ELSEVIER

Cognition 65 (1998) 103–135

COGNITION

Two dogmas of conceptual empiricism: implications for hybrid models of the structure of knowledge

Frank C. Keil*, W. Carter Smith, Daniel J. Simons, Daniel T. Levin

Department of Psychology, Uris Hall, Cornell University, Ithaca, NY 14853, USA

Abstract

Concepts seem to consist of both an associative component based on tabulations of feature typicality and similarity judgments and an explanatory component based on rules and causal principles. However, there is much controversy about how each component functions in concept acquisition and use. Here we consider two assumptions, or dogmas, that embody this controversy and underlie much of the current cognitive science research on concepts. Dogma 1: Novel information is first processed via similarity judgments and only later is influenced by explanatory components. Dogma 2: Children initially have only a similarity-based component for learning concepts; the explanatory component develops on the foundation of this earlier component. We present both empirical and theoretical arguments that these dogmas are unfounded, particularly with respect to real world concepts; we contend that the dogmas arise from a particular species of empiricism that inhibits progress in the study of conceptual structure; and finally, we advocate the retention of a hybrid model of the structure of knowledge despite our rejection of these dogmas. © 1998 Elsevier Science B.V.

Keywords: Dogmas; Conceptual empiricism; Hybrid models

1. Introduction

As this issue's opening article by Rips and Sloman makes clear, mature concepts have two central components: one that is largely associative and one that is explanatory. Proposals for such a hybrid stretch back at least to Locke and are very much with us today (Neisser, 1967; Sloman, 1996). The general acceptance of the hybrid model of thought and knowledge unifies the papers in this special issue of *Cognition*. Yet, controversy remains about how these components are involved in cognitive

* Corresponding author. Tel.: +1 607 2556365; fax: +1 607 2558433; e-mail: FCK1@Cornell.edu

PII S0010-0277(97)00041-3

development and in adult learning and use of concepts. These controversies are fueled by two distinct, but related assumptions (or Dogmas) that dominate most views of the origin and structure of concepts.

Dogma 1: Any new category is understood by first processing similarity and only later by considering causal or explanatory principles. The explanation-based component only arises given sufficient opportunity for cognitive reflection.

Dogma 2: Infants and young children initially represent categories using the association component of concepts. Only with development does the more abstract, explanatory component of concepts emerge.

Both of these dogmas presuppose that learning is perceptually driven, progressing from the simple processing of sensory features to forming complex representations and abstract thoughts; this presupposition eventually dooms them as reasonable models of knowledge acquisition. Indeed, we argue that the two dogmas fail both under principled considerations and under the bright light of experimental data. Yet, these dogmas are implicit in most models of concept development and of adult processing of categories. If these dogmas fail as explanations of category and concept acquisition and use, why, then, are they so prevalent and widespread? We show how empiricist biases influencing the choice of stimuli used to study concepts and the interpretations of data from these experiments have perpetuated and sometimes even reified these dogmas. First, however, we must briefly review the need for both the similarity and the explanatory components of concepts.

2. Do similarity and explanation account for separate aspects of concepts?

Many aspects of human thought seem to rely on relatively automatic processing, driven largely by the statistical properties of instances. Some aspects of our categorization, induction, and concept acquisition clearly are influenced by tabulations of feature frequencies (e.g. how frequently a feature occurs among members of one category versus others) and probabilistic comparisons to stored exemplars (see Smith and Medin, 1981 for a summary of earlier work). Indeed, the emergence and successes of the Roschean view of concepts in the 1970s suggested that probabilistic representations might sufficiently explain all of categorization, displacing earlier rule-based ‘classical’ models (Rosch and Mervis, 1975). These successes were particularly seductive because such models were able to account for much of categorization without delving into ‘cause’ and ‘explanation’. Indeed, because the typicality of a feature often is directly related to its causal importance in determining category membership, models of categorization based on feature typicality alone seemed able to eliminate the explanatory component of concepts. For example, whiteness is both highly typical of and causally central to polar bears. Causally central features are those that are closely tied to other critical properties and are linked to the origins or essence of the object. ‘White’ is causally central to polar bears because the property is closely linked to survival in snow covered environments. Any other color would impair a polar bear’s ability to hunt. Less central properties (e.g. a bear’s tail shape) could more easily be variable without affecting the underlying nature or the basic behaviors of the bear. The causal mechanisms (evolutionary in this case) underlying the color of the polar bear explain why whiteness is a typical property. Given the substantial correlation between typicality and causal structure for real world categories, experimenters may miss a critical aspect of our concepts unless they empirically separate these components.

In many other cases, however, the properties most typically associated with members of a kind are not the most causally or explanatorily central. For example, virtually every washing machine ever encountered is white, yet we know that whiteness is irrelevant in determining that an object is a washing machine. Thus, although we do use typicality in categorization and induction, in many cases something other than mere similarity is needed, at least partly because some typicality information is not relevant. That 'something else' involves intuitions about why things are similar—about the causal forces underlying the similar properties themselves. Moreover, even variable properties can be causally central if the patterns of variation are linked to causal features of underlying category membership. For example, the variable colors of chameleons help determine category membership. Given that chameleons can change their colors, we rarely if ever encounter two chameleons that share identical coloring. Thus, the color variability for an individual chameleon leads to variability across chameleons as well. The variability itself is the feature that ties individuals to the larger category. Successful categorization requires us to ignore the typicality structure and rely on causal explanations for how color is involved in chameleon ecology. The same point can be made in cases where variability is even more lawfully and predictably related to causal interactions with the environment. For example, the color of iron varies considerably as a consequence of its surface being oxidized into rust. This variation is intimately linked to the interaction of iron's chemical nature and the environment. Even social categories such as 'teenager' can lead to variable, but nonetheless often depressingly predictable patterns of appearance. A concept of 'teenager' relies on causal accounts of how teens interact with their peers and adult society.

In some domains, the difference between similarity and rule-based categorization is even more striking. Categories such as odd numbers, prime numbers, and triangles are governed by formal, precise rules that transcend associative relations. Adults know and understand these logical constraints, but still easily and consistently judge some exemplars of these categories to be better than others (Armstrong et al., 1983). That is, even categories that are clearly defined by explicit rules can show typicality- or frequency-like effects. In other domains as well, we can clearly separate judgments based on typicality and on rules. For example, we can all appreciate that kinship terms such as 'uncle' have a dual representation. The typical uncle is a man about the same age as your parents who brings you gifts or sends cards on major holidays and often talks about things he and your mother or father did together as children. The rule-defined uncle is an individual of any age who shares a biological relationship to your mother or father. Whether or not he visits or brings gifts is irrelevant. Likewise, we know about the kinds of typical events that make one New Years' Eve similar to another, but we also know a rule defining that holiday.

Although the contrast between rules and typicality seems stark, we cannot fairly argue that these components are so discrete and independent that they function with little or no interaction. In fact, the cases described highlight the danger of generalizing from only one class of things. With nominal kinds such as triangle and uncle, the components may be fairly separable. However, for natural kinds, the interdependence of typicality and explanation is much more evident. If natural kinds are the most prevalent and typical type of category, such interdependence might well be the norm.

3. Natural kinds: the necessity of the hybrid

Knowledge of natural kinds clearly illustrates the need for a hybrid model of concepts. Surely the first human knowledge was largely of natural kinds in a world empty of artifacts. Thus, it is safe to assume that our first ways of meaningfully carving up the world and understanding its structure centered around our knowledge of natural kinds. Today, our knowledge about natural kinds is closely linked to the contemporary sciences, but how did these sciences ever get started? What cognitive processes account for the success of science in making predictions and increasing our understanding of the natural world? This question underlies much of the modern history and philosophy of science (Salmon, 1989). It should also be one of the foundational questions of cognitive science.

Science works because it links observations of associated variables with explanations of the causal mechanisms underlying those associations. Natural science is a union of the associative and the explanatory, of similarity and cause. For example, consider the class of metals occupying column 1B in the periodic table (i.e. copper, gold and silver) which tend to share many physical properties (reactivity, volatility, and malleability, etc.). We attribute these physical similarities to the nature of their electron shells: the causal mechanisms underlying the observable similarities depend on the behaviors of the electrons. Yet, even if the physical properties of these metals can be explained in terms of their electron shells, questions about the causal mechanisms responsible for the behavior of electrons remain. Diligently, we turn to particle physicists to find an explanation. They attribute the similarities to causal interactions of subatomic particles like quarks. We can proceed in this manner (e.g. what explains the behavior of quarks?) until eventually, this chain of linked mechanisms runs out and we must resort to less precise notions of what properties are causally relevant: ‘at the end of every explanatory regress we must perforce shift from causal mechanisms to causal powers’ (Harré, 1988, p. 142). Causal powers refer to our knowledge about the dispositions of entities to engage in some kinds of causal interactions but not others (see Harré and Madden, 1975; Harré, 1988). The distinction between causal mechanisms and causal powers mirrors the distinction between specific theories and framework theories (Wellman, 1990), where the frameworks constrain the kinds of information likely to be incorporated into a specific theory; that is, they highlight which properties are likely to be relevant to a particular causal explanation (see Brown, 1990; Gelman, 1990). With the accumulation of new observations, specific theories and explanations develop and sometimes are reinterpreted in different theoretical terms. These new observations often take the form of correlational evidence. Thus, further theory growth relies on an ever-growing similarity database. But this database is itself at least partially constrained by the assumption that similarities and correlations occur for a reason. We either seek to incorporate the observations into existing specific theories, or occasionally, if anomalous observations begin to accumulate, we attempt to significantly revise those theories (Kuhn, 1962).

For natural kinds then, there are at least three levels of conceptual understanding: precise knowledge of the causal mechanisms by which properties and kinds interact, notions of what kinds of properties are causally central in a domain (as well as what kinds of causal patterns and dispositions might be associated with a kind), and finally, a database of properties that tend to co-occur within a kind (i.e. the 'similarity' component). In terms of the hybrid structure of concepts, the explanatory component consists of both knowledge of precise mechanisms and framework notions of causal powers and patterns. This explanatory part draws on the 'similarity database.' New additions to that database, however, do not arise from a theory-neutral stance, but are in turn constrained by the explanatory component. That is, some similarities will be more readily noticed than others because they are more congruent with available systems of explanation. Each component, therefore, supports and helps guide the other. Indeed, without such mutual constraints each component would be dysfunctional.

Why do appeals to raw, theory-neutral similarity fail? Why cannot progressively higher-order tabulations of similarity fully account for our conceptions of natural kinds? The reason is simple: there are too many possible frequencies and correlations to be tabulated. Indeed the set of such possible tabulations is theoretically infinite. As has now been discussed in many places (Goodman, 1972; Keil, 1981; Murphy and Medin, 1985), depending on the features chosen for comparison, any two arbitrarily chosen objects can be maximally similar to each other. For example, a cloud and a white feather might be considered quite similar if color and ability to float in air are the comparison features, whereas they could not be more different if size and functional affordances are the comparison features. In the real world, every instance of every object we encounter has an indefinitely large number of features, any of which could be tabulated. Any attempt to store information about objects in terms of feature frequencies would fail without constraints on the inputs to such tabulations. Moreover, the whole notion of features as atomic primitives is suspect given that they can either proliferate or evaporate depending on the theoretical perspective adopted (see Wisniewski and Medin, 1994 for discussion). Thus, the idea of a theory-neutral set of primitive features is a fiction that, when acknowledged, makes the need for an explanatory component all the more obvious.

Empirical evidence bolsters this plausibility argument. In learning artificially constructed categories, in which the features are theory-neutral, adults are poor at detecting feature correlations even when doing so would enhance concept learning (Murphy and Wisniewski, 1989). Intuitive theories help people to assimilate statistical correlations (Wright and Murphy, 1984), and such intuitive notions of causation allow false beliefs to persist even when the statistical evidence contradicts the belief (e.g. Chapman and Chapman, 1967, 1969; Wright and Murphy, 1984). Thus, bottom-up statistical patterns do not always drive reasoning: we often use high level schema to impose interpretations on statistical patterns (see also Nisbett et al., 1983). Finally, in seeking to attribute the cause of an effect to some factor, adults prefer information about specific, underlying mechanisms over information about the covariation of surface events (Ahn et al., 1995).

In short, models of knowledge representation based solely on associative learning among properties are hopeless without some way of sharply limiting the set of properties to be associated. Empiricist approaches to knowledge representation cling to the assumption that such constraints will be imposed by the perceptual system. Accordingly, our sensory systems must provide a limited set of perceptual primitives so that exhaustive tabulations over those primitives are computationally possible. This assumption is nothing more than an article of faith and has never been even faintly satisfied. Even with a stock of a few hundred primitive features (e.g. colors, shapes, textures, sizes, and surface patterns), the number of potential associations to notice becomes massive. Add temporal and spatial factors to the tabulations (e.g. one feature is noticed at roughly the same time and place as another) and the computational task explodes once again.

Without constraints on the ‘tabulation space’ sufficiently more powerful than those imposed by the sensory system, associations alone cannot explain categorization of natural stimuli. In much of human cognition, those additional constraints come through rules, explanatory relations, and notions of mechanism. In other species, they may be less explicit and rule like, but still are radically different from pure association. If concepts were learned solely through association, those properties that co-occur more frequently should be more readily associated. Yet, classic work on phobias and taste aversion suggest that a single co-occurrence may indelibly link two properties or events. Many species form ecologically important links between properties of objects in the world in ways that run against the frequency of occurrence. Rats will associate nausea with the taste of a food ingested 8 hours earlier rather than with a physical trauma occurring minutes earlier (Garcia and Koelling, 1966; Rozin and Kalat, 1971). Also, monkeys more readily associate fear with snakes than with rabbits or flowers even when they have had equal exposure to both classes of entities (Cook and Mineka, 1989, 1990). These preferred linkages cannot be explained easily by different saliency of features; instead, the animal is predisposed to be especially sensitive to those associations that make sense for that organism’s ecological niche.

Humans show similar constraints in their understanding and categorization of natural kinds; higher-order criteria determine the particular frequencies and correlations we choose to notice and remember. When learning about a new kind of animal, we expect certain property relations to be causally and explanatorily useful. These expectations can override similarity-based information. Although the nature of these constraints is still unclear (they could take a range of forms from causal intuition to precise rules), our concepts of natural kinds are a blend of frequency information and constraints on the particular relations we choose to consider. In fact, we could not have concepts of natural kinds without both components, any more than science could proceed without both theoretical expectations to guide our inquiry and feature tabulations to gather as yet uninterpretable information. Hybrid concepts are necessary to adequately represent any natural kind.

The hybrid model of concepts occupies the middle ground in the continuum offered by Rips and Sloman and acknowledges that both similarity and rules are needed without reducing one to the other. Adopting a hybrid model neither entails a particular representational format for each component nor specifies the nature of the interaction between the components. We still need to know more about how the two components interact and how each is represented. Over many different versions of the hybrid model of concepts, the two dogmas outlined above are the primary assumptions that have shaped current theory and research. Our main purpose here is to place these dogmas in a broader context and to discuss why they must fail for any model of natural kind concepts.

4. Dogma 1—New categories are initially similarity-based: only with time do we come to apply rules and explanations

Although, at first glance, it seems plausible that adults acquire new knowledge via an association to theory shift, research on causal attribution and the development of domain specific expertise in adults directly contradicts this characterization. This research shows that causal theories often exert a powerful influence on the acquisition process itself and furthermore that associationistic principles are not eventually replaced by explanatory knowledge but remain important even in experts' classification schemes.

Traditional models of causal attribution propose that people use information about the covariation of factors and effects to determine what caused an event (e.g. Kelley, 1967; Cheng and Novick, 1990, 1992). But an examination of the information search strategies people use to develop explanations for events challenges this view (Ahn et al., 1995). In seeking to learn why an event occurred, people solicit information about hypothesized causal mechanisms much more often than they seek information about co-occurring events (Ahn et al., 1995). For example, when trying to discover why John had a car accident on Route 9 last night, people tend to ask questions that either propose specific underlying mechanisms (e.g. 'Was John drunk?' 'Did John's brakes fail?') or that seek more information about the elements described (e.g. 'Was there something wrong with John's car?' 'Was there something peculiar about last night?'). Generally, they do not ask questions about the frequency with which John has car accidents, the frequency of car accidents last night, or the frequency of accidents on that particular road. So, in learning about a novel event, people do not often seek information about covarying factors, nor do they seem to use such information when given it directly. However, as Ahn et al. (1995) point out, these findings do not diminish the importance of covariation information: such information may be used to generate hypotheses or to confirm the presence of a particular underlying mechanism.

Covariation models of causal attribution, with their emphasis on common and distinctive ‘factors’ are analogous to similarity-based accounts of concept formation, in which categories are learned by attending to common versus distinctive features. A still-thriving tradition of research on adult concept learning is based on the premise that raw computations of similarity relations adequately describe an early state of concept learning. Much of this research uses simple, semantically meaningless stimuli that vary continuously on two dimensions to test predictions of universal exemplar models of category structure. For example, the Generalized Context Model (GCM; Nosofsky, 1984, 1986) predicts classification and recognition behavior by assessing the similarity between a test stimulus and stored representations of other, previously acquired exemplars (i.e. the two categories of artificial stimuli used in a given experiment). This similarity measure is a weighted function of the features across which a test stimulus is compared to stored exemplars. Stimuli that are very similar to a large number of stored exemplars from either category will provoke a recognition response, whereas stimuli that are more similar to exemplars from one category than those of another will be classified accordingly. This ability to predict performance across a number of tasks using simple generalizable principles is one of the most important aspects of the GCM model and other similarity based models.

Nosofsky (1992) provides the groundwork for a plausible interface between these generalizable similarity computations and a more rule-based component of concepts by suggesting that exemplar models be understood in terms of representation-process pairs (Anderson, 1976) whereby representations are given separate status from the processes that work on them. Accordingly, similarity-based exemplar spaces represent a stable ground upon which various processes (e.g. summed similarity) operate to produce predictable patterns of behavior. This kind of model implies that associative principles are fundamental in organizing concepts, but that theory laden processes can supplement these basic representations to produce a complete explanation of more complex concepts (such as real world categories). The most concrete mechanism purported to govern the interaction of the associative and explanatory components is selective attention. Both the GCM and the earlier Context Theory (Medin and Schaffer, 1978) suggest that selective attention changes the weighting given to each stimulus feature in different task contexts. Thus, selective attention takes into account the relevance of particular features to particular classification goals. Therefore, a resulting hybrid model of concepts might borrow this distinction between representations and processes by positing theory and explanation as processes akin to selective attention that operate a posteori on stable associative networks by determining which features will be most heavily weighted in a given comparison. Advocates of this kind of model appeal to a combination of basic perceptual filters and task-specific principles to constrain the properties represented and thus render the similarity comparison process more plausible than a brute force associative engine (Goldstone, 1994; Medin et al., 1993).

Such models directly instantiate Dogma 1 whereby learning begins as an accumulation of a series of exemplars encoded according to domain-general associative principles and is gradually enriched by attentive processes that select critical features more carefully, which, in turn, drives the emergence of more sophisticated, perhaps causal, principles relating those features. This progression seems to account for the learning of categories of artificial stimuli devoid of real-world meaning. The use of such stimuli effectively eliminates the possibility that existing explanatory knowledge can contribute to the organization of novel categories. Thus, this instantiation of Dogma 1 seems to follow by necessity: it is difficult to imagine how a rule-based ‘process’ can function until a sufficient number of exemplars have been coded associatively into a sufficient ‘representation’.

However, when items in a novel category do evoke real world knowledge, the importance of causal theories and explanations to even the earliest stages of the acquisition of new concepts becomes evident. Unlike the arbitrary categories described above, natural kinds and even humanmade implements have features that tend to be causally interwoven. When features are causally interrelated, explanatory knowledge speeds both category learning and the identification of features (Murphy and Allopena, 1994). In Murphy and Allopena’s study, participants learn to classify novel animals with either interrelated or non-interrelated features. For example, one animal might have ‘pointed ears’ and ‘spots’ and another animal might ‘eat meat,’ and have ‘sharp teeth.’ Having pointed ears does not causally entail the presence of spots or any other surface marking. However, ‘eating meat’ likely entails the presence of ‘sharp teeth’. Species that eat meat have generally evolved sharp teeth. Participants who learn categories with causally interrelated features are faster and more accurate at identifying isolated features than subjects who learn non-interrelated features. Moreover, causally interrelated, but infrequently mentioned features are identified just as quickly and are rated as being just as typical as frequently mentioned features. In contrast, participants who learn non-interrelated features rate infrequently mentioned features as atypical and identified such features more slowly than frequently mentioned features (Murphy and Allopena, 1994). This effect is not likely to be merely a consequence of earlier stored correlations about the features of animals. Even if one taught a new, plausible causal mechanism, it should produce the same effects in a replication of the Murphy and Allopena study.

When the features available for learning novel instances of natural categories are interrelated, people use knowledge of these interrelations to guide learning a new category without having to attend to raw feature frequencies. But when the features are not interrelated, people must resort to strategies using only the raw statistical frequencies of features. Thus, knowledge about a domain can aid in learning about new categories in a domain.

Knowledge from one domain can also be used during the initial learning of another domain, as shown by research on the acquisition of expertise. Not only do novices immediately use domain-specific causal theories, they also recruit theories from other domains when analyzing new information. Learners seem compelled to immediately organize new information using causal theories. Yet similarity is still central to concept learning; it highlights associations needing further explanation. Thus, explanation and association both play important roles throughout the course of learning.

In one of the best studied areas of expertise, medicine, there is ample evidence that novices use causal knowledge right from the start. The presence of such theories is perhaps most evident in the reasoning errors made by both novices and experts. Beginning medical students assimilate and misconceptualize new medical knowledge because they rely on causal theories learned in school and everyday life. For example, students may recruit knowledge about kitchen plumbing when trying to learn properties of the circulatory system (Feltovich et al., 1989). As a result, they fail to understand the contribution of the elasticity of arteries in amplifying the power of the heart to pump blood because kitchen pipes are rigid. Even experts in medicine sometimes mistakenly assume that an enlarged heart implies overstretching of its muscle filaments because they incorrectly recruit knowledge of skeletal muscles which can become stretched enough to lose efficiency (Feltovich et al., 1989).

In addition to recruiting knowledge from other domains novices rely heavily upon causal theories within the new domain. They may even recruit causal theories more frequently than experts. For example, Boshuizen and Schmidt (1992) asked students and experts to diagnose case summaries, and found that students engaged more frequently in explicit biomedical reasoning. When asked, experts could produce more biomedical information, but only relatively infrequently did they spontaneously refer to this information in their protocols. The acquisition of medical expertise begins with biomedical knowledge (i.e. rules and theories), and only later is enriched by associations with the examples gained from practical experience (Boshuizen and Schmidt, 1992). Even in cases where the medical school curriculum emphasizes an exemplar-based approach, specific examples are almost always connected with pre-existing, causal, biomedical knowledge (Norman and Schmidt, 1992).

Experts may even use associations more rapidly than novices. Novices appear to use a seemingly unwieldy process of backward chaining in which reasoning progresses from initial hypotheses to the observation of disease features and back to a subsequent revision of the diagnosis if the features do not match the hypothesis (Patel and Groen, 1991). Experts, on the other hand, appear to use forward chaining in which reasoning jumps directly from features to hypotheses or disease schemas (Patel and Groen, 1991). This procession from features to hypotheses is typical of an associative process, yet it is observed more frequently in later stages of the learning process, not in the initial stages.

Although the medical knowledge literature supports neither a shift from exemplars to causal theories nor a shift from perceptual to conceptual processing, other contrasts between experts and novices have been taken to support Dogma 1. The best known example comes from the classification of physics problems. Chi et al. (1981) found that experts classify problems according to deep structural principles (e.g. the kinds of physical laws applicable to solving the problem) but that novices classify based on surface features (e.g. the kinds of objects involved in the physical mechanism under examination). Expert and novice math problem solvers (as defined by Scholastic Aptitude Test scores) also seem to differ in their use of association (Novick, 1988). Expert students showed positive transfer from an example word problem to a new one which shared a structural component—the particular mathematical solution procedure. Novices, on the other hand, attempted to transfer based on a match between the surface features of the old and new problems and therefore often selected the wrong solution procedure.

In both examples, however, the things identified as ‘surface’ features are more complex than primitive perceptual attributes: ‘surface features’ include non-concrete terms and structural relations between objects. For the physics problems, surface structure was defined as ‘(a) the objects referred to in the problem (e.g. a spring, an inclined plane); (b) the literal physics terms mentioned in the problem (e.g. friction, center of mass); or (c) the physical configuration described in the problem (i.e. relations among physical objects such as a block on an inclined plane)’ (Chi et al., 1981, p. 125). In Novick’s (1988) mathematical problem solving task, two problems were said to share surface similarity if they involved the same kind of activity (e.g. gardening). The characterization of such abstract, relational features as mere surface features seems to lead inherently to an underestimation of the depth of processing in which novices engage. Although it might be the case that novices encode problems more shallowly than experts, it is equally plausible that novices are relying on a different (possibly more domain-general) set of causal principles to code the problems. If this were true, then describing novice’s reasoning as ‘shallow’ and more instance-bound would be a mischaracterization.

For example, college students who learn to solve a physics problem involving a constant change in the speed of some object will often spontaneously transfer the appropriate equation from this base problem to a target problem concerning a constant change in the rate of population growth. These same participants are less likely to use this equation to solve a problem concerning a constant change in the rate of increase in attendance at an annual fair (Bassock, 1996). Because participants represent speed and population growth as continuously changing quantities, whereas they represent increases in attendance at an annual fair as a discretely changing quantity, transfer between the former problems is less difficult than transfer from the physics problem to the attendance problem (Bassock, 1996). Thus, novices may engage in sophisticated causal reasoning and mental-model building even if they do not happen to exploit the particular mathematical principles or structural correspondences considered important by the experimenter. Furthermore, the activities described in a problem often *do* signal a particular formal procedure: expert problem solvers readily use such content, or ‘surface,’ cues to select appropriate solution schemas (Blessing and Ross, 1996).

In any problem solving situation, even novices exclude certain features from consideration. For example, in assessing the force required to push a block up an inclined plane, what novice would seriously consider the color of the block, the gender of the pusher, or whether the plane ends in the back of a Ryder or a Hertz rental truck? Rather, the typical novice would employ a whole series of implicit and explicit theories about the properties of the plane’s surface, the tilt of the plane, and the weight of the object. A strict empiricist might argue that the process of importing this information is one of sophisticated association. Perhaps the learner has directly experienced the process of pushing something along a surface, walking up a plane, and carrying heavy and light objects. Therefore, each problem element might be associated with a relevant experience. Such a model would, of course, have to be modified to include more abstract associations than simple perceptual features. It would require representations of prior problem solving experiences with particular ways of selecting features. Such an approach is, of course, not new to learning theory in either its behaviorist (e.g. the concept of a discriminative stimulus) or neural net instantiation (e.g. adaptive resonance theory; see Carpenter and Grossberg, 1987).

However, by allowing more abstract associations, this model cannot explain behavior solely through the interaction of a domain-general learning system with the structure of the environment. One of learning theory's major goals is to explain behavior without positing complex representations or domain-specific rules. The need for constraints on feature selection forces learning theory to adopt just this type of complex, domain-specific apparatus. Whether feature selections are driven by domain-specific theories or discriminative stimuli, both require modifications to supposedly objective associations in order to allow the domain-general apparatus to accommodate domain-specific regularities. Such modifications imply that existing knowledge, whether it be association weights, connections, propositions, or images, cannot be ignored when modeling even the simplest interactions with a stable environment. To claim that a carefully structured neural network that has been fed precisely selected information can model behavior only using domain general learning rules is to purposely black-box all of the domain-specific processes that would be necessary to get the net to explain anything about real world behavior.

In summary, the development of expertise in adults does not support Dogma 1. Not only do novices immediately recruit domain specific causal theories or theories from other domains when acquiring new concepts, but experts continue to rely on a blend of similarity-based computations and explanatory reasoning. Association and similarity remain important throughout the course of learning, neither predominating nor becoming relegated to the periphery. Clearly, similarity, perceptual or otherwise, can play an important role in isolating potentially important correlations when theory runs out or an initial hypothesis is contradicted. But similarity will rarely be the sole basis for conceptualizing: in real world contexts, expertise develops from a symbiosis of preconceived theory and experience with new material.

In addition to the expertise differences explicitly adopted by Dogma 1, a second, less frequently mentioned assumption was prominent in many earlier discussions of concepts and is still tacitly presupposed in contemporary discussions. The assumption rests on the following line of reasoning: (a) The similarity-based component of concepts is simpler than the explanation-based aspect (this assumption is clearly linked to the belief that novices are driven by similarity); (b) Simpler aspects of cognition are performed more rapidly than complex aspects; (c) Simple, quickly formed aspects of cognition often serve as preprocessing for more complex aspects; (d) therefore, in real-time use of concepts (e.g. categorization and induction), the similarity-based component is processed first, and only later does the explanation-based component come into play.

The idea that initial 'rough and ready' or 'quick and dirty' processing relies more heavily on similarity was proposed explicitly in the 1970s (Smith et al., 1974) and has been left largely unchallenged by contemporary work. Yet, this assumption may not describe real-world cognition. Even in the briefest glance at a scene, viewers could have theory-based expectations about which features count. For example, a brief glimpse of a whale may lead to the perception of a fish or of a mammal depending on which features the viewer considers given a prior theory that emphasizes some features over others. Of course ignorance may lead to mistakes (e.g. observers may misclassify bats as birds and dolphins as fish), but mistakes reflect the use of different theory-based expectations not a complete reliance on similarity. Our perceptions are often influenced by what was causally and explanatorily relevant in the past. Only with stimuli so arbitrary and decontextualized that no prior information could be relevant will the 'similarity first' rule of processing apply.

Although the conclusion that initial processing of new entities is driven solely by perceptual primitives and the laws of association may be correct under such decontextualized circumstances, we cannot generalize this conclusion to cases of conceptual apprehension in the real world. In some intriguing studies, researchers have found that people with psychological disorders show Stroop-interference effects for words related to their own individual diagnosis (for review, see Williams et al., 1996). These studies suggest that a priori thoughts or emotions can create schemas that affect early stages of perceptual processing. Research on natural scene processing provides further evidence suggesting that the initial steps in encoding objects and scenes involve mutual constraint satisfaction in which scene context facilitates object recognition (Biederman et al., 1982; Rayner and Pollatsek, 1992). Other research on rapid identification of pictures shows it is possible to detect an object that does not belong to a target category even when each image is visible for only 114 ms (Intraub, 1981). To detect an exemplar that is not a member of a specified category requires participants to activate category knowledge at a higher level than perceptual similarity; perceptual similarity is less useful in this case because the participant does not know what features define the target. Therefore, on-line expectations about upcoming events can direct early visual processing in searching for something as abstract as the non-occurrence of a given category of object. In both of these cases, early object recognition processes are sensitive to the surrounding scene context and current task demands, which implicate the use of more abstract knowledge than perceptual primitives. We are not resurrecting aspects of the 'new look' in perception (e.g. Bruner and Goodman, 1947; Bruner and Postman, 1947), a movement that foundered in a methodological morass, but we are suggesting that abstract conceptual information is involved in the interpretation and use of information from the earliest stages of processing.

5. Dogma 2—Children initially rely on associations and only later begin to use theories to constrain their concepts

The notion has long persisted that children have different patterns of thought and different concepts than adults. Such apparent differences suggest the need for a transition between undeveloped and mature thinking. The transition most frequently proposed to account for these differences is one in which similarity precedes and acts as the cradle for rules and explanation. In this view, knowledge emerges through ever more complex associations over perceptual primitives. The child is seen as initially nothing more than a frequency and correlation detector, and the massive collections of associations eventually yield impressions of causality, rules, and explanation. In other words, young children are initially instance bound, but gradually develop principled ways of understanding. Many theorists have argued for a shift from concrete, instance-based knowledge to abstract, rule-based concepts. For example, Vygotsky (1962) argued that the development from similarity to rules parallels the acquisition of language; language serves as a tool for stating and using rules. Others argued for a shift from concrete representations to abstract ones (Werner and Kaplan, 1963) or from accidental features to essential ones (Inhelder and Piaget, 1969).

Despite the prevalence of the concrete to abstract shift in models of development, no one has ever been able to describe how senses of explanation, mechanism, and cause gradually arise out of statistical operations over primitives. Models based on the second dogma offer no plausible account of how the second part of the hybrid might emerge from the first. Threshold models in which associations of sufficient strength become causal cannot account for our ability to reject causal explanations for some cases of perfect, but indirect or arbitrary correlations. Thus, the shadow of a flagpole at any given time of day may be perfectly correlated with its height, but no one assumes that the shadow of the flagpole causes it to have a particular height. In the 1950s in the US, hair length was almost perfectly correlated with gender, but no one thought that gender caused hair to be of different lengths or that hair length somehow caused an individual to be a certain gender. In contrast to such threshold models, which do not capture our ability to posit mediating variables or to reject nonsensical explanations, recent papers have argued that theory often precedes, or operates in concert with considerations of similarity (e.g. Simons and Keil, 1995).

The second dogma also encounters the problem discussed earlier in the context of adult concepts of natural kinds. If children merely tabulate feature frequencies, how do they know which frequencies and correlations to encode? Even if the features considered are constrained to a relatively small set of perceptual primitives, the number of possible relations to tabulate seems overwhelming. Thus, in principle, children could not acquire a rule-based concept of a natural kind from similarity relations alone. Not surprisingly, then, research with human infants increasingly suggests a sensitivity to features and relations that are anything but the sorts of perceptual primitives embraced by those adopting Dogma 2. Young infants appear to perceive causal relations over mere contiguities (Leslie, 1995), to evaluate the rationality of an agent's actions in relation to its inferred goals (Gergely et al., 1995), to intermodally integrate information at birth (Gibson and Walker, 1984), to imitate high order relational patterns (Meltzoff, 1988), and to discriminate biological from non-biological motion (e.g. Fox and McDaniel, 1982). All of these abilities seem to require representations of interactions among features or entities. Apparently, infant concepts are not so impoverished that they are trapped by similarity of perceptual features.

Several recent findings from work with preschool-aged children also raise serious questions about the developmental progression described by Dogma 2 (Simons and Keil, 1995; Wellman and Gelman, 1992). For example, children's inductions about novel, unobservable properties seem to be based preferentially on information about category membership as opposed to information about perceptual similarity (Gelman and Markman, 1986, 1987). In such tasks, preschool children are shown an animal or artifact which is given a label. They are then told about an unobservable property or part of that object and are asked which of several other objects have that part. The test objects vary in their perceptual similarity to the original object and in whether they have the same label as the original. Thus, the task directly pits category membership against perceptual similarity. Dogma 2 would suggest that preschoolers form categories using perceptual similarity alone, and that only later do children begin to use more abstract rules to form categories. If so, then inductions of unseen properties should be based on the similarity of the original and queried object. However, even 2-year-old children override perceptual similarity and attribute unobservable parts to members of the same category (Gelman and Coley, 1990; Gelman and Markman, 1986). These children also seem to know which properties should be inferred on the basis of category membership; they know that category membership is more central than perceptual similarity. Although these findings do not directly speak to the developmental origins of preschooler's concepts, they do suggest that children's categories are based on more than perceptual similarity. The findings suggest that 'even before children can make use of subtle perceptual cues to determine category membership, they readily use category labels as the basis of their inferences' (Gelman and Coley, 1990, p. 803).

Other evidence suggests that preschoolers develop knowledge of unobservable internal parts in the opposite order to that suggested by Dogma 2. Children initially have abstract expectations which guide their search for more concrete information (Simons and Keil, 1995). In these studies, children were shown a target animal or machine along with a set of potential 'insides': the insides of an animal, of a machine, a pile of blocks, and a pile of rocks. On each trial, children were asked to pick the insides that belonged with the target animal or machine. Even preschool children expected the insides of animals and machines to differ, but they did not have specific expectations for the physical appearance of the insides. Children consistently chose different insides for animals and machines. However, their mistakes across several studies were particularly revealing. In one study, when they picked the wrong insides for animals, they picked the rocks rather than the machine insides. In a second study, when they chose the wrong insides for machines, they tended to choose the blocks rather than the animal insides. Although children lack concrete knowledge of the insides of animals and machines, they have abstract expectations about the differences between these categories and use the expectations to guide their inductions about unfamiliar properties. By 8 years, children have acquired enough concrete knowledge that they rarely choose the wrong insides (Simons and Keil, 1995). Abstract expectations about the differences between animals and machines guide children's search for more concrete details; concepts of animals and machines become more like those of adults by refining broad initial expectations with specific concrete details. Although studies of children's knowledge of insides do not directly assess the relationship between perceptual similarity and rule based components of concepts, they do suggest that young children are guided by rule-like expectations for the differences between superordinate level categories. Preschoolers know that only natural kind insides belong in animals and only artificial insides belong in machines. They can use this abstract distinction to draw inferences without having to rely exclusively on perceptual groupings.

Some have argued that demonstrations like these are based on an excessively impoverished notion of perceptual similarity. As a result, early categorization is falsely portrayed as non-perceptual, and thus by necessity conceptual, when a richer view of perceptual similarity makes clear the primacy of perception in development (e.g. Jones and Smith, 1993). We need not enter that debate here (but for related commentary see Barsalou, 1993; Gelman and Medin, 1993; Mandler, 1993; Mervis et al., 1993; and Smith and Jones, 1993), but we should note that these more elaborated forms of 'perceptual similarity' have little in common with the notion put forth in Dogma 2. Instead, these richer versions of perceptual similarity include highly relational and abstract structural properties and are heavily influenced by context and relevant prior experience. There may be good reasons for attempting to define a more principled perceptual/conceptual boundary, but they are beyond the scope of this article.

If infants are sensitive to more complex forms of information from the start, they must be biased to notice certain feature patterns over others. What, then, is the nature of these biases and how do they transcend the laws of a raw similarity system? One developmental model posits very crude initial sensitivity to relational patterns that pick out some limited domain of information such as the human face. These crude sensory biases isolate face-relevant information and then learning occurs via domain-general mechanisms (Johnson et al., 1991). Thus, infants might first pick out and attend to human faces by relying on the crude triangular configuration of the eyes and mouth (see Johnson, 1992). Such information may be sent to one region of the cortex which, other than being the endpoint for face related configural information, has no a priori specialization for face processing. This modern variant of empiricism accounts for distinct causal explanations by adopting minimal prior biases (possibly embodied in the neural architecture of the perceptual system) which help to separate distinct domains into separate representational systems.

An alternative approach argues for initial sensitivity to many complex patterns at both perceptual and cognitive levels. These might range from the particular temporal patterning of social interactions to having expectations about the mechanics of physical objects (Spelke et al., 1992). At a minimum, this approach requires early sensitivity to informational patterns which uniquely specify important kinds in the world. Positing sensitivity to complex patterns or invariants in the environment does not require that infants have theories, conscious or otherwise, or even sets of beliefs, but it does require a learning system that is tuned to particular kinds of causal patterning and that uses such tuning to build and structure knowledge in a domain. Despite these arguments and evidence against Dogma 2, it continues to resist extinction. One reason for its persistence arises from the tendency of some to mistake other developmental patterns for versions of Dogma 2.

6. Mistakenly interpreting developmental trends as evidence for Dogma 2

In recent years, more sophisticated versions of Dogma 2 have arisen that are not so easily rendered implausible. Rather than simply positing a shift from similarity to rules, recent models argue for more subtle transitions. For example, the acquisition of word meanings and their corresponding concepts seems to follow a 'characteristic-to-defining shift' (Keil and Batterman, 1984; Keil, 1989). Other similar developmental trends include a shift from processing surface to deep features (Chi, 1992) and a gradual transition from noting properties, to detecting relations between properties, and finally to detecting relations between relations (Gentner and Toupin, 1988). In all of these cases, the first kind of knowledge seems more like similarity and the second more like rules and/or explanations. Real developmental changes have inspired all of these claims (as well as all earlier ones), but a closer look at the nature of children's knowledge suggests that the changes have little to do with the pattern described by Dogma 2.

Consider, for example, the shift from characteristic to defining features. Young children regularly identify instances as members of a category on the basis of shared typical or characteristic features even when 'defining' ones are absent. For example, young children claim that a gift-giving, male, unrelated adult is an uncle, but an adolescent who is their father's brother is not. These children seem to be tabulating the features most frequently associated with uncles and neglecting any use of the kinship rule. Older children accurately reject such cases, but accept ones that have almost no characteristic features when a defining feature is present (Keil, 1989).

A closer look at the characteristic to defining shift, however, reveals that it is not truly a shift from similarity to rules, or from association to definitions. Even the youngest children never rely solely on tabulations of the most perceptually salient features associated with instances and adults rarely rely exclusively on rules or explanations. The features children tabulate in demonstrations of a characteristic to defining shift are always constrained; they limit their consideration even of characteristic features to those likely to be relevant. For example, even if every 'uncle' a child encounters is seen wearing glasses, it is unlikely that the child would weigh this feature as heavily as certain social, behavioral, and personality traits that are characteristic of uncles. Typicality is always harnessed to the particular explanatory framework currently active; which typical features are taken to be 'characteristic' depends critically on the child's cognitive 'point of view'.

More broadly, younger children can differ from older ones in two ways that give the illusion of reinforcing Dogma 2: (1) they know less about the world; and (2) they may have different biases about what information they regard as explanatorily relevant. In the first case, children might well be ignorant of specific causal mechanisms or rules underlying a phenomena and therefore may be forced to rely on notions of causal frameworks much sooner than adults. Furthermore, given the necessary vagueness of those frameworks, they may in turn have to rely more heavily on similarity to structure the information they encounter. However, they do not and could not ever use similarity exclusively. At most, there is a shift in the ratio of their use of similarity-based vs. explanation-based knowledge. But by neglecting the less precise aspects of explanatory knowledge (such as notions of causal potency), researchers might mistakenly attribute to children a full association to explanation shift. In the second case, when trying to assimilate new information, a child may implicitly adopt an inappropriate theory (e.g. psychological versus biological; see Carey, 1985) that makes it appear as if she is ignoring critical features and simply tabulating features indiscriminately when, in fact, she is using a theory, albeit one inappropriate to the domain.

7. Causal potency of properties: an initial study

If explanatory knowledge constrains initial tabulations of similarity, how is such knowledge represented? The early explanatory component of concept representation might consist of a discrimination between causally central and causally peripheral properties even when specific mechanisms are unknown. Such knowledge must distinguish between causal centrality and typicality of properties (even though they are often closely related). Knowledge of this distinction is easy to demonstrate in adults. For example, adults view the property ‘curvedness’ as equally typical of bananas and boomerangs, but do not regard it as equally central to determining category membership (Medin and Shoben, 1988). Although a straight ‘banana’ could still be a banana, a straight ‘boomerang’ would be a stick. Research on the causal centrality of properties typically asks participants to consider a hypothetical situation in which a typical property (curvedness) is not present, and then to judge whether this counterfactual vitiates the inclusion of an instance in a category. We have used such counterfactual questions to examine the properties that adults believe are causally central to categories such as animals, artifacts and non-living natural kinds. Importantly, judgments of centrality are distinct from judgments of typicality; the most typical properties are not always thought to be causally central (Keil and Smith, 1996).

Although several researchers have considered the distinction between typicality and centrality with adults, the development of notions of causal centrality has remained largely unexplored. Children under 6 years of age usually do not greet sets of counterfactual questions with much pleasure. Often, they simply refuse to entertain the counterfactual (but see Harris et al., 1996). However, by describing a novel category and posing less structurally complex questions, we have been able to explore intuitions about causal centrality in children as young as 5 years.

In the following study, we read stories about a novel kind of animal (‘glicks’) and a novel kind of machine (‘nilards’) to 5-, 7-, and 9-year-olds. Each story described six property types: size, weight, color, surface markings, number of important inside parts, and appearance of functional outside parts. Subsequently, we asked children if other instances of the described category had to share the same property as those described in the story, or whether something could still belong to the category if it differed on a particular property type (e.g. if a something had a different color than the glicks described, could it still be a glick?). Because both the machine and animal stories described the same property types, we were able to assess to degree to which children believed that a particular property type was relevant to a decision about category membership for different ontological kinds[1].

8. Method

8.1. Participants

Sixty-eight children participated: 22 5-year-olds, 24 7-year-olds, and 22 9-year-olds (mean ages, 5 years 6 months, 7 years 6 months, 9 years 6 months). Children were tested individually outside their preschool classroom. Each session lasted roughly 10 min.

8.2. Materials and procedure

The stimuli consisted of two animal stories and two machine stories. Each child heard one animal and one machine story, with half of the children of each age hearing a ‘glick’ and ‘nilard’ pair and the other half hearing a ‘bleek’ and ‘jullet’ pair. The artificial labels, ‘glick’ and ‘bleek,’ each referred to a novel kind of animal; ‘nilard’ and ‘jullet’ each referred to a novel kind of machine. The order of presentation of the stories was counterbalanced across subjects. Each story contained a set of property types that could in fact apply to either an animal or machine, and only an introductory statement told the children the category (animal or machine). Because each child heard an animal story and a machine story, the stories were not word for word duplicates in terms of the properties described: although both stories indicated the color of the ‘glick’ or ‘nilard’, in one story this color was black while in the other story this color was, for example, yellow. This design was necessary to avoid situations in which a child might confuse the two stories; however both stories contained the same six property types. (See Tables 1 and 2 for examples).

After hearing a story, children responded to a series of questions about the properties mentioned in the story. Specifically, they were asked whether the object they heard about would be the same kind of thing if a target property were changed. For example, they might be asked ‘Do you think that all Glicks have to have black stripes on their backs, or could something still be a Glick even if it didn’t have black stripes on its back?’ We explored six different property types presented in a pseudo-random order for each child: surface markings (e.g. stripes), the number of internal parts, the shape of external parts, color, size, and weight.

To make sure children understood the task, two control questions were asked: ‘Could something still be a Glick even if it didn’t have dirt on its tail?’ and ‘Could something still be a Glick if it was made out of butter?’ All but one child answered these control questions correctly (‘yes’ for the first question and ‘no’ for the second); this child’s data were not included in the analyses.

[1]To avoid confusion, we distinguish between property types and particular properties. That is, color is a property type, while ‘red’ is a particular property or instantiation of that property type.

Table 1
Animal story and questions read to 5-, 7-, and 9-year-olds

There is a kind of animal called a Glick. Have you ever heard of a Glick? I certainly had not, but the other day, I was walking through the woods, and I saw one. Do you want to hear what the Glick looked like? Well, it was brown, had black stripes on its back, was about this big [gesture] and weighed about 10 pounds. It also had 26 really important parts on the inside of it. The Glick I saw really liked to eat berries, and had four parts on the outside of it that it used to pick apart the berries. It also was sitting in a tree with 16 branches and had a little bit of dirt on its tail.	
Do you think Glicks really have to have dirt on their tails, or could something still be a Glick even if it didn't have dirt on its tail?	[Dirt]
Do you think Glicks have to be the same size as the one I saw, or could something still be a Glick even if it was a different size?	[Size]
Do you think that Glicks have to be brown, or could something still be a Glick even if it was a different color?	[Color]
Do you think that all Glicks have to have black stripes on their backs, or could something still be a Glick even if it didn't have black stripes on its back?	[Surface markings]
Do you think that all Glicks have to weigh the same as the Glick I saw, or could something still be a Glick even if it weighed something different?	[Weight]
Do you think that all Glicks have the same number of important inside parts, or could something still be a Glick even if it had a different number of important inside parts?	[# Of inside parts]
Do you think that all Glicks have the same kind of parts on the outside of it or could something still be a Glick and have different looking parts on the outside?	[Shape of outside parts]
Could something be made of butter and still be a Glick?	[Butter]

8.3. Results

A belief that the hypothetical object with the changed property was still a member of the labeled category was coded as '0', and a belief that changing the property precluded category membership was coded as '1'. These scores correspond to how relevant a child thought a particular attribute was to determining an object's category membership. And, when averaged across children, they provide the proportion of children of each age group who believed that property to be central.

The mean responses for each age group are represented graphically in 'spider web' plots (see Figs. 1–3). Each axis radiating from the center represents one property, and distance from the origin indicates the degree to which children judged the property to be central (i.e. that it had to be shared by all members of the labeled category). For example (see Fig. 1), whereas about 75% of the children (averaging across age) thought that all animals called 'glicks' had to have outside parts that looked the same, only 20% of the children thought that all machines called 'nilards' had to share similar looking outside parts.

A 3 (Age: 5, 7, or 9 years) by 2 (Kind: animal or machine) by 6 (Property: stripes, color, size, weight, # of internal parts, shape of external parts) mixed-design ANOVA with Kind and Property as repeated measures compared children's

Table 2
Machine story and questions read to 5-, 7-, and 9-year-olds

There is this kind of machine called a Nilard. Have you ever heard of a Nilard? I certainly had not, but the other day, I was walking through the hardware store, and I saw one. Do you want to hear what the Nilard looked like? Well, it was yellow, and had green stripes on it, was about this big [gesture] and weighed about 50 pounds. It also had 30 really important parts on the inside of it. The Nilard I saw is used to dig holes in the ground, and had 5 parts on the outside of it that would dig those holes. It was sitting on a shelf and had some dust on it.	
Do you think Nilards really have to be dusty, or could something still be a Nilard even if it wasn't dusty?	[Dirt]
Do you think Nilards have to be the same size as the one I saw, or could something still be a Nilard even if it was a different size?	[Size]
Do you think that Nilards have to be yellow, or could something still be a Nilard even if it was a different color?	[Color]
Do you think that all Nilards have to have green stripes on their backs, or could something still be a Nilard even if it didn't have green stripes on its back?	[Surface markings]
Do you think that all Nilards have the weigh the same as the Nilard I saw, or could something still be a Nilard even if it weighed something different?	[Weight]
Do you think that all Nilards have the same number of important inside parts, or could something still be a Nilard even if it had a different number of important inside parts?	[# Of inside parts]
Do you think that all Nilards have the same kind of parts on the outside of it or could something still be a Nilard and have different looking parts on the outside?	[Shape of outside parts]
Can something be made of butter and still be a Nilard?	[Butter]

responses to the property changes for each story[2]. The analysis revealed significant main effects only of Property ($F_{(5,58)} = 21.94$, $P < 0.0001$) and Kind ($F_{(1,62)} = 22.69$, $P < 0.0002$) and a significant Property × Kind interaction ($F_{(5,58)} = 21.87$, $P < 0.0001$)[3]. No other main effects or interactions reached significance, although the main effect of Age approached significance ($F_{(2,62)} = 2.481$, $P = 0.092$). Post-hoc analyses confirmed that each age group showed this same pattern of results. For each age group a 2 (Kind) by 6 (Property) repeated measures ANOVA revealed significant main effects of, and a significant interaction between, kind and property: For 5-year-olds, Kind $F_{(1,21)} = 3.80$, $P = 0.065$, Property $F_{(5,17)} = 4.83$, $P = 0.006$, Property by Kind $F_{(5,17)} = 6.39$, $P = 0.002$; for 7-year-olds, Kind $F_{(1,21)} = 7.09$, $P = 0.015$, Property $F_{(5,17)} = 6.97$, $P = 0.001$, Property by Kind $F_{(5,17)} = 7.91$, $P < 0.0006$; for 9-year-olds, Kind $F_{(1,20)} = 13.81$, $P = 0.001$, Property $F_{(5,16)} = 16.37$, $P < 0.0001$,

[2]Three participants (two 7-year-olds and one 9-year-old) were excluded from these overall analyses because their responses to one of the property questions were ambiguous and were coded as missing values. These subjects were included in analyses of the individual properties for which they gave unambiguous responses.

[3]We do not report MSEs here because they are meaningless for repeated factors having more than two levels. In mixed designs such as ours, the standard univariate tests of significance, with their accompanying neatly partitioned sums of squares, are invalid due to violations of the assumption of compound symmetry (see Winer, 1971). We thank Richard Darlington for pointing this out.

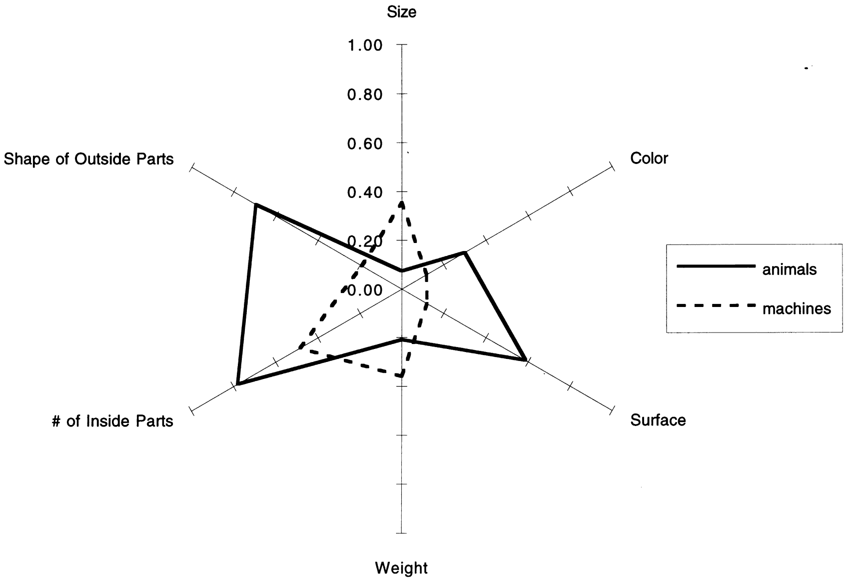

Fig. 1. Mean proportion of children (averaging across ages) who judged that all members of a novel category of animals or machines must share a given property. Each axis of the radial graph represents one of the properties about which children were asked. Distance from the center corresponds to the proportion of children (out of 68 total) who indicated that all members of a particular category had to share the same property (e.g. having black stripes) as the exemplar described in the story. Thus, high values on an axis indicate the relative centrality of the designated property to category membership. The dashed plot line shows the response profile for the novel machine category and the solid plot line shows the response profile for the novel animal category.

Property by Kind $F_{(5,16)} = 8.23$, $P < 0.0006$. In addition, we conducted follow-up analyses to examine the effect of category for each property individually. For each property, we subtracted responses to the machine story from responses to the animal story to get a single difference score for each participant. Six separate one-way ANOVAs compared the three age groups (5, 7, and 9 years) for each of the six properties. No significant main effects were found for any property (color, $F_{(2,65)} = 2.31$, $P > 0.10$; surface markings, $F_{(2,65)} = 1.66$; weight, $F_{(2,65)} < 1$; inside parts, $F_{(2,65)} < 1$; outside parts, $F_{(2,65)} = 1.67$; size, $F_{(2,65)} < 1$). Thus, children of all ages tested were sensitive to differences among the properties they were asked to consider and to the differential relevance of particular attributes to animals vs. machines.

In order to directly assess the significance of the differences in the profiles for animals and machines, we averaged across the age groups and conducted McNemar tests for each property[4]. Here, the McNemar test considers cases in which children responded that changing that property casts doubt on category membership for one kind but not on the other (e.g. all ‘glicks’, a novel animal, have to have stripes, but all ‘nilards,’ a novel machine, do not). Children significantly differentiated animals and

[4]The McNemar test compares responses on two related, dichotomous variables based on a χ^2-distribution. The standard χ^2-test is inappropriate when a single participant would be counted in multiple cells.

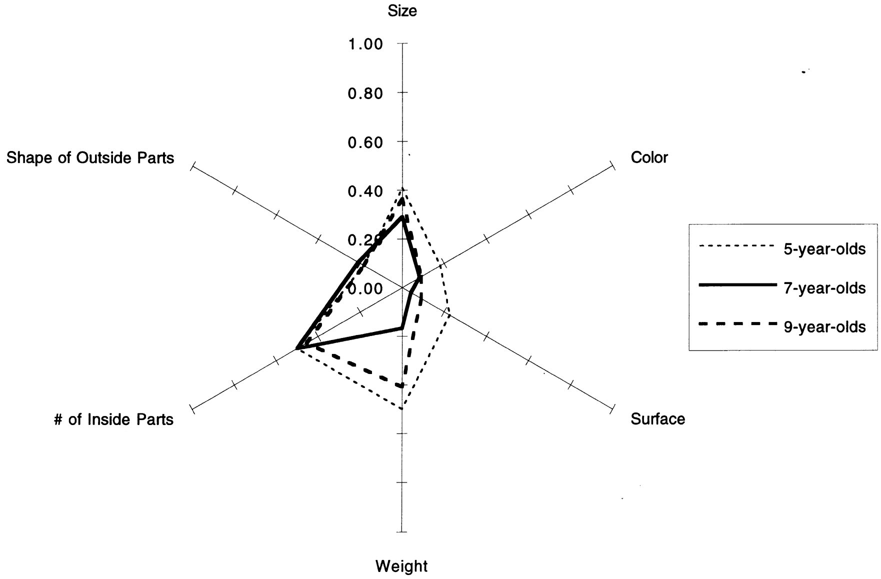

Fig. 2. Proportion of children at each age tested (5-, 7-, and 9-years-old) who judged that all members of a novel category of machines must share a given property. Small dashed line shows response profile of 5-year-olds; solid line shows response profile of 7-year-olds; and large dashed line shows response profile of 9-year-olds. To interpret this radial graph, see detailed description of similar graph in Fig. 1.

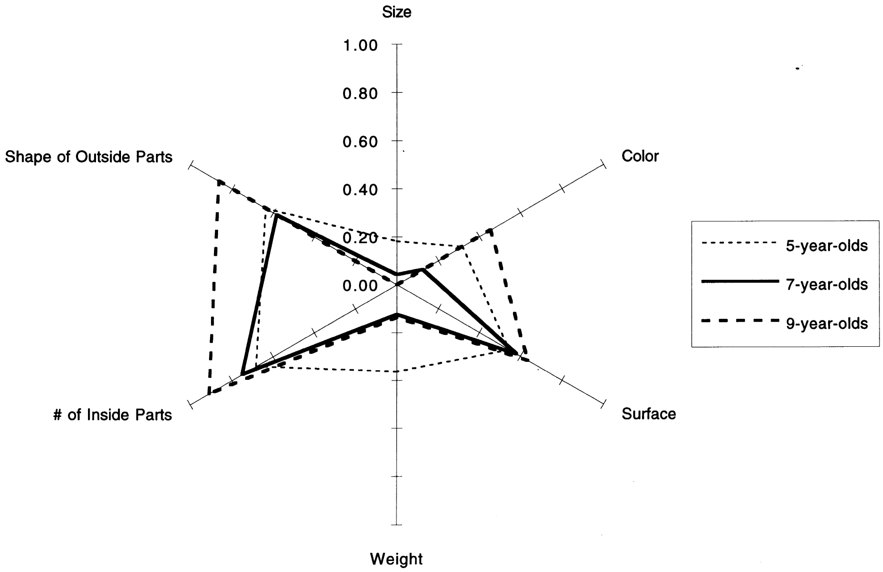

Fig. 3. Proportion of children at each age tested (5-, 7-, and 9-years-old) who judged that all members of a novel category of animals must share a given property. Small dashed line shows response profile of 5-year-olds; solid line shows response profile of 7-year-olds; and large dashed line shows response profile of 9-year-olds. To interpret this radial graph, see detailed description of similar graph in Fig. 1.

Table 3
Number of children (n = 68 total) who judged that a property change was, or was not, relevant to membership in an Animal vs. a Machine novel category

Changed property	Irrelevant to both animal and machine categories	Relevant to both animal and machine categories	Relevant to *animal* but not to machine categories	Relevant to *machine* but not to animal categories	McNemar χ^{2b}
Kind of outside parts[a]	19	10	34	2	30.44***
# Of inside parts	10	32	20	6	7.54**
Color	41	2	18	7	4.84*
Surface markings	26	7	34	1	31.11***
Size	44	3	2	19	13.76**
Weight	39	8	5	16	5.76*

[a] Three participants gave ambiguous responses and are not counted for this property.
[b] Whenever the expected frequency of a McNemar cell is below 5, the McNemar χ^2 tends to be inflated; in these cases, we relied on the exact binomial formula to obtain a P-value. χ^2 degrees of freedom=1.
$^{***}P < 0.0001$; $^{**}P < 0.005$; $^{*}P < 0.05$.

machines for all six properties. Changing the shape of outside parts had a greater effect on animals than machines (i.e. changing the shape of outside parts disrupted category membership more for animals than for machines). Similarly, changing the number of inside parts, the surface markings, and the color had a greater effect on animals than machines. Changing the size and the weight mattered more for machines than for animals. See Table 3 for the frequencies of the four possible response patterns.

9. Discussion

This study revealed that children as young as 5 years recognize the differential importance of specific property types for animals and machines. The majority of children indicated that all members of a novel animal category must share the same color, surface markings, number of inside parts and appearance of outside parts[5]. The same children thought that members of a novel machine category could vary along these attributes yet still belong to the category. In contrast, children indicated that members of an animal category could vary more in size and weight than could members of a machine category.

Interestingly, the pattern of responses shows no systematic developmental changes (as confirmed by the post hoc analyses; see Figs. 2 and 3). This consistency

[5] As long as the outside parts of a particular machine perform the same function as other members of its class, the actual appearance of those parts might vary; but the same is not likely to be true for members of an animal class: if the outside parts look different despite serving the same function for the animal, we would probably think that the animal belonged to a different class. For example, both squirrels and nuthatches dig holes to cache food, but they do not have similarly shaped parts for digging those holes. In contrast, different lamps may have highly dissimilar parts for turning them on and off, yet we would hardly consider this difference significant enough to warrant classifying lamps on the basis of the appearance of their on/off switches.

indicates that children of different ages had the same notions about which types of properties were most central to each category. Across a substantial developmental period notions of centrality for these sets of properties remain remarkably stable. Given the similarity of response profiles across ages 5, 7, and 9, there is no basis to presume that children even younger than those tested here would have radically different knowledge about the fundamental differences in the causal relevance of particular properties to particular kinds of things.

We have suggested that children were responding on the basis of *causal* notions. Children might not, however, have clear notions of the causal reasons for the differential importance of, say, color, to animals versus machines. Instead, such notions might develop from an accumulation of observations about property variation rates within different kinds. If so, the children in this study may have based their responses not so much on an understanding of the causal importance of property types to particular ontological kinds but rather on this accumulation of frequency-based information[6]. The children's protocols suggest otherwise. Although children varied in their verbal precocity, those children who did justify their responses often did so in causal terms. For example, a large proportion of 5-year-olds spontaneously appealed to notions of 'growth' to explain why all 'glicks,' a kind of animal, did not have to be the same size. We did not conduct a systematic protocol analysis (partly because younger children are notoriously bad at articulating explanations), but we suggest that if children were asked to choose between a covariation explanation and a causal explanation (similar to Ahn et al., 1995) for these results, they would overwhelmingly prefer the causal account.

These data may not, then, rule out interpretations that appeal to frequency tabulations, but they do pose a challenge to such accounts. If younger children are insensitive to causal potency and must rely solely on correlational evidence, then when do they start using causal notions—do causal notions click into place once a certain threshold of correlational evidence has accumulated? But then, where do such causal notions come from? Certainly children are not explicitly taught why, for example, surface markings are relevant to what an animal is but not to what a machine is. Furthermore, correlational accounts, in which information accrues from concrete instances in the absence of later-developing abstract notions, would predict that local explanations would develop first only to be replaced later by abstract notions of how the mechanisms of growth, inheritance, and human intention apply broadly to animals or artifacts in general. That is, children should first generalize locally before generalizing observations about properties to all animals as a broad class. But in our study, the animal is a novel one, and the description gives no information that would allow local generalizations—indeed we forced children to consider 'glicks' as a specific type of animal, a very broad level of classification (i.e. as opposed to a specific type of mammal, or pet, or insect). Finally, such accounts would seem to have particular difficulty accounting for children's belief that all animals of the same kind should have the same number of inside parts, given that children very likely have not ever seen, much less counted, the inside parts of various animals.

[6]We thank an anonymous reviewer for pointing this out.

These results also accord with a related task in which two groups of 8-year-olds were familiarized with a single set of features, half of which characterized one category, half of which characterized a contrasting category (Barrett et al., 1993). One group of children was told that the features described two types of animals, while the other group was told that those same features described two types of tools. Subsequently, children's classification of test items (with novel combinations of the familiar features) differed, depending on whether or not the features had been introduced as belonging to two kinds of animals or two kinds of tools. Thus, these older children's knowledge was structured and guided by intuitive notions of the differential importance of certain features to animals vs. tools.

As suggested by the infant research reviewed earlier, these basic notions of causal centrality may emerge early, possibly as early as sensitivity to the typicality of properties. We are not suggesting that infants innately 'know' patterns of causal centrality for different kinds (although even that possibility has not been ruled out). Instead, we are suggesting that the ability to perceive and learn causal patterns may be just as fundamental as the ability to learn typicality and frequency distributions; typicality and causal centrality may go hand in hand in development.

10. Differential knowledge of local causal mechanisms

Knowledge about things such as living kinds and machines does develop. A framework of causal understanding of the differential centrality of properties for different kinds is only that—a framework. Within this framework, a great deal of cognitive growth and conceptual change occurs. Although the results of our study revealed no developmental differences, dramatic developmental changes can occur even for the properties we examined. If we queried details about the causal mechanisms underlying a given property, undoubtedly, we would see developmental changes. However, our questions required only a preliminary understanding of general causal patterns. We did not ask children, for instance, about the inner workings of a gasoline engine, the adaptive value of claws vs. hooves, or how being white helps a polar bear. This distinction between detailed knowledge of local causal mechanisms and less precise frameworks is an important one, as we emphasized earlier. Failure to acknowledge this distinction has sometimes led to mistaken endorsements of Dogma 2 and may foster false controversies regarding developmental change when some research focuses on local mechanism knowledge and other research emphasizes general explanatory frameworks (see Simons and Keil, 1995 for discussion).

We have recently begun to examine more carefully which aspects of causal knowledge emerge early in development to form a relatively fixed explanatory framework and which show change with increasing age. A full account of this distinction will require an extensive set of studies, but the general point can be illustrated through thought experiments. The study reported in this paper demonstrated that, during a developmental period marked by many changes in detailed knowledge of living kinds and artifacts, a general framework of causal knowledge is likely to exist early on, side by side with a sensitivity to typicality information. This somewhat counterintuitive prediction required an initial empirical demonstration. One other, more intuitive pattern will almost certainly be found, although the details of the developmental course can only be revealed through future studies:

For sufficiently complex mechanisms, even most adults will not have a precise causal understanding (e.g. Au, in press). Thus, we might expect there to be relatively stable causal intuitions across lay people and little developmental change in these intuitions, but more dramatic, though highly specific, expert-novice differences. That is, only the most sophisticated experts would have such knowledge, and even for them it would likely be incomplete. An extreme example would be properties that no one outside of the leading laboratories has ever encountered. But even when almost everyone has heard of a property, only a few experts may understand how it is involved in specific mechanisms. For example, color is a salient property for both machines and animals, but only expert chemists might be able to explain the chemical composition of paint pigments or the biological underpinnings of skin and fur pigments. Furthermore, such experts are likely to be the only ones who could explain how the organic and inorganic compounds in animals and machines produce colors differently. Similarly, only certain ecologists would be able to explain in detail how a particular pattern of coloration distinct to a species has adaptive value. For each of the properties used in our study, we could pose similar sorts of questions about mechanisms that only a tiny fraction of adults would know. Moreover, it would be almost trivial to demonstrate that some aspects of knowledge about any particular mechanism take a lifetime of devoted study to acquire.

Aside from such cases of extreme expertise, lay knowledge of both causal mechanisms and typicality often will develop during childhood. Without further careful studies, we cannot easily predict the developmental time course for particular properties, but few would doubt that many such cases exist. For example, children of different ages would likely give quite different answers to questions about the local mechanisms at work in such domains as biology, chemistry, and physics: ‘Is the process that makes things rust like the process that makes people get grey hair, or is it more like the process that makes batteries run down?’ In our preliminary studies we have found that answers to such detailed questions about color, size, and weight may change with age. For example, 5-year-olds do not distinguish between an animal and a machine when asked which is more likely to change its weight during the course of a day, but older children and adults think animals are more likely to show weight fluctuations.

11. Dogma 2: summary

Children do learn as they get older, but the sorts of things they are capable of learning may not undergo qualitative changes such as a shift from similarity to causal frameworks or rules. Young children surely know fewer detailed mechanisms in domains ranging from the germ theory of disease to the actions of levers and fulcrums. They will also know less about typicality distributions of disease types and machine types. They may also acquire some new frameworks of causal understanding. However, our study suggests that such causal frameworks can be stable across large periods of development. Even when children default to typicality to distinguish classes of entities, they do not necessarily undergo a shift from strictly similarity-based reasoning to explanation-based reasoning. The illusion of this shift results from our neglect of their more abstract forms of non-associative knowledge.

12. The unnecessary commitments of empiricism

We have not yet explained why the two dogmas have influenced so much research in cognitive science. One powerful reason may be a tendency within cognitive science to adopt the several unnecessary commitments of the empiricist approach to the acquisition of knowledge. The strongest version of empiricism adopts the view that all knowledge and constraints on learning are acquired through the senses. This extreme view is untenable. Even the British empiricists (e.g. Locke) understood the need for constraints on knowledge acquisition; the sensory organs are clearly tuned to different kinds of information. Rather, they argued that there are no specialized knowledge acquisition devices tuned to different kinds of information. Instead a single general learning capacity accounts for all knowledge acquisition (e.g. Hume's account of learning about causation). Most nativists objected not to the absence of innate beliefs or knowledge but to the lack of domain-specific knowledge acquisition systems. They believed that different systems were biased to acquire different kinds of knowledge, that knowledge is constrained both by the sensory organs and by cognitive biases as well.

If these were the only differences between nativists and empiricists (and in fact they seem to be the only reliable, principled differences), then neither group should be particularly disposed towards the two dogmas. However, many empiricists tend to make additional assumptions that do not logically follow from their fundamental position but which lead to the dogmas. Specifically, they assume that the domain-general learning mechanism relies on association across a stock of perceptual primitives provided by the sensory apparatus. Accordingly, they believe that new associations must build on earlier associations such that higher-order relations can only be represented after all of their constituents are firmly in place. Following Hume, researchers often assume that causal relations can only be appreciated after they are built from pre-existing correlational constituents.

There is, however, no obvious reason why a domain-general learning device must rely on associations among sensory primitives. In fact, the strong form of this empiricist model is clearly false. Infants perceive intermodally at birth and seem to perceive causal relations as early as they have been tested (e.g. Leslie, 1982, 1984; Leslie and Keeble, 1987). In addition, newborns can imitate the facial expressions and body movements of their parent, revealing a sensitivity to particular kinds of complex, relational information from the start. This message, of course, was the key theme of the Gibsonian view of perceptual development. The Gibsons argued that we are all capable, at any age, of picking out complex relational patterns in the environment, especially those that form invariants that could be used to guide action (Gibson, 1966, 1979; Gibson, 1991). Their position was not clearly nativist or empiricist because it did not commit to a domain-general or domain-specific system for detecting invariants. We see no reason why a similar approach cannot be extended to the apprehension of relations that are more cognitive in nature (an extension that was not proposed by the Gibsons). Perhaps a domain-general learning system could be sensitive to and immediately encode such relations as cause, containment, and temporal precedence. If so, we could quickly learn abstract relational patterns that have immediate behavioral consequences. Such a learning system would fit with empiricist principles to the extent that it applied equally well to all sorts of information. There appear to be some innovative and clever attempts to further such a position. Mandler's early image schemas could possibly be interpreted in this way (see Mandler, this issue). For example, containment may be a fundamental image schema accessible to young infants (Baillargeon et al., 1995).

The nativist perspective argues that different aspects of the mind are innately tuned or optimized for picking up different kinds of high-level relational information. Thus, one aspect may be optimized for understanding physical, mechanical causation and another for social causation. The difference between empiricist and nativist accounts is closely linked to the number of distinct learning systems. If there are only two different learning systems that are optimized for two very large domains (e.g. social relations and mechanical forces), the differences between empiricist and nativist accounts are less dramatic. Alternatively, if there are thousands of distinct systems with highly local biases and prejudices, the differences between nativists and empiricists are large but potentially uninteresting as each highly local bias appears trivial and hardly the basis for a system of knowledge or explanatory insight. The middle ground between these extremes, with about a dozen different learning systems, may best capture the heterogeneous causal structure of the world without sacrificing coherency.

Our purpose here is not to determine whether particular kinds of knowledge acquisition problems are best solved by empiricist or nativist approaches, nor is it to argue that only one of these approaches fits with hybrid views of concepts. We do suggest, however, that a certain brand of empiricism seems to lead naturally to the two dogmas. Unfortunately, that brand seems to dominate many current models of how we acquire and use concepts.

Acknowledgements

F.K. was supported by NIH grant R01-HD23922 and D.S. was supported by a Jacob K. Javits fellowship. Many thanks to Bethany Richman for helping to design the study and for collecting all of the data. Thanks also to the children, parents, and teachers at the Ithaca area daycares and schools who made this research possible. We thank three anonymous reviewers for comments that helped to strengthen and clarify our arguments in several areas.

References

Ahn, W., Kalish, C.W., Medin, D.L., Gelman, S.A., 1995. The role of covariation versus mechanism information in causal attribution. Cognition 54, 299–352.

Anderson, J.R., 1976. Language, Memory, and Thought, Earlbaum, Hillsdale, NJ.

Armstrong, S., Gleitman, L., Gleitman, H., 1983. What some concepts might not be. Cognition 13, 263–308.

Au, T., Romo, L., in press. Mechanical causality in children's "folkbiology". In: Medin, D.L., Atran, S. (Eds.), Folkbiology. MIT Press, Cambridge, MA.

Baillargeon, R., Kotovsky, L., and Needham, A., 1995. The acquisition of physical knowledge in infancy. In: Sperber, D., Premack, D., Premack, S.J., Causal Cognition: A multidisciplinary debate, Clarendon Press, Oxford.

Barrett, S.E., Abdi, H., Murphy, G.L., Gallagher, J.M., 1993. Theory-based correlations and their role in children's concepts. Child Development 64, 1595–1616.

Barsalou, L.W., 1993. Challenging assumptions about concepts. Cognitive Development 8, 169–180.

Bassock, M., 1996. Using content to interpret structure: effects on analogical transfer. Current Directions in Psychological Science 5, 54–58.

Biederman, I., Mezzanotte, R.J., Rabinowitz, J.C., 1982. Scene perception: detecting and judging objects undergoing relational violations. Cognitive Psychology 14, 143–177.

Blessing, S.B., Ross, B.H., 1996. Content effects in problem categorization and problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition 22, 792–810.

Boshuizen, H.P.A., Schmidt, H.G., 1992. On the role of biomedical knowledge in clinical reasoning by experts, intermediates and novices. Cognitive Science 16, 153–184.

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

Bruner, J.S., Goodman, C.C., 1947. Value and need as organizing factors in perception. Journal of Abnormal Social Psychology 42, 33–44.

Bruner, J.S., Postman, L., 1947. Emotional selectivity in perception and reaction. Journal of Personality 16, 69–77.

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

Carpenter, G.A., Grossberg, S., 1987. A massively parallel architecture for a self-organizing neural pattern recognition machine. Computer Vision, Graphics, and Image Processing 37, 54–115.

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.

Cheng, P.W., Novick, L.R., 1990. A probabilistic contrast model of causal induction. Journal of Personality and Social Psychology 58, 545–567.

Cheng, P.W., Novick, L.R., 1992. Covariation in natural causal induction. Psychological Review 99, 365–382.

Chi, M.T.H., 1992. Conceptual change within and across ontological categories: Examples from learning and discovery in science. In: Giere, R. (Ed.), Cognitive models of science: Minnesota studies in the philosophy of science, University of Minnesota Press, Minneapolis, MN.

Chi, M.T.H., Feltovich, P.J., Glaser, R., 1981. Categorization and representation of physics problems by experts and novices. Cognitive Science 5, 121–152.

Cook, M., Mineka, S., 1989. Observational conditioning of fear to fear-relevant versus fear-irrelevant stimuli in rhesus monkeys. Journal of Abnormal Psychology 98, 448–459.

Cook, M., Mineka, S., 1990. Selective associations in the observational conditioning of fear in rhesus monkeys. Journal of Experimental Psychology: Animal Behavior Processes 16, 372–389.

Feltovich, P.J., Spiro, R.J., Coulson, R.L., 1989. The nature of conceptual understanding in biomedicine: the deep structure of complex ideas and the development of misconceptions. In: Evans, D.A., Patel, V.L. (Ed.), Cognitive Science in Medicine: Biomedical Modeling, MIT press, Cambridge MA, pp. 113–172.

Fox, R., McDaniel, C., 1982. The perception of biological motion by human infants. Science 218, 486–487.
Garcia, J., Koelling, R.A., 1966. The relation of cue to consequence in avoidance learning. Psychonomic Science 4, 123–124.
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, S.A., Coley, J.D., 1990. The importance of knowing a dodo is a bird: categories and inferences in 2-year-old children. Developmental Psychology 26, 796–804.
Gelman, S.A., 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.
Gelman, S.A., Medin, D.L., 1993. What's so essential about essentialism? A different perspective on the interaction of perception, language, and conceptual knowledge. Cognitive Development 8, 157–168.
Gentner, D., Toupin, C., 1988. Systematicity and surface similarity in the development of analogy. Cognitive Science 10, 277–300.
Gergely, G., Nádasdy, Z., Csibra, G., Bíró, S., 1995. Taking the intentional stance at 12 months of age. Cognition 56, 165–193.
Gibson, E.J., 1991. An odyssey in learning and perception, MIT Press, Cambridge, MA.
Gibson, E.J., Walker, A.S., 1984. Development of knowledge of visual and tactual affordances of substance. Child Development 55, 453–460.
Gibson, J.J., 1966. The senses considered as perceptual systems, Houghton-Mifflin, Boston.
Gibson, J.J., 1979. The ecological approach to visual perception, Houghton-Mifflin, Boston.
Goldstone, R.L., 1994. The role of similarity in categorization: providing a groundwork. Cognition 52, 125–157.
Goodman, N., 1972. Problems and Projects, Bobbs-Merrill, New York.
Harré, R., 1988. Modes of explanation. In: Hilton, D.J. (Ed.), Contemporary science and natural explanation: commonsense conceptions of causality, Harvester Press, Brighton, Sussex.
Harré, R., Madden, E.H., 1975. Causal powers: a theory of natural necessity. Totowa, NJ: Rowman and Littlefield.
Harris, P.L., German, T., Mills, P., 1996. Children's use of counterfactual thinking in causal reasoning. Cognition 61, 233–259.
Inhelder, B., Piaget, J., 1969. The early growth of logic in the child, classification and seriation. New York: W.W. Norton.
Intraub, H , 1980. Rapid conceptual identification of sequentially presented pictures. Journal of Experimental Psychology: Human Perception and Performance 7, 604–610.
Johnson, M.H., 1992. Imprinting and the development of face recognition: From chick to man. Current Directions in Psychological Science 1, 52–55.
Johnson, M.H., Dziurawiec, S., Ellis, H., Morton, J., 1991. Newborns' preferential tracking of face-like stimuli and its subsequent decline. Cognition 40, 1–19.
Jones, S.S., Smith, L.B., 1993. The place of perception in children's concepts. Cognitive Development 8, 113–139.
Keil, F.C., 1981. Constraints on knowledge and cognitive development. Psychological Review 88, 197–227.
Keil, F.C., 1989. Concepts, kinds, and cognitive development, MIT Press, Cambridge, MA.
Keil, F.C., Batterman, N., 1984. A characteristic-to-defining shift in the development of word meaning. Journal of Verbal Learning and Verbal Behavior 23, 221–236.
Keil, F.C., Smith, W.C., 1996. Is there a different 'basic' level for causal relations? Paper presented at the 37th annual meeting of the Psychonomic Society (November), Chicago, IL.
Kelley, H.H., 1967. Attribution theory in social psychology. Nebraska Symposium on Motivation 15, 192–238.
Kuhn, T.S., 1962. The structure of scientific revolutions, University of Chicago Press, Chicago.
Leslie, A.M., 1982. The perception of causality in infants. Perception 11, 173–186.

Leslie, A.M., 1984. Spatiotemporal continuity and the perception of causality in infants. Perception 13, 287–305.

Leslie, A.M., 1995. A theory of agency. In: Sperber, D., Premack, D., Premack, A.J. (Eds.), Causal Cognition: A multidisciplinary debate, Oxford: Clarendon Press, pp. 121– 141.

Leslie, A.M., Keeble, S., 1987. Do 6-month-olds perceive causality? Cognition 25, 265–288.

Mandler, J.M., 1993. On concepts. Cognitive Development 8, 141–148.

Medin, D.L., Shoben, E.J., 1988. Context and structure in conceptual combination. Cognitive Psychology 20, 158–190.

Medin, D.L., Schaffer, M.M., 1978. Context theory in classification learning. Psychological Review 85, 207–238.

Medin, D.L., Goldstone, R.L., Gentner, D., 1993. Respects for similarity. Psychological Review 100, 254–278.

Meltzoff, A.N., 1988. Infant imitation after a 1-week delay: long-term memory for novel acts and multiple stimuli. Developmental Psychology 24, 470–476.

Mervis, C.B., Johnson, K.E., Scott, P., 1993. Perceptual knowledge, conceptual knowledge, and expertise: comment on Jones and Smith. Cognitive Development 8, 149–156.

Murphy, G.L., Allopena, P.D., 1994. The locus of knowledge effects in concept learning. Journal of Experimental Psychology: Learning, Memory, and Cognition 20 (4), 904–919.

Murphy, G.L., Medin, D., 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: Tiberghien, G. (Ed.). Advances in cognitive science: Vol. 2. Theory and applications, Ellis Horwood, Chichester, UK, pp. 23–45.

Neisser, U., 1967. Cognitive Psychology. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R.E., Krantz, D.H., Jepson, C., Kunda, Z., 1983. The use of statistical heuristics in everyday inductive reasoning. Psychological Review 90, 339–363.

Norman, G.R., Schmidt, H.G., 1992. The psychological basis of problem based learning. Academic Medicine 67, 557–565.

Nosofsky, R.M., 1984. Choice, similarity, and the context theory of classification . Journal of Experimental Psychology: Learning, Memory and Cognition 10, 104–114.

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

Nosofsky, R.M., 1992. Similarity scaling and cognitive process models. Annual Review of Psychology 43, 25–53.

Novick, L.R., 1988. Analogical transfer, problem similarity, and expertise. Journal of Experimental Psychology: Learning, Memory and Cognition 14, 510–520.

Patel, V.L., and Groen, G.J., 1991. The general and specific nature of medical expertise: a critical look. In: Ericsson, K.A., Smith, J. (Ed.), Toward a general theory of expertise, Cambridge University Press, New York, pp. 93–125.

Rayner, K., Pollatsek, A., 1992. Eye movements and scene perception. Canadian Journal of Psychology 46, 342–376.

Rosch, E., Mervis, C.B., 1975. Family resemblances: studies in the internal structure of categories. Cognitive Psychology 7, 573–605.

Rozin, P., Kalat, J.W., 1971. Specific hungers and poison avoidance as adaptive specializations of learning. Psychological Review 78, 459–486.

Salmon, W.C., 1989. Four decades of scientific explanation, University of Minnesota Press, Minneapolis, MN.

Simons, D.J., Keil, F.C., 1995. An abstract to concrete shift in the development of biological thought: the insides story. Cognition 56, 129–163.

Sloman, S., 1996. The empirical case for two systems of reasoning. Psychological Bulletin 119, 3–22.

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

Smith, E.E., Shoben, E.J., Rips, L.J., 1974. Structure and process in semantic memory: a featural model for semantic decisions. Psychological Review 81, 214–241.

Smith, L.B., Jones, S.S., 1993. Cognition without concepts. Cognitive Development 8, 181–188.
Spelke, E.S., Breinlinger, K., Macomber, J., Jacobson, K., 1992. Origins of knowledge. Psychological Review 99, 605–632.
Vygotsky, L.S., 1962. Thought and language (E. Hanfmann and G. Vakar, Trans.), MIT Press, Cambridge, MA.
Wellman, H.M., 1990. The child's theory of mind, MIT Press, Cambridge, MA.
Wellman, H.M., Gelman, S.A., 1992. Cognitive development: Foundational theories of core domains. Annual Review of Psychology 43, 337–375.
Werner, H., Kaplan, B., 1963. Symbol formation: An organismic-developmental approach to language and the expression of thought, Wiley, New York.
Williams, J.M.G., Mathews, A., MacLeod, C., 1996. The emotional Stroop task and psychopathology. Psychological Bulletin 120, 3–24.
Winer, B.J., 1971. Statistical principles in experimental design, 2nd ed., McGraw Hill, New York.
Wisniewski, E., Medin, D.L., 1994. On the interaction of theory and data in concept learning. Cognitive Science 18, 221–281.
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.