance for any stimuli at any developmental level. For example, Wellman and Gelman (1988) argue that very young children believe in "essences" too and ignore appearance in domains in which they have sufficient knowledge. Carey, Gelman, and Keil each argue that essences that transcend perceptual appearance are characteristic of peoples' beliefs about naturally occurring objects but not manmade objects. By these arguments, artifacts are characterized by an impoverished theoretical structure and thus their category structures are more controlled by superficial perceptual similarity than are the category structures of natural kinds. Greer and Sera (1990; see also, Mandler, Bauer and McDonough, in press) argue in contrast that there is no principled distinction between artifacts and natural kinds. They report that people have causal theories about the internal parts and workings of complex artifacts such as computers and radios and will maintain that these objects' identities go unchanged despite radical changes in outward appearance.

The research agenda, the theoretical disputes, are all defined in terms of perception versus conception. Is this the most useful definition of the research question? When we come upon an object in the world, when we come face-to-face with a dog or a chair, we use sophisticated perceptual processes to determine what it is (e.g. Biederman, 1985). But if we are to take the data and claims about essences seriously, our concepts of objects are largely unrelated to and largely unaffected by the perceptual systems that have evolved to recognize real objects in the real world.

C. Category Induction

When asked to make inductive inferences from what is known about objects in one category to objects in another, conceptual structure and not perceptual similarity is again said to be the preeminent force. The typical empirical study again pits perceptual and conceptual solutions. The state-of-the field is cogently summarized by a thought experiment derived from the work of Carey and Gelman. Imagine a real monkey, a mechanical monkey, and a real snake. We tell you that the real monkey has a duogleenan inside. Is the mechanical monkey or the real snake more likely to also have a duogleenan inside? A real monkey and a mechanical monkey look alike, but the real monkey and real snake share "deeper" properties that place them in the same superordinate categories.

Carey (1985) and Gelman and Markman (1986, 1987, see also, Gelman, 1988) have shown that people (sometimes as young as three years of age) make inductive inferences about the nonperceptual insides of objects in accord with superordinate category structure. In two important studies, Gelman and Markman (1986, 1987) presented children with triads of pictures as in Figure 7.1. The crow and bat were black and the flamingo was colored pink. By Gelman and Markman's analysis, the crow and bat were similar overall whereas the flamingo and the crow belong to the same conceptual category. The children were told the names of each object (bird, bird, and bat) and then were told that the crow laid eggs. They were asked which other object, bat or flamingo, also laid eggs.

Previous research has shown that young children often freely classify objects like these by perceptual similarity (Tversky, 1985). Gelman and Markman found, however, that children made category inductions along conceptual lines. Gelman and Markman concluded from these results that words access conceptual structure and that categories -- even for very young children -- are organized by a rich theoretical structure that goes beyond appearance. We have no quarrel with this conclusion.

We do, however, question the ancillary assumption that perceptual properties have *no* role in conceptual structure. Gelman and Markman's assumption that perception does not matter is highlighted by their design of experiments and choice of stimuli. Gelman and Markman chose stimuli that distorted the similarity relations that actually exist in the world. The bat and crow are drawn and colored to emphasize the few features that real bats and crows have in common and to de-emphasize the many real differences that exist between

Figure 7.1. A sample triad of stimuli used by Gelman & Markman (1986).

bats and crows and to also de-emphasize the many real similarities that exist between crows and flamingos. Gelman and Markman selected their stimuli for a good reason -- to provide a strong test of the hypothesis that children would use conceptual structure even in the context of strong countering perceptual similarity. Their goal was to choose between two hypotheses -- conceptual structure controls performance or perceptual structure controls performance. Our question is whether this is a theoretically sensible goal.

Consider two more thought experiments that are variants on Gelman and Markman's study. In the first thought experiment, the triad of stimuli looks like that in Figure 7.2. Objects A and B are clearly most similar overall. The "conceptual" information that is provided is that inside A is a bird, inside B is a bat, and inside C is another bird. Now, if the object in A lays eggs, which other object, B or C, is also likely to lay eggs? Given the impoverished character of the perceptual information, one would be wise to generalize according to the labels provided. This experiment if actually conducted, would provide a perfectly good test of whether children can ignore perceptual information. Like Gelman and Markman's original studies, perceptual similarity is pitted against conceptual similarity; and also like Gelman and Markman's stimuli, *the perceptual stimuli provide little information about how the objects really look*. Although the stimuli in Figure 7.2 might provide an adequate test of the perception-versus-conception
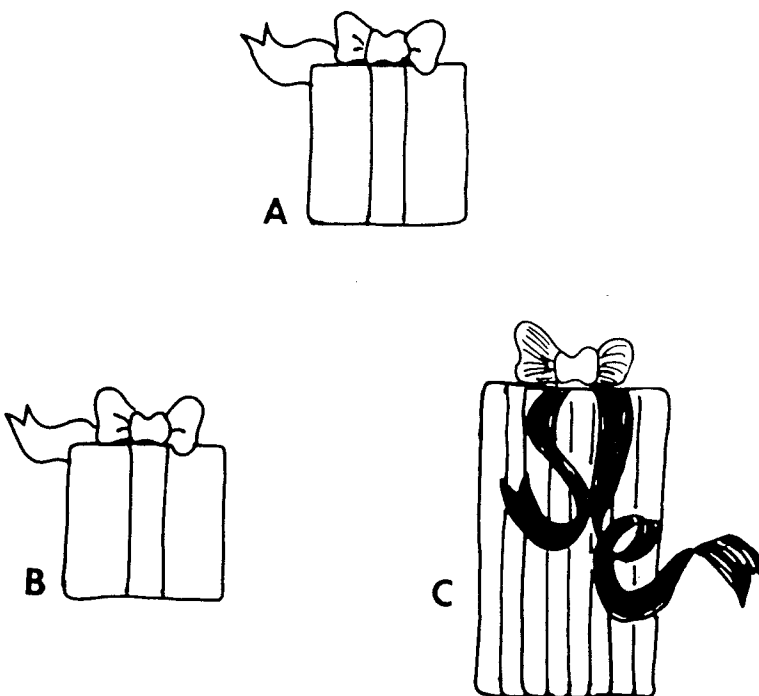
Figure 7.2. Three stimuli for a thought experiment.

hypothesis, they provide a poor test of the uses of perceptual and conceptual information in real category and concept formation.

Now consider a second version of Gelman and Markman's study. In this version, the stimuli consist of a real living flamingo, a living crow, and a living bat. The objects are not named. Nothing is said about them except the fact that one (the crow) lays eggs. In the richness of the real perceptual information, there seems little doubt that egg-laying would be generalized from one bird to the next. Of course, these stimuli are inappropriate to the task of determining whether children use perceptual *or* conceptual information in category induction because the rich perceptual structure and the conceptual structure converge. However, these stimuli might be useful if our research goal is to determine how children form the categories they do. The point is that determining how real categories are formed requires an understanding of perceptual similarity.

III. IN DEFENSE OF PERCEPTUAL SIMILARITY

Defending the role of perceptual similarity in conceptual structure requires a clear statement of just what perception and conception are. We take

perception to refer to the structure of immediate experience. We take conceptual structure to refer to knowledge that is (or can be made to be) explicit -- that is knowledge that is easily talked about and thought about. These definitions appear to the same as those of Keil (1989) and Mandler (1990). However, where we differ is that we believe perceived similarity embodies and reflects much implicit knowledge. Many theorists (e.g. Keil, 1989) write about perceptual similarity as if it were given in the stimulus. This naive realism assumes that the perceived similarity between two objects is constant and unchanging just as the physical measurements of the properties of objects are constant and unchanging. But perceived similarity is the result of psychological processes and it is highly variable. The central process that changes perceptual similarity is attention. The perceived similarity of two objects changes with changes in selective attention to specific perceptual properties.

The importance of attention in categorization has been recognized by other investigators. Murphy and Medin (1985) persuasively pointed out that when people categorize objects and reason about categories, they shift attention in principled ways among sets of perceptual features. For example, people seem to use different perceptual properties when they classify an object as an animal versus when they classify that same object as a pet. Murphy and Medin (see also, Keil, 1989) suggest that people's explicit causal theories or beliefs about objects organize and drive the selection of relevant perceptual features. The idea is that although perception is involved, perception alone -- animal similarity -- is not enough. According to Murphy and Medin, the real force in categorization is the conceptual structure that organizes and interprets perception. That is, by their view, conceptual structure is the cause and changing attention weights on specific dimensions is the effect.

Our view is quite the opposite: the dynamic nature of perceptual similarity is a causal force in the development of conceptual beliefs.

A. Perceptual Similarity is Dynamic

Among those who study perceptual similarity, there is one agreed upon fact, the perceptual similarity between any two objects varies considerably. Perceptual similarity varies with the attributes attended to (Shepard, 1964; Nosofsky, 1984). The dynamic nature of similarity is evident in the empirical research on perception and perceptual categorization and in formal theories of similarity (Goldstone, Medin, & Gentner, in press; Nosofsky, 1984; and Tversky, 1977). In formal theories, similarity is some function of some weighted

combination of features and attributes. As the feature weights change, so similarity changes.

For example, in Smith's (1989) model of perceptual classification as in Nosofsky (1984) and Shepard (1987), similarity is calculated as an exponential decay function of the distance between stimuli in the psychological space. The similarity between two objects, O_i and O_j, then is

$$S_{ij} = e^{-d_{ij}}$$

Distance, d_{ij}, is defined as the sum of the weighted dimensional differences

$$d_{ij} = \sum_{K=1}^{N} W_K [O_{ik} - O_{jk}]$$

where O_{ik} minus O_{jk} is the difference between objects i and j on dimension k, N is the number of dimensions, W_k is the weight given dimension K and $O \leq W_k \leq 1.00$ and

$$\sum_{K=1}^{N} W_k = 1.00 .$$

The present point is that this formula *defines* perceptual similarity; perceptual similarity is *just this* (or some calculation like it). And if perceptual similarity is some weighted combination of dimensional similarities, then perceptual similarity necessarily varies with the magnitude of the difference between stimuli on the dimensions *and with the dimension weights.*

Consider the triad of stimuli on the left hand side of Figure 7.3. These stimuli are represented in terms of their coordinates on two varying dimensions. Object A is identical to Object B on dimension X but differs from B considerably on dimension Y. Object B differs from Object C by a small amount on both dimensions. Using the equations above, the right hand side of Figure 7.3 shows the *perceived similarities* between objects A and B, objects B and C, and objects A and C as the dimension weights change from perfect selective attention to

Figure 7.3. Left -- A triad of stimuli represented in terms of their coordinate values in a 2-dimensional space. Right -- the similarity between pairs in the triad as a function of selective attention to dimension x.

dimension Y ($W_y = 1.00$, $W_x = 0$) through equal weighting of the two dimensions, ($W_x = .50$, $W_y = .50$) to perfect selective attention to dimension X given dimension X ($W_x = 1.00$, $W_y = 0$). As is apparent, the absolute value of the similarity of Objects A and B changes considerably with changes in the feature weights. Moreover, there is a dramatic change in the relative similarities of AB, AC and BC. Which two objects in the triad are the most similar depends on which features are attended to.

   If we extend these notions to Markman and Gelman's stimuli in Figure 7.1, a perceptual interpretation of their results is possible. Under one set of feature weights, say one that emphasized overall shape and color, the crow is perceptually more similar to the bat than the flamingo. Under another set of feature weights, for example, one that emphasized beaks or feet, the crow and the flamingo might be the *perceptually* more similar pair. Importantly, then, a demonstration that children perceive a bat and crow to be similar in a classification task and a crow and flamingo to be similar in the category induction task *need* not mean that children shifted from perceptual similarity to conceptual similarity when asked to make inductions. They may only have shifted the perceptual feature weights. Perceptual similarity may have played the key role in both judgments.

   Nosofsky (1984) represents the effects of shifting feature weights on perceptual similarity in terms of the stretching of the psychological similarity space in one direction or another. Borrowing this idea, we represent hypothetical

Figure 7.4. A theoretical depiction of the similarity relations between Gelman and Markman's stimuli under two different sets of feature weights.

similarity relations between Gelman and Markman's stimuli in Figure 7.4. For ease of discussion, the many-dimensionalspace is compressed into two dimensions. The X axis represents some combination of variation along overall shape and color and the Y axis represents variation on some combination of head and feet features. The three dots represent the location in that space of Gelman and Markman's crow, bat, and flamingo. The distances between the dots represent their perceptual similarity to each other. Figure 7.4 top illustrates the case for an equal weighting of all the dimensions; for these particular drawings, the bat and

crow are the most perceptually similar objects among the three. Figure 7.4 bottom represents the similarity relations when attention is focussed on the head and the feet. The space is stretched so that the crow and flamingo are now closer, that is, more similar than the bat and crow.

One might argue that such stretched similarity spaces are not the original similarity or "animal" similarity meant by Keil and Quine. These theorists meant "raw" similarity unconfounded by knowledge or whatever mechanism pushes feature weights and similarity around. Consistent with these views, Smith (1989) defined raw similarity as the similarity relations that result from the equal weighting of all perceptual features. This overall similarity seems rightly raw in that it is the *default* similarity, the one that dominates perception whenever limits are placed on performance and/or there is no previous experience. Even for adults, an equal weighting of dimensions dominates when processing time is limited, when stimuli are complex, and when there is lack of relevant knowledge (see J. D. Smith & Kemler-Nelson, 1984; Smith, 1981). This is also the similarity that dominates 2- and 3-year-olds' perceptual categorization in most task circumstances (Smith, 1989).

Thus, the equal weighting of all dimensions does seem to be the point of origin for perceptual similarity in the sense of being the zero state of the system. However, overall similarity is neither mechanistically nor developmentally distinct from a stretched similarity space. There is but one system, one manner of calculating perceptual similarity -- the one that is mathematically described by the equations (see Smith, 1989, for an empirical demonstration of this claim). The possibility of changing feature weights is not an add-on to some more primitive form of similarity. Separate features and changing feature weights are inherent · in the very nature and process of perceptual similarity. This claim is supported by the facts of animal perception and by the facts of perceptual similarity in human infancy.

Shifts in feature weights are characteristic of perceptual similarity in non-human animals - i.e., in real "animal similarity" (Mackintosh, 1965; Sutherland, 1964). The phenomenon of cue-blocking (or overshadowing) provides a particularly good example. Consider an illustrative cue-blocking experiment. In original learning, the organism is trained to respond to red objects for a reward but not to respond to blue ones. The shapes of the objects vary and do not matter in the response-reward contingencies. If animals formed conceptual rules, a reasonable one would be "red wins, blue loses, and shape doesn't matter." If we

Figure 7.5. An example of a stretched similarity space in which color is weighted more than shape in the calculation of similarity.

stopped the experimental procedure here and tested the organism with novel shapes and colors, the pattern of response would fit the rule. The organism would respond to all and only red objects and ignore their shape. Figure 7.5 depicts the effects of learning on the perceptual similarity space. Original learning takes the organism from a similarity space in which colors and shapes matter equally (top) to one in which color is (virtually) all that matters (bottom).

The phenomenon of cue-blocking shows that this interpretation is correct. In a cue-blocking experiment, a redundancy phase is added. In this phase, shape is correlated with color. For example, all the red objects are now triangles and all the blue ones are circles. The organism is given some lengthy set of these redundancy trials. The question is whether, given the previous original learning, the animal notices the redundancy. If during original learning, the organism learned to selectively attend to color and ignore shape such that the similarity space has been severely stretched in one direction, then the organism should not notice the redundancy. The results show that in tasks such as these animals do not learn about the added redundancy. After the redundancy phase, the organism responds to all and only red objects and ignores shape just as if the redundancy

training had not occurred. Selective attention is clearly a force in "animal similarity." The similarity space for pigeons and rats is easily stretched.

Another literature that suggests that "raw" similarity may rarely be realized is the infant habituation literature. Consider what happens to perceived similarity in a prototypic habituation study (for real and similar studies, see, e.g., Bornstein, 1985). In a typical experiment, infants might be shown repeated examples of objects that are alike on some dimension but vary on a second dimension. For example, they might be shown red squares, red circles, and red triangles. Looking time is recorded until the babies no longer look much at the red squares, circles, and triangles. Then, the infants are shown new objects -- either one that differs on the previously constant (or relevant) dimension (e.g., a green square) or one that offers a new value on the previously varying (or irrelevant) dimension (e.g., a red cross). The standard result, given repeated examples of the same color during habituation, is that infants look more to the novel color (green square) than the novel form (red cross) during dishabituation. Importantly, if infants were shown repeated examples of the same shape but varying colors during habituation, they would look more at the novel shape than the novel color.

Presumably, the mechanism behind such habituation-dishabituation effects is short-term changes in features weights. If the habituation stimuli are all red but with varying shapes, the common redness of the habituation stimuli must somehow stretch the similarity space along the color dimension (or the varying shapes cause a shrinkage of the space along the shape dimension). The habituation phase causes a change in the perceived similarity of objects; in our example, the habituation phase made differences in color count more in the calculation of perceptual similarity than differences in form. Original similarity, the perceptual similarity of infancy, is inherently dynamic; it moves about because of the mechanism through which it is computed. The mechanism is one in which feature weights may vary from context to context. The use of habituation as a method to study infant perception depends on the reality of malleable feature weights in psychological similarity.

If perceptual similarity shifts so readily in the artificial laboratory contexts of operant conditioning and habituation experiments, might not perceptual similarity shift systematically and *meaningfully* in response to the contingencies and correlations that exist between perceptual features in the real world?

Figure 7.6. A distorted region of a similarity space as might result from increased attention to a particular conjunction of features (the relevant conjunction is the compressed area).

B. Knowledge and Perceptual Similarity

The relevance of perceptual similarity for conceptual structure stems from the role of real world experience on feature weights. One kind of knowledge that pushes feature weights around and stretches the similarity space is implicit knowledge about relations between perceptual features. As Rosch (1973) argued, perceptual features do not vary orthogonally in the world. They come in causally related clusters. Birds with web feet tend to have bills. Objects with dog-like feet tend to have dog-like heads. Evidence from laboratory experiments indicates that both adults (Medin et al., 1982) and older infants (Younger & Cohen, 1983; Younger, 1990) are sensitive to such correlations. This empirical evidence indicates that experience with correlations causes increased attention to the *combinations* of features that enter into correlations.

Figure 7.6 represents how such a correlation might organize a similarity space. The grid lines represent "equal" physical distances. Perceptual similarity is represented by real Euclidean distance in the space. In the depicted space, objects that possess a particular conjunction of values are closer together or perceptually more similar than objects that possess other particular conjunctions of values. Thus the compacted upper right of the figure might be the region of the stimulus space in which objects with dog-like feet *and* dog-like heads fall. If this figure were an accurate depiction of perceived similarity, it would mean that the similarity between objects with both dog-feet and dog-heads is greater than the

Figure 7.7. Stimuli and mean propostion name extensions to test stimuli in Jones, Smith & Landau (1990).

similarity between objects with both dog-feet *and* pig-heads. Do young children, in their everyday interactions in the world, learn about such correlations in ways that systematically and to good purpose distort the similarity space?

1. The Case of Eyes and Texture

Recently, Jones, Smith, and Landau (1990) discovered a co-relation effect in children's novel word extensions that suggests a powerful role for combinations of perceptual features in natural category formation. Jones et al. found that the perceived similarity of objects for 36-month-olds changed dramatically with the addition of a constant perceptual property. In the experiment, children were shown a novel object and it was named. For example, they might be shown the object at the top left of Figure 7.7 and told that it was a Dax. Children were then asked what other test objects were also Daxes. As shown in the left column of the figure, one test object differed from the exemplar only in shape, one only in texture, and one only in size. While the exemplar (the experimenter-named) Dax remained in view, children were asked separately whether each one of the test

objects was also a Dax. The proportion of times the children agreed that the exemplar's name was also the name of the test object is given next to each test object. As the results attest, these children called the test objects by the same name as the exemplar when *the object was identical in shape to the exemplar object*. Shape changes from the exemplar mattered considerably more than changes in texture and size.

The right side of Figure 7.7 shows the stimuli for a second (between-subjects) condition. In this condition, toy eyes were affixed to each object. The addition of a constant stimulus property (the eyes) radically changed children's judgments as is evident in the proportions of name extensions. When eyes were added, the exemplar's name was extended to a test object only if the test object was the *same shape and same texture* as the exemplar.

These results demonstrate the synergistic relation that must exist between perceptual properties and knowledge. The presence of eyes activates children's knowledge about kinds of categories. That knowledge then guides attention, changing the weights accorded to the various dimensions, and thus shifting the perceptual similarities among the objects.

The kind of co-relation that exists between eyes and texture is different from the correlations between properties that have been studied by Medin et al. and Younger and Cohen. The kind of correlation studied by previous investigators is between specific properties, for example, between feathers, webbed feet, and bills. The kind of co-relation that underlies 3-year-olds novel word extensions does not seem to be between eyes and a *specific* texture.

In the world, if having eyes is correlated with a specific texture, it is with one that is soft and pliable. In the Jones et al. study, the Dax exemplar was made of wood. In one experiment, the texture-change test item was a sponge, in another, it was a soft cloth. If the presence of eyes had signalled the importance of particular textures, then the children should have been more likely to accept these softer Daxes in the Eyes condition than in the No-eyes condition even though they had a different texture than the exemplar. However, the results are just the opposite of this specific-feature-correlation hypothesis. Children in the No-eyes condition readily called these soft objects Daxes. However, children in Eyes condition rejected these items. They required objects with eyes to have the same wooden texture as the exemplar to be called a Dax. Apparently, the presence of eyes indicated that texture -- whatever that texture might be -- is criterial in categorization.

This kind of co-relation -- between a specific property (eyes) and a *dimension* of variation (texture) -- should prove a powerful aid in learning new category structures, one that would supplement the use of correlations between specific properties (e.g., wings and feathers). A learned correlation between the specific properties of wings and feathers is a local distortion of the similarity space as shown in Figure 7.6. Such a local distortion would help the categorizer distinguish objects that do or do not belong to an already-learned category.

A learned correlation between a cue and dimension of variation (like that between eyes and texture), however, causes a more widespread distortion of similarity. The entire similarity space would be stretched in one direction exaggerating texture differences in the context of eyes relative to the context of no eyes. Such an expansion of the space should help the child organize multiple categories of birds, and bats and walruses, and of real monkeys, mechanical monkeys, and snakes because the objects in these categories differ from one another in the texture of their surfaces. The idea, then, is that the similarity space reflects the fact that surface texture differences matter more for categorizing all objects with eyes than such differences matter for categorizing objects without eyes. A general expansion of the similarity space along the texture dimension for objects with eyes would increase the similarity of real crows and flamingos (which both have feathered surfaces) and decrease the similarity of crows and bats (since crows are feathered and bats are furry) relative to eyeless counterparts.

If we imagine multiples of local and dimension-wide distortions of the similarity space -- distortions resulting from real-world correlations between specific properties, from co-relations between material kind and kind of motion, eyes and texture, eyes and kind of motion, shape and motion, and so on -- then what emerges is a bumpy and irregular similarity space that *organizes itself into multiples of categories at multiple levels* in context dependent ways.

We suggest that the result of spending time in the world -- of looking at it, hearing it, and feeling it -- is a structured and context dependent similarity space. The result is a dynamic similarity space that fits the structure of the world. Rosch (1973) suggested that many natural categories are given in the world in correlated features sets. Her proposal seems right. The reason that the thousands of categories named by concrete nouns emerge so rapidly in the first three years of life and are conceptually well understood by four or five years of age is that they are given in a context dependent perceptual similarity.

This idea that perceptual similarity *reflects* the co-relations between perceptual properties as they exist in the world presumes a particular mechanism.

We presume what most formal theories of perceptual similarity and perceptual categorization posit -- that learned *atheoretical* associations shift the features weights and the computed similarities in systematic ways that reflect the structure of the world (see Nosofsky, 1986; Gluck & Bower, 1988; Medin & Schaffer, 1978).

There is another explanation of the effect of eyes on 3-year-olds' categorizations -- one built on represented causal theories. For example, 3-year-olds' judgments could be organized by knowledge structures of the sort: objects with eyes are natural kinds, members of the same natural-kind class have the same genetic structure, the same genetic structure produces the same material substance; the same material substance produces the same surface texture; therefore, texture matters for objects with eyes. Such a complicated set of beliefs seems implausible to us but more to the point, we have evidence that such beliefs are not necessary for the co-relation effects reported by Jones et al.

We (Heise & Smith, 1990) have demonstrated that the same phenomenon observed in children's novel word extensions occurs in adults' learning of arbitrary categories. Billman (1989) also has reported compelling evidence that makes the point. In our study, adults learned two different classification systems in a single experiment. For example, subjects might be given red bugs and blue bugs and learn to classify the objects into four groups such that the red bugs were partitioned into two subgroups by shape (circular vs. dropshape) and the blue bugs were partitioned into two subgroups by number of legs (2 versus 4). Given adults proclivity for selective attention and the formation of criterial property categories (Medin, Wattenmaker and Hampson, 1987), one might expect adults in this task to learn only to attend to shape and/or number of legs. Transfer tasks, however, showed that adults learned a more general co-relation between color and the other dimensions. In the transfer task, adults were given new sets of bugs and asked to freely classify them into groups. Unlike typical transfer tasks, then, we did not ask how well adults put new items into the just learned categories, we asked how well they formed new categories of new items -- on their own -- and without instruction.

Adults' transfer performance in this free classification task suggests they had learned an implicit "rule" of the sort: red bugs are distinguished into subkinds by their shape but blue bugs are distinguished into subkinds by their number of legs. Given the training described above, adults classified red bugs by their shape even when the novel bugs had shapes never seen before and they classified the blue bugs by their number of legs even when the numbers of legs were different from those in the learning phase. This acquisition of a context dependent shift in

feature weights was obtained with various transpositions of the cuing and transfer dimensions.

These results suggest that context-dependent distortions of the similarity space are easily set up. Moreover, the adults in these experiments had no pre-experimental naive causal theories about how one dimension should be related to another in classifying cartoon bugs. The subjects also do not appear to have developed such causal theories in the course of the experiment. Clearly, causal theories or well-developed belief systems are not necessary for context dependent and useful shifts in attention to perceptual properties.

2. The Case of Lexical Form-Class

The case of eyes and texture indicates that one perceptual property may influence attention to another perceptual property and systematically alter the perceived similarities of objects. But knowledge of perceptual properties is only one force on perceptual similarity. Other *non*perceptual forces (though probably not explicit knowledge) also play a role. For example, young children, shift their attention among dimensions in novel word extension tasks as a function of the syntactic form class of the word.

In the Jones, Smith, and Landau (1990) study, (see also, Landau, Smith, & Jones, 1988), young children's selective attention to shape with eyeless stimuli and their attention to shape and texture with eyed stimuli occurred only in a word extension task. When children were asked to freely classify the same stimuli, they showed no differential attention to the dimensions of size, shape, and texture. The presence of a novel word, therefore, organized the feature weights and shifted them from the default value of equal attention to all dimensions. In those studies, the novel word was a count noun.

Recently, Smith, Jones, and Landau (1990) examined 36-month-olds' shifting attention to shape and color when the novel word was an adjective as well as when it was a count noun. In one study, half the children were presented with an exemplar object and told it was "a dax" and half the children were presented with the same exemplar object but were told it was "a dax one." The exemplar was made of wood, possessed a zig-zag shape and was colored a glittery (and highly reflective) combination of silvery gold. The critical test items consisted of (1) two unique objects that were the same shape but different color and (2) two unique objects that were the same color but different shape. Children were shown each test object one at a time and asked whether it was a "dax" in the noun

Table 7.1. Mean proportion extensions to critical test trials in the Adjective and Noun conditions in Smith, Jones & Landeau (1990). (Standard deviations are in parentheses.)

	Noun		Adjective	
	Same Color Diff. Shape	**Same Shape Diff. Color**	**Same Color Diff. Shape**	**Same Shape Diff. Color**
No Cave	.20 (.23)	.85 (.11)	.51 (.31)	.78 (.29)
Cave	.29 (.22)	.94 (.08)	.73 (.18)	.27 (.20)

condition or a "dax one" in the adjective condition. The exemplar was in view through out the entire procedure.

The results are given in Table 7.1 in the row labelled "No cave". Children in the noun condition attended to shape calling same shape items a *Dax*; children did *not* call same color but different shape items a *Dax*. Children also showed a weak shape bias in the adjective condition. However, in the adjective condition, individual children's performances differed from each other. Most children maintained that the same-shape objects were "dax ones." But some children believed same-color objects were "dax ones" and some children responded haphazardly. The shiny color of the exemplar was highly noticeable and indeed was spontaneously commented on by *all* the children. Nonetheless, children did not systematically interpret the novel adjective as referring to this novel and salient feature. These results show that the presence of a novel word, particularly a count noun, organizes attention to shape.

The results of a second experiment showed that children's attention in the adjective condition could be as highly organized as their attention in the noun condition -- if additional forces directed their attention among the varying dimensions. The procedure in this second experiment was identical to that in the first with the exception that the stimuli were presented inside a small dark cave. The child and experimenter looked through the opening of the cave to view the exemplar and each test item. A spotlight in the cave illuminated the stimuli. The cave and spotlight had the effect of heightening the already salient glittery color.

The results from this second experiment are also given in Table 7.1 in the row labelled "Cave". Again, in the noun condition, children extended the novel word to new items according to shape alone. But in the adjective condition, children attended to color and called objects "dax ones" only if they possessed the same glittery color as the exemplar.

These results show that both lexical knowledge and local forces conspire to control children's attention to dimensions and the perceived similarity of objects. Presumably, in the course of learning language, children have learned that shape is the principle determiner of membership in the categories labelled by count nouns (see Landau et al., 1988). The act of naming objects thus may serve as a cue to stretch the similarity space along the shape dimension. Apparently, however, adjective categories are not well constrained by kind of property and thus the context of an adjective does not systematically distort the similarity space. But local, context-specific, *ad hoc* forces can work synergistically with the adjective context to forcefully organize attention and thereby constrain possible interpretations of the novel adjective.

We do not find these results surprising. They fall right out of the mechanism that underlies perceptual similarity. However, since it seems that many kinds of knowledge shift feature weights -- associative connections between perceptual properties, knowledge about syntactic form classes, local *ad hoc* effects, some might question whether distortions of similarity are truly *perceptual* phenomenon. Is the shifty nature of perceptual similarity too contaminated by (implicit) knowledge to be rightly considered perceptual? We believe there is little to be gained from pursuing this question. The critical point is that perceived similarity is dynamic and shifts in meaningful ways.

## C. Conceptually Relevant Perceptual Properties

The power of perceptual similarity and shifting attention among features becomes more impressive when we consider abstract perceptual properties. There are lots of perceptual properties that are not easy to talk about but that may shape our conceptions of the world from very early in development. We are thinking of the perceptual properties that distinguish biological from nonbiological motion (e.g., Berthental, Proffitt, Spertner & Thomas, 1985; R. Gelman, Spelke, & Meck, 1983; Bullock, R. Gelman, & Baillageon, 1982), that allow us to predict what novel objects can and cannot move alone (Massey & R. Gelman, 1988), and to determine whether an action by one object on another was intentional (Leslie,

1984; 1988; Spelke, 1982). There is a growing research effort to discover such perceptual properties that must underlie, indeed organize, conceptual categories.

Briefly, we would like to consider one perceptual property that may underlie the contrast between naturally occurring objects and manufactured objects. Our hypothesis is that manufactured and naturally occurring objects differ *in a perceptible way* in their surface gradients. Consider the two stars shown at the top of Figure 7.8. They have the same global shape. But if one of these is a manufactured holiday star and the other a living starfish, it is the one on the right that was made in a factory and the one on the left that was found at the beach. Naturally occurring objects tend to have an intricate surface gradient that is distinct from that of manufactured objects.

Mandelbrot (1983) described the surfaces of natural objects as fractal curves. These surfaces are self-similar across changes in scale. Thus, if we
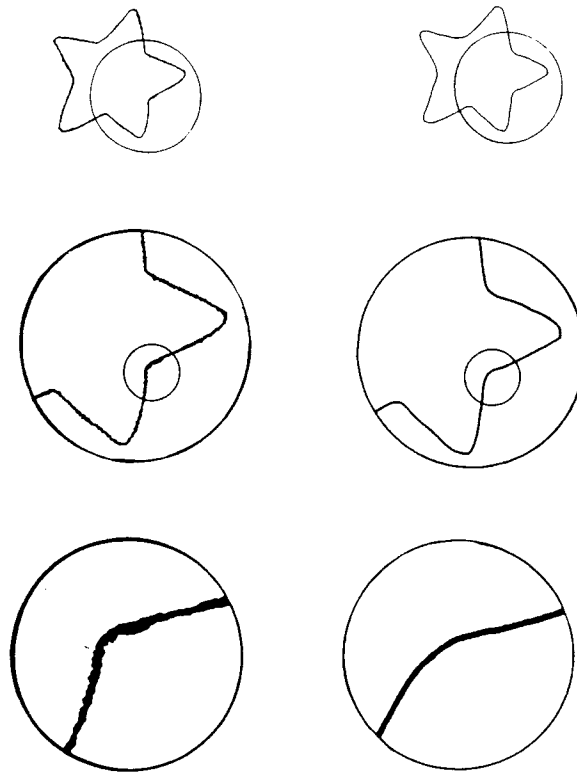
Figure 7.8. Two objects with the same global shape but different surface gradients and enlargements of portions of those objects.

magnified the starfish, we would see bumps on the surface. If we magnified those bumps, we would see bumps on the bumps and so on. We may speculate, following Mandelbrot, that the complex and seemingly irregular surface gradient of natural objects stems from the fact that they grow and are caused by a multiple of converging forces whose effects accrue in time. In contrast, the scale-specific structure of manufactured objects presumably derives from the fact that they were made at a particular scale level and at a particular point in time.

Are children sensitive to the surface gradient differences between naturally occurring and manufactured objects? Diana Heise, Susan Rivera and I have preliminary evidence that 12-month-old children are. In this ongoing research, the task is preferential looking. In the baseline experiment, the stimuli were toy vehicles and life-like model animals. Model animals, of course, are manufactured and so it might be argued that they are poor stimuli for a study of the perceptual properties afforded by naturally occurring objects. However, the perceptual properties that make one model animal seem life-like and another toy-like or robot-like are presumably the same ones along which naturally occurring and (typical) manufactured objects differ.

Figure 7.9. Mean looking time on the first and last of 14 familiarization trials showing Animals or Vehicles and on the test trial in the Between-category (AN-VEH/VEH-AN) and Within-category (AN-AN/VEH-VEH) conditions when the stimuli were intact toy animals and vehicles.

In the experiment, the babies were repeatedly shown in alternation two toys and then after familiarization were shown a novel test toy. In two Between-Category conditions, the novel test toy was from a different superordinate class than the familiarization toys. In the Within-Category conditions, the novel test toy was from the same class as the familiarization stimuli. For example, children in the Animal/Between-Category condition might see a goat and cow during familiarization and a boat during test. Children in the corresponding Animal/Within-Category condition would see a goat and cow during familiarization and an elephant during test. The results in Figure 7.9 show clear evidence that infants discriminate animals from vehicles. Looking time during test goes up more in the Between-Category conditions than in the Within-Category conditions.

There are a number of perceptual properties along which babies could (and probably do) distinguish toy animals and toy vehicles that include eyes, mouths, global shape, and by our hypothesis, surface gradient. We examined the infants' ability to use surface gradient cues in a subsequent experiment by removing all potential cues except surface gradient and asking if babies could still distinguish animals from vehicles. Specifically, we cut up the toys into approximately 2" x 2" pieces, removing any pieces with eyes, mouths, or parts of


Figure 7.10. Mean looking time on the first and last of 8 familiarization trials showing Animals or Vehicles and on the test trial in the Between-category (AN-VEH/VEH-AN) and Within-category condition (AN-AN/VEH-VEH) conditions when the stimuli were cut-up toy animals and vehicles.

faces (and all wheels from the vehicles). We separately mounted the remaining pieces of each toy in a circle for viewing by the infant and repeated the baseline experiment with these cut-up stimuli. The results shown in Figure 7.10 indicate that children discriminated cut-up animals from cut-up vehicles. Contrary to the first experiment, children looked *more* at the Within-category test stimulus than the Between-category stimulus. Although a preference for the familiar as opposed to novel was shown with the "parts," the results, nonetheless, show babies discriminate animal textures from vehicle textures. Thus, there do appear to be perceptual properties that distinguish animals from vehicles. Along at least one perceptible property, surface gradient, animals and vehicles do not look alike. Perhaps, then, the babies in Mandler & Bauer's study were simply classifying the objects by what they looked like.

Our finding that babies are sensitive to a texture cue that may distinguish animals from vehicles has implications beyond Mandler & Bauer's results. The findings provide a useful starting point for thinking about the relation between perceptual and conceptual structure. The results show a perceptual dimension along which many naturally occurring and manufactured objects may be distinguished and to which infants are sensitive. Moreover, it is the nature of the human apparatus for perceptual similarity to link up this dimension with co-occurring properties and in this way to selectively enhance attention to this dimension in the stimulus contexts in which it most matters -- for example, in the context of stimuli with eyes, or stimuli that move in a certain biological way.

Our claim here is not that surface gradient is a perfect or fool-proof cue by which to distinguish naturally occurring from manufactured objects. There are stones so smooth from the continual action of the waves that they are as perfect spheres as ball bearings and there are talented artists who can fool the eye with their creations. Our idea, though, is that surface gradient is an important force in starting, maintaining, and grounding conceptual structure to reality.

This concludes our defense of perceptual similarity as a major player in category development. Perceptual similarity is context dependent. It is shifted about by implicit knowledge of relations between perceptual properties. And, the perceptual properties that can be emphasized or de-emphasized in similarity judgments may, like the presence of eyes and like the surface gradient differences between animals and vehicles, be highly conceptually relevant -- of a kind worth developing a causal theory to explain.

IV. PERCEPTUAL SIMILARITY AND CAUSAL THEORIES

Perceptual and conceptual similarity are not, by our view, the same thing. Conceptual similarity does not reduce to perceptual similarity. But perceptual and conceptual similarity are also not independent. We will consider first how perceptual and conceptual similarity are different. Then, we will consider the causal dependencies between the two.

A. Perceptual and Conceptual Similarity are Not the Same

When we look at objects, what we *perceive* reflects much implicit knowledge --- a myriad of associations between perceptual features, linguistic, social, and physical context. Perceptual similarity is thus dynamic and smart but it is not smart in the same way that conceptual structure is. Perceptual similarity is not a set of causal beliefs and explanations; it embodies causal relations but it does it not represent them. Thus, when children look at objects with eyes they perceive the similarities and differences in texture and their perception is smart and adaptive. But there is in the perception no causal understanding of why textures matter for objects with eyes.

Rozin (1976) in his landmark paper on intelligence distinguished between welded implicit knowledge of the sort that allows birds to migrate and transportable explicit knowledge such as that used by navigators on ships. Rozin characterized implicit knowledge as rigid and inflexible. Our view of the implicit knowledge that is embodied in the dynamics of perceptual similarity does not fit this description. Rozin also characterized implicit knowledge as being welded to particular contextual factors that are necessary for the knowledge to show itself. The knowledge behind perceptual categorization would seem to be welded in the sense of being tightly determined by the actual perceptual properties present, past specific learning, and context. Explicit knowledge, true intelligence according to Rozin, is knowledge that is more context free -- that may be *voluntarily* transported across contexts to new problem spaces. Conceptual structure may be smart in this way.

This distinction between implicit perceptual structure and explicit, transportable, conceptual knowledge is attested to by the difficulty one has in *consciously* controlling perceived similarities. Whereas conceptual similarity may be strongly influenced by being told some new fact, perceptual similarity may be moved along only a little by "mere talk." Being told that a purple, three-legged beast with moose horns is a cow, doesn't make the unfortunate animal *look* like

a cow (although one might, if given that information by experts, claim to believe the odd beast to be a cow despite its looks). The separateness of perception and our easily talked about conceptual beliefs is evident in linguistic hedges (Lakoff, 1972) and statements such as, "It looks like a fish but it's really a mammal." This obvious distinction between perceptual knowledge, which is not transparent to conscious thought processes, and explicit conceptual knowledge probably underlies the temptation to empirically pit perception against conception in developmental experiments.

B. Perceptual and Conceptual Structure are Causally Related

The separateness and different status of perceptual and conceptual structure does not mean they are not causally related. There are causal relations that go in both ways. Explicit conceptual knowledge can push perception only so far (no amount of information will convince a person looking at a spoon that it is a tomato). However, conceptual knowledge can cause individuals to seek out perceptual dimensions that make sense of conceptual distinctions. This effect of conceptual knowledge on perception is aptly demonstrated in Heise's ongoing dissertation research. In that research, children are presented with real world correlates of the picture stimuli used by Gelman and Markman (1986). One triad consists of a ball of yellow wool yarn, a ball of yellow acrylic yarn, and raw unwoven wool. The two balls of yarn look very much alike. However when told that the wool yarn and the raw wool are both wool, 5 year old children (and adults) often look more closely and say such things as "Yes, these two are less shiny, that's how you can tell"). This example is telling on two counts. First, it shows how a conceptual distinction may invite the search for a perceptual distinction. Second, it shows how the human cognitive system attempts to keep conceptual and perceptual structures in line with each other. Conceptual growth does not cause us to abandon perception but it may cause us to take another look.

Although explicit conceptual knowledge may sometimes be a causal force in perceptual comparison, the causal relation between perception and conception goes mostly one way --- with perception the cause and conception the effect. Perceptual structure *ties* our conceptual beliefs to reality; it is the mapping function that takes our representations and beliefs and gives them meaning. Perceptual structure is the *data* that our conceptual theories are about and just as the data are the final arbiter of truth in science so is perception the unassailable constrainer of concepts in cognitive development.

Perceptual structure is what conceptual beliefs are about. The atheoretical, uninterpreted associations and correlations between perceptual features makes a bumpy terrain of the perceptual landscape. Our explicit conceptual beliefs try to make conscious sense of this terrain. *Because* our conceptual beliefs are about making sense of the world, perceptual and conceptual structure will generally agree.

However, our explicit conceptual beliefs, even at maturity, probably do not do a very good job of making explicit the implicit interrelations and co-relations that structure and organize perception. Perceptual knowledge is deep, highly entrenched, highly embedded and probably much smarter than conceptual knowledge (see Gentner & Rattermann, 1990, for a similar argument) this is so despite the fact that the knowledge embodied in perception is severely limited in its transportability across cognitive domains. The evidence on how people readily catch thrown balls without understanding the physics is but one example of the smart but welded nature of perceptual knowledge (see, e.g., Bingham, Schmidt, E. Rosenblum, 1989). We expect the same contrast between "smart perception" and "less smart" conception is true for categories. It is perceptual knowledge that is deep and it is the conceptual structure that is superficial and inaccurate as it depends on cultural and perhaps mistaken scientific beliefs (see Jeyifous, 1986; Lakoff, 1987; Putnam, 1975 for examples).

C. How Conceptual Structure Depends on Perceptual Structure

In his recent book, Keil (1989) examines Boyd's proposal about causal homeostasis as a model for human conceptual structure. The idea is that categories are not just clusters of correlated properties. Instead, there are sets of causally important properties that are contingently clustered. The correlated properties are not the results of happenstance. The laws of physics, biology, and behavior (the laws scientists seek to discover) are such that properties co-occur (and vary together) *because* of causal "homeostasis." The presence of some properties favors or causes the presence of others.

How is this structure in the world related to the psychological structure of categories? It is informative here to contrast our position with Keil's proposal. Following Quine, Keil considers an initial "original similarity" state that is represented as in Figure 7.11A. The perceptual primitives are represented by the circles, triangles, squares, and ellipses. These primitives are connected by pretheoretical associations that are formed when specific features co-occur in the world. What happens with development, according to Keil, is that causal beliefs

are overlaid on this perceptual structure. These beliefs are shown by the heavier arrows in Figure 7.11B. (Keil suggests that 7.11B may in fact more closely represent the starting point for development than 7.11A since some domain-specific causal theories may be "innate".) With increasing age, theories are elaborated as in 7.11C. Early in development, Keil suggests, causal beliefs may interpret the perceptual associations between primitives but with development, they become more systematic, elaborate, and differentiated.

Figure 7.11. Taken from Keil (1989). A: Networks of features linked by associations; B,C: The emergence of theory indicated by arrows.

We concur with the main points of this outline of a theory. Yet we find its portrayal of perceptual structure disquieting. Again, static perceptual structure provides at best a minor constraint on conceptual structure. In Keil's description, perceptual structure only weakly reflects (in its associative connections) the causal homeostasis that exists between properties in the world. The causal homeostasis

CONTEXT   **THE PERCEPTUAL LANDSCAPE**

Figure 7.12. Perceptual landscapes that vary with context and causal belief systems (arrows) that attempt to explain them.

that Boyd writes about finds its psychological home in Keil's view in people's elaborated and systematic causal beliefs.

We offer as an alternative to Keil the outline of a theory of perceptual context C_1, is shown the perceived similarity of objects along some set of and conceptual structure depicted in Figure 7.12. What this figure attempts to depict is the changing perceptual landscape with changing context and people's theories about the perceived similarities. The dynamic perceptual landscape is represented by the context dependent and variously distorted grids. Conceptual theories are represented by the arrows that point out relations between distortions in the perceptual landscape and between distortions and context. So, next to dimensions in that context. The perceived similarities in the space are not uniform because of learned correlations, context-dimension associations, and language-dimension associations. We assume all these associations are acquired from being in the world. Context C_2 shows the distortions of *this same similarity* space in another context that activates another set of associations that shift the dimension weights and perceived similarity in other directions. Thus, in contrast to Keil, we do not present one perceptual landscape because there is no one landscape. There is no static original similarity. Instead the perceptual landscape emerges in particular contexts with particular coalitions of features. Morever, because of the nature of perceptual similarity and the context-specific shifts in feature weights, these dynamic perceptual landscapes reflect the causal structure of the physical and social world. Boyd's causal homeostasis finds its psychological home in this theory in the bumps and holes and ridges of perceived similarity.

People's explicit causal beliefs about categories and relations between objects are graphically illustrated by the arrows. By our view, people build causal connections between points in the perceptual landscape --- incorporating cultural and scientific knowledge --- in order to interpret and make sense of the complex perceptual terrain. Some causal beliefs may be local and about interdependencies in a single contextually determined terrain. Other naive theories may be grander and try to capture the dynamics of the changing emphasis and de-emphasis of features across contexts. It is unlikely that these naive theories adequately reflect or explain the rich structure of the perceptual landscape or the causal homeostasis in the world.

The difference between our representation and Keil's, then, is that the perceptual structure is richer and far more complicated. Perception provides the nutrients on which conceptual structure grows. Perception is the reality on which the conceptual system operates. Indeed, it is because we are perceptually in

contact with the complex world that we need an elaborated conceptual structure that simplifies that complexity.

V. STRUCTURE AND PROCESS

The chapters of this book have all grappled with the question of the relation between structure and process. Although, we have not directly confronted the issue, the distinction between structure and process may be the root of the problem in understanding perception-conception interactions. The idea of the structure-process distinction is that cognition is composed of two parts: structure which is the represented knowledge and process which operates on the structures. In this view of cognition, structures are the nouns of thoughts, *the things*, and processes are the verbs, the actions. The problem with this metaphor is that perceptual and conceptual structures are designated as things as if that they were static and unchanging. This metaphor is particularly troublesome when thinking about development because development is about change (see, Smith, 1990; Smith & Sera, 1990).

How might cognitive development go forward if knowledge is a set of "things" --- if we acquire pieces of knowledge like new pieces of clothing? One possibility is that children just keep adding new pieces to their knowledge set. Another possibility is that children abandon or ("trade up") their use of one structure in favor of another. For example, they might abandon perceptual structure in favor of conceptual structure. We have argued that this particular solution to developmental changes in categorization does not fit the data.

There is an alternate to partitioning cognition into structure and process. The alternate view is that there is only process. By this view, perceptual and cognitive structures are not thing-like representations that reside in the child but are rather the emergent properties of complex processes. This alternative view is, of course, the one that underlies the theoretical endeavors known as connectionism (e.g., Rumelhart & McClelland, 1986) and dynamical systems theory (Thelen, 1989; Smith & Sera, 1990). It is the assumption that guides our work and forms our belief in perceptual similarity as dynamic. Perceptual similarity --- perceptual structure --- is dynamic because perceptual structure is not a thing; it is the emergent result of attentional processes.

VI. CONCLUSIONS

We have argued for a new research agenda -- one that does not pit perception against conception but instead asks how perceptual and conceptual structure are related and how they organize together to propel development forward. We have specifically argued for greater attention to the dynamics of perceptual similarity and perceptual structure. Perceptual similarity is a much maligned force in cognition. Philosophers such as Nelson Goodman (1972) have dismissed similarity as too slippery and variable a concept to have any explanatory power. By our view, the power of perceptual similarity stems directly from its (systematically) variable nature.

ACKNOWLEDGEMENTS

REFERENCES

Bertenthal, B. I., Proffitt, D. R., Kramer, S. J., & Spetner, N. B. (1987). Infants' encoding of kinetic displays varying in relative coherence. Developmental Psychology, 23, 171-178.

Biederman, I. (1985). Human image understanding. Computer Vision, Graphics, & Image Processing, 32, 29-73.

Billman, D. (1989). Systems of correlations in rule and category learning: Use of structured input in learning syntactic categories. Language and Cognitive Processing, 4, 127-155.

Bingham, G. P., Schmidt, R. C. and Rosenblum, L. D. Hefting for a maximum distance throw: A smart perceptual mechanism. Journal of Experimental Psychology: Human Perception and Performance, 15, 507-528.

Bornstein, M. H. (1985). Habituation of attention as a measure of visual information processing in human infants: Summary, Systematization, and Syntheses. In G. Gottlieb and N. Krasnegor (Eds.) Measurement of audition and vision in the first year of postnatal life: A methodological overview. Norwood, N. J.; Ablex, 253-300.

Brown, A. L. (1990). Domain-specific principles affect learning and transfer in children. Cognitive Science, 14, 107-134.

Bruner, J. S. & Olver, R. R. (1963). Development of equivalence transformations in children. Monograph of the Society for Research in Child Development, No. 28.

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), The developmental psychology of time. New York: Academic Press.

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

Flavell, J. H. (1970). Concept development. In P. H. Mussen (Ed.), Carmichael's manual of child psychology: Vol 1. New York: Wiley.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and the 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 and J. A. Sloboda (Eds.), The acquisition of symbolic skills. London: Plenum.

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

Gelman, S. & Markman, E. M. (1986). Categories and induction in young children. Cognition, 23, 183-209.

Gelman, S. A. & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. Child Development, 58, 1532-1541.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony (Eds.), Similarity and analogical reasoning (pp 199-241). New York: Cambridge University Press.

Gentner, D. & Rattermann, M. J. (1990). Language and the career of similarity. In S. Gelman (Ed.), Language and conceptual development.

Goodman, N. (1972). Problems and projects. Indianapolis: Bobbs-Merrill.

Goldstone, R. L., Medin, D. L., & Gentner, D. (in press). Relational similarity and the non-independence of features in similarity judgments. Cognitive Psychology.

Greer, A. E. & Sera, M. D. (1990). Artifacts vs. natural kinds: An empirical investigation. (Unpublished manuscript).

Gluck, M. A. & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. Journal of Experimental Psychology: General, 117, 227-247.

Heise, D. & Smith, L. B. (1990). Learning co-relations. (Unpublished manuscript).

Jeyifous, S. (1986). Antimodemo: Semantic conceptual development among the Yoruba. Doctoral dissertation, Cornell University.

Jones, S. S., Smith, L. B., & Landau, B. (1990). Object properties and knowledge in early lexical learning. (Under review)

Johnson, M. (1987). The body in the mind: The bodily basis of meaning, imagination, and reasoning. Chicago: University of Chicago Press.

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

Kemler Nelson, D. G. (1990). When experimental findings conflict with everyday observations: Reflections on children's category learning. Child Development, 61, 606-610.

Lakoff, G. (1972). Hedges: A study in meaning criteria and the logic of fuzzy concepts. In Papers from the Eighth Regional Meeting of the Chicago Linguistic Society. Also in Journal of Philosophical Logic (1973) 2, 458-508.

Lakoff, G. (1987). Women, fire, and dangerous things: What categories reveal about the mind. Chicago: University of Chicago Press.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The Importance of Shape in Early Lexical Learning Cognitive Development, 3, 299-321.

Leslie, A. (1984). Infant perception of a manual pick-up event. British Journal of Developmental Psychology, 2, 19-32.

Leslie, A. (1988). The necessity of illusion: Perception and thought in infancy. In L. Weiskrantz (Ed.), Thought without language. Oxford: Clarendon.

Mackintosh, N. J. (1965). Selective attention in animal discrimination learning. Psychological Bulletin, 64, 124-150.

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

Mandelbrot, B. B. (1983). The fractal geometry of nature. San Francisco: Freeman.

Mandler, J. M. (1988). How to build a baby: On the development of an accessible representational system. Cognitive Development, 3, 113-136.

Mandler, J. M. (1990). How to build a baby: Part 2. (Unpublished manuscript)

Mandler, J. M. & Bauer, P. J. (1988). The cradle of categorization: Is the basic level basic? Cognitive Development, 3, 237-264.

Mandler, J. M., Bauer, P. J., & McDonough, L. (in press). Separating the sheep from the goats: Differentiating global categories. Cognitive Psychology.

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

Medin, D. & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou and A. Ortony (Eds.), Similarity and analogical reasoning. New York: Cambridge University Press.

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

Medin, D. L., Ahn, W., Bettger, J., Florian, J., Goldstone, R., Lassaline, M., Markman, A., Rubinstein, J., & Wisniewski, E. (1990). Safe take offs - soft landings. Cognitive Science, 14, 169-178.

Medin, D. L., Altom, M. W., Edelson, S. M. & Freko, D. (1982). Correlated symptoms and simulated medical classification. Journal of Experimental Psychology: Learning, Memory and Cognition, 8, 37-50.

Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family Resemblance, conceptual cohesiveness and category construction. Cognitive 4, 19, 242-299.

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

Murphy, G. L. & Medin, D. L. (1985). Role of theories in conceptual coherence. Psychological Review, 92, 289-316.

Nosofsky, R. M. (1984). Choice, similarity, and the context of classification. Journal of Experimental Psychology: Learning, Memory, & Cognition, 10, 104-114.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, 115, 39-57.

Piaget, J. (1929). The child's conception of the world. New York: Harcourt Brace.

Putnam, H. (1975). The meaning of meaning. In H. Putnam (Ed.), Mind, language, and reality, Vol. 2. London: Cambridge University Press.

Quine, W. V. O. (1960). Word and object. Cambridge, MA: MIT Press.

Quine, W. V. O. (1977). Natural kinds. In S. P. Schwartz (Ed.), Naming, necessity, and natural kinds. Ithaca, NY: Cornell University Press.

Rosch, E. (1973). Natural Categories, Cognitive Psychology, 4, 328-350.

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

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, M. D., & Boyes-Braem, P. (1976). Basic objects in natural categories. Cognitive Psychology, 8, 382-439.

Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. Progress in Psychobiology and Physiological Psychology, 6, 245-280.

Rumelhardt, D. E., McClelland, and the PDP Research Group (1986). Parallel Distributed Processing. Volume 1: Foundations, Cambridge, MA: MIT Press.

Shepard, R. W. (1987). Toward a universal law of generalization for psychological science. Science, 237, 1317-1323.

Shepard, T. W. (1964). Attention and the metric structure of the stimulus space. Journal of Mathematical Psychology, 1, 54-87.

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

Smith, J. D. & Kemler-Nelson, D. G. (1984). Overall similarity in adults' classification: The child in all of us. Journal of Experimental Psychology: General, 113, 137-159.

Smith, L. B. (1981). The importance of the overall similarity of objects for adults' and children's classification. Journal of Experimental Psychology: Human Perception and Performance, 1, 811-824.

Smith, L. B. (1989). A model of perceptual classification in children and adults. Psychological Review, 96, 125-144.

Smith, L. B. & Sera, M. (1990). A developmental analysis of the polar structure of dimensions. (Under review)

Smith, L. B., Jones, S. S., & Landau, B. (1990). Nouns, adjectives, and perceptual properties in very young children's lexical strategies. (Under review).

Spelke, E. S. (1982). Perceptual knowledge of objects in infancy. In J. Mehler, M. Garrett, and E. Walker (Eds), Perspectives on mental representation. Hillsdale, NJ: Erlbaum.

Sugarman, S. (1983). Children's early thought. Cambridge, England, Cambridge University Press.

Sutherland, N. S. (1964). The learning or discriminations by animalls. Edeavour, 23, 148-152.

Thelen, E. (1989). Self-organization in developmental processes: Can systems approaches work; In M. Gunnar & E. Thelen (Eds.), Systems and Development: The

Minnesota Symposia on Child Psychology, vol.22, pp. 11-117. Hillsdale NJ: Erlbaum.

Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.

Tversky, B. (1985). Development of taxonomic organization of named and pictured categories Developmental Psychology, 21, 1111-1119.

Vygotsky, L. S. (1934/1962). Thought and language. In E. Hanfmann and G. Vakar, trans. Cambridge, MA: MIT Press.

Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. Cognitive Development, 5, 123-150.

Wellman, H. M. & Gelman, S. A. (1988). Children's understanding of the nonobvious. In R. Sternberg (Ed.), Advances in the psychology of intelligence, Vol. 4. Hillsdale, NJ: Erlbaum.

Werner, H. (1948). Comparative psychology mental development, 2nd Ed. New York: International Universities Press.

Wohlwill, J. F. (1962). From perception to inference: A dimension of cognitive development. Monographs of the Society for Research in Child Development, 27, 87-112.

Younger, B.A. (1990). Infants' detection of correlations among feature categories. Child Development, 61, 614-620.

Younger, B. A. & Cohen, L. B. (1983). Infant perception of correlations among attributes. Child Development, 54, 858-867.