From Perceptual Categories to Concepts: What Develops?

Vladimir M. Sloutsky

Center for Cognitive Science, The Ohio State University

Abstract

People are remarkably smart: They use language, possess complex motor skills, make nontrivial inferences, develop and use scientific theories, make laws, and adapt to complex dynamic environments. Much of this knowledge requires concepts and this study focuses on how people acquire concepts. It is argued that conceptual development progresses from simple perceptual grouping to highly abstract scientific concepts. This proposal of conceptual development has four parts. First, it is argued that categories in the world have different structure. Second, there might be different learning systems (subserved by different brain mechanisms) that evolved to learn categories of differing structures. Third, these systems exhibit differential maturational course, which affects how categories of different structures are learned in the course of development. And finally, an interaction of these components may result in the developmental transition from perceptual groupings to more abstract concepts. This study reviews a large body of empirical evidence supporting this proposal.

Keywords: Cognitive development; Category learning; Concepts; Conceptual development; Cognitive neuroscience

1. Knowledge acquisition: Categories and concepts

People are remarkably smart: They use language, possess complex motor skills, make nontrivial inferences, develop and use scientific theories, make laws, and adapt to complex dynamic environments. At the same time, they do not exhibit evidence of this knowledge at birth. Therefore, one of the most interesting and exciting challenges in the study of human cognition is to gain an understanding of how people acquire this knowledge in the course of development and learning.

A critical component of knowledge acquisition is the ability to use acquired knowledge across a variety of situations, which requires some form of abstraction or generalization.

Correspondence should be sent to Vladimir M. Sloutsky, Cognitive Development Laboratory, Center for Cognitive Science, The Ohio State University, 208C Ohio Stadium East, 1961 Tuttle Park Place, Columbus, OH 43210. E-mail: sloutsky.1@osu.edu

Examples of abstraction are ample. People can recognize the same object under different viewing conditions. They treat different dogs as members of the same class and expect them to behave in fundamentally similar ways. They learn words uttered by different speakers. Upon learning a hidden property of an item, they extend this property to other similar items. And they apply ways of solving familiar problems to novel problems. In short, people can generalize or form equivalence classes by focusing only on some aspects of information while ignoring others.

This ability to form equivalence classes or categories is present in many nonhuman species (see Zentall, Wasserman, Lazareva, Thompson, & Rattermann, 2008 for a review); however, only humans have the ability to acquire concepts—lexicalized groupings that allow ever-increasing levels of abstraction (e.g., Cat $\rightarrow$ Animal $\rightarrow$ Living thing $\rightarrow$ Object). These lexicalized groupings may include both observable and unobservable properties. For example, although prelinguistic infants can acquire a category ''cat'' by strictly perceptual means (Quinn, Eimas, & Rosenkrantz, 1993), the concept ''cat'' may include many properties that have to be inferred rather than observed directly (e.g., ''mating only with cats, but not with dogs,'' ''being able to move in a self-propelled manner,'' ''having insides of a cat,'' etc.). Often such properties are akin to latent variables—they are inferred from patterns of correlations among observable properties (Rakison & Poulin-Dubois, 2001). These properties can also be lexicalized, and when lexicalized, they allow nontrivial generalizations (e.g., ''plants and animals are alive'' or ''plants and animals reproduce themselves''). Although the existence of prelinguistic concepts is a matter of considerable debate, it seems rather noncontroversial to define those lexicalized properties that have to be inferred (rather than observed) as *conceptual* and lexicalized categories that include such properties as *concepts*.

Concepts are central to human intelligence as they allow uniquely human forms of expression, such as many forms of reasoning. For example, counterfactuals (e.g., ''if the defendant were at home at the time of the crime, she could not have been at the crime scene at the same time'') would be impossible without concepts. According to the present proposal, most concepts develop from perceptual categories and most conceptual properties are inferred from perceptual properties.[1] Therefore, although categories comprise a broader class than concepts (i.e., there are many categories that are not lexicalized and are not based on conceptual properties), there is no fundamental divide between category learning and concept acquisition.

Most of the examples presented in this study deal with ''thing'' concepts (these are lexicalized by ''nominals''), whereas many other concepts, such as actions, properties, quantities, and conceptual combinations are left out. This is because nominals are often most prevalent in the early vocabulary (Gentner, 1982; Nelson, 1973) and entities corresponding to nominals are likely to populate the early experience. Therefore, these concepts appear to be a good starting point in thinking about conceptual development.

The remainder of the study consists of four parts. First, I consider what may develop in the course of conceptual development. Second, I consider some of the critical components of category learning: the structure of input, the multiple competing learning systems, and the asynchronous developmental time course of these systems. Third, I consider evidence

for interactions among these components in category learning and category representation. And, finally, I consider how conceptual development may proceed from perceptual groupings to abstract concepts.

2. The origins of conceptual knowledge

In an attempt to explain developmental origins of conceptual knowledge, a number of theoretical accounts have been proposed. Some argue that input is dramatically underconstrained to enable the acquisition of complex knowledge and some knowledge has to come a priori from the organism, thus constraining future knowledge acquisition. Others suggest that there is much regularity (and thus many constraints) in the environment, with additional constrains stemming from biological specifications of the organism (e.g., limited processing capacity, especially early in development). In the remainder of this section, I review these theoretical approaches.

2.1. Skeletal principles, core knowledge, constraints, and biases

According to this proposal, structured knowledge cannot be recovered from perceptual input because the input is too indeterminate to enable such recovery (Gelman, 1990). This approach is based on an influential idea that was originally proposed for the case of language acquisition but was later generalized to some other aspects of cognitive development, including conceptual development. The original idea is that linguistic input does not have enough information to enable the learner to recover a particular grammar, while ruling out alternatives (Chomsky, 1980). Therefore, some knowledge has to be innate to enable fast, efficient, and invariable learning under the conditions of impoverished input. This argument (known as the Poverty of the Stimulus argument) has been subsequently generalized to perceptual, lexical, and conceptual development. If input is too impoverished to constrain possible inductions and to license the concepts that we have, the constraints have to come from somewhere. It has been proposed that these constraints are internal—they come from the organism, and they are a priori and top-down (i.e., they do not come from data). A variety of such constraints have been proposed, including, but not limited to, innate knowledge within ''core'' domains (Carey, 2009; Carey & Spelke, 1994, 1996; Spelke, 2000; Spelke & Kinzler, 2007), skeletal principles (e.g., Gelman, 1990), ontological knowledge (Keil, 1979; Mandler, Bauer, & McDonough, 1991; Pinker, 1984; Soja, Carey, & Spelke, 1991), conceptual assumptions (Gelman, 1988; Gelman & Coley, 1991; Markman, 1989), and word-learning biases (Markman, 1989; see also Golinkoff, Mervis, & Hirsh-Pasek, 1994).

However, there are several lines of evidence challenging (a) the explanatory machinery of this account with respect to language (Chater & Christiansen, this issue) and (b) the existence of particular core abilities (e.g., Twyman & Newcome, this issue). Furthermore, although the Poverty of the Stimulus argument is formally valid, its premises and therefore its conclusions are questionable. Most importantly, very little is known about the information value of input with respect to knowledge in question. Therefore, it is not clear whether

input is in fact as impoverished as it has been claimed. In addition, there are several lines of evidence suggesting that input might be richer than it is expected under the Poverty of the Stimulus assumption.

First, the fact that infants, great primates, monkeys, rats, and birds all can learn a variety of basic-level perceptual categories (Cook & Smith, 2006; Quinn et al., 1993; Smith, Redford, & Haas, 2008; Zentall et al., 2008) strongly indicates that perceptual input (at least for basic-level categories) is not impoverished. Otherwise, one would need to assume that all these species have the same constraints as humans, which seems implausible given vastly different environments in which these species live.

In addition, there is evidence that perceptual input (Rakison & Poulin-Dubois, 2001) or a combination of perceptual and linguistic input (Jones & Smith, 2002; Samuelson & Smith, 1999; Yoshida & Smith, 2003) can jointly guide the acquisition of broad ontological classes. Furthermore, cross-linguistic evidence suggests that ontological boundaries exhibit greater cross-linguistic variability than could be expected if they were fixed (Imai & Gentner, 1997; Yoshida & Smith, 2003). Therefore, there might be enough information in the input for the learner to form both basic-level categories and broader ontological classes. There is also modeling work (e.g., Gureckis & Love, 2004; Rogers & McClelland, 2004) offering a mechanistic account of how coherent covariation in the input could guide the acquisition of broad ontological classes as well as more specific categories.

In short, there are reasons to doubt that input is in fact impoverished, and if it is not impoverished, then a priori assumptions are not necessary. Therefore, to understand conceptual development, it seems reasonable to shift the focus away from a priori constraints and biases and toward the input and the way it is processed.

2.2. Similarity, correlations, and attentional weights

According to an alternative approach, conceptual knowledge as well as some of the biases and assumptions are a product rather than a precondition of learning (see Rogers & McClelland, 2004, for a connectionist implementation of these ideas). Early in development, cognitive processes are grounded in powerful learning mechanisms, such as statistical and attentional learning (French, Mareschal, Mermillod, & Quinn, 2004; Mareschal, Quinn, & French, 2002; Rogers & McClelland, 2004; Saffran, Johnson, Aslin, & Newport, 1999; Sloutsky, 2003; Sloutsky & Fisher, 2004a; Smith, 1989; Smith, Jones, & Landau, 1996).

According to this view, input is highly regular and the goal of learning is to extract these regularities. For example, category learning could be achieved by detecting multiple commonalities, or similarities, among presented entities. In addition, not all commonalities are the same—features may differ in salience and usefulness for generalization, with both salience and usefulness of a feature reflected in its attentional weight. However, unlike the a priori assumptions, attentional weights are not fixed and they can change as a result of learning: Attentional weights of more useful features increase, whereas these weights decrease for less useful features (Kruschke, 1992; Nosofsky, 1986; Opfer & Siegler, 2004; Sloutsky & Spino, 2004; see also Hammer & Diesendruck, 2005).

There are several lines of research presenting evidence that both basic-level categories (e.g., dogs) and broader ontological classes (e.g., animates vs. inanimates) have multiple perceptual within-category commonalities and between-category differences (French et al., 2004; Rakison & Poulin-Dubois, 2001; Samuelson & Smith, 1999). Some researchers argue that additional statistical constraints come from language in the form of syntactic cues, such as count noun and mass noun syntax (Samuelson & Smith, 1999). Furthermore, cross-linguistic differences in the syntactic cues (e.g., between English and Japanese) can push ontological boundaries in speakers of respective languages (Imai & Gentner, 1997; Yoshida & Smith, 2003). Finally, different categories could be organized differently (e.g., living things could be organized by multiple similarities, whereas artifacts could be organized by shape), and there might be multiple correlations between category structure, perceptual cues, and linguistic cues. All this information could be used to distinguish between different kinds. As children acquire language, they may become sensitive to these correlations, which may affect their attention to shape in the context of artifacts versus living things (Jones & Smith, 2002).

This general approach may offer an account of conceptual development that does not posit a priori knowledge structures. It assumes that input is sufficiently rich to enable the young learner to form perceptual groupings. Language provides learners with an additional set of cues that allow them to form more abstract distinctions. Finally, lexicalization of such groupings as well as of some unobservable conceptual features could result in the acquisition of concepts at progressively increasing levels of abstraction. In the next section, I will outline how conceptual development could proceed from perceptual groupings to abstract concepts.

2.3. From percepts to concepts: What develops?

If people start out with perceptual groupings, how do they end up with sophisticated conceptual knowledge? According to the proposal presented here, conceptual development hinges on several critical steps. These include the ability to learn similarity-based unimodal categories, the ability to integrate cross-modal information, the lexicalization of learned perceptual groupings, the lexicalization of conceptual features, and the development of executive function. The latter development is of critical importance for acquiring abstract concepts that are not grounded in similarity. Examples of such concepts are unobservables (e.g., *love*, *doubt*, *thought*), relational concepts (e.g., *enemy* or *barrier*), as well as a variety of rule-based categories (e.g., *island*, *uncle*, or *acceleration*). Because these concepts require focusing on unobservable abstract features, their acquisition may depend on the maturity of executive function.

This developmental time course is determined in part by an interaction of several critical components. These components include the following: (a) the structure of the to-be-learned category, (b) the competing learning systems that might subserve learning categories of different structures, and (c) developmental course of these learning systems. First, categories differ in their structure. For example, some categories (e.g., most of natural kinds, such as *cat* or *dog*) have multiple intercorrelated features relevant for category membership. These

features are jointly predictive, thus yielding a highly redundant (or statistically dense) category. These categories often have graded membership (i.e., a typical dog is a better member of the category than an atypical dog) and fuzzy boundaries (i.e., it is not clear whether a cross between a dog and a cat is a dog). At the same time, other categories are defined by a single dimension or a relation between or among dimensions. Members of these categories have very few common features, with the rest of the features varying independently and thus contributing to irrelevant or ''surface'' variance. Good examples of such sparse categories are mathematical and scientific concepts. Consider the two situations: (a) increase in a population of fish in a pond and (b) interest accumulation in a bank account. Only a single commonality—exponential growth—makes both events instances of the same mathematical function. All other features are irrelevant for membership in this category and can vary greatly.

Second, there might be multiple systems of category learning (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998) evolved to learn categories of different structures. In particular, a compression-based system may subserve category learning by reducing perceptually rich input to a more basic format. As a result of this compression, features that are common to category members (but not to nonmembers) become a part of representation, whereas idiosyncratic features get washed out. In contrast, the selection-based system may subserve category learning by shifting attention to category-relevant dimension(s) and away from irrelevant dimension(s). Such selectivity may require the involvement of brain structures associated with executive function. The compression-based system could have an advantage for learning dense categories, which could be acquired by mostly perceptual means. At the same time, the selection-based system could have an advantage for learning sparse categories, which require focusing on few category-relevant features (Kloos & Sloutsky, 2008; see also Blair, Watson, & Meier, 2009, for a discussion).

The involvement of each system may also affect what information is encoded in the course of category learning, and, subsequently, how a learned category is represented. In particular, the involvement of the compression-based system may result in a reduced yet fundamentally perceptual representation of a category, whereas the involvement of the selection-based system may result in a more abstract (e.g., lexicalized) representation. Given that many real-life categories (e.g., dogs, cats, or cups) are acquired by perceptual means and later undergo lexicalization, there are reasons to believe that these categories combine perceptual representation with a more abstract lexicalized representation. These abstract lexicalized representations are critically important for the ability to reason and form arguments that could be all but impossible to form by strictly perceptual means. For example, it is not clear how purely perceptual representation of constituent entities would support a counterfactual of the form ''If my grandmother were my grandfather…''

And third, the category-learning systems and associated brain structures may come online at different points in development, with the system subserving learning of dense categories coming online earlier than the system subserving learning of sparse categories. In particular, there is evidence that many components of executive function critical for learning sparse categories exhibit late developmental onset (e.g., Davidson, Amso, Anderson, & Diamond, 2006). If this is the case, then able learning and representation of dense categories should

precede that of sparse categories. Under this view, ''conceptual'' assumptions do not have to underlie category learning, as most categories that children acquire are spontaneously dense and can be acquired implicitly, without a teaching signal or supervision. At the same time, some of these ''conceptual'' assumptions could be a product of later development.

The current proposal of conceptual development has three parts (see Sections 3–5). In Section 3, I consider in detail components of category learning: category structure, the multiple competing learning systems, and the potentially different maturational course of these systems. I suggest that categories in the world differ in their structure and consider ways of quantifying this structure. I then present another argument that there might be different learning systems (subserved by different brain mechanisms) that evolved to learn categories of differing structures. Finally, I argue that these systems exhibit differential maturational course, which affects how categories of different structures are learned in the course of development. Then, in Section 4, I consider an interaction of these components. This interaction is important because it may result in the developmental transition from perceptual groupings to abstract concepts. These arguments point to a more nuanced developmental picture (presented in Section 5), in which learning of perceptual categories, cross-modal integration, lexicalization, learning of conceptual properties, the ability to focus and shift attention, and the development of lexicalized concepts are logical steps in conceptual development.

3. Components of category learning: Input, learning system, and the learner

3.1. Characteristics of input: Category structure

It appears almost self-evident that categories differ in their structure. Some categories are coherent: Their members have multiple overlapping features and are often similar (e.g., cats or dogs are good examples of such categories). Other categories seem to be less coherent: Their members have few overlapping features (e.g., square things). These differences have been noted by a number of researchers who pointed to different category structures between different levels of ontology (e.g., Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) and between animal and artifact categories (Jones & Smith, 2002; Jones, Smith, & Landau, 1991; Markman, 1989). Category structure can be captured formally and one such treatment of category structure has been offered recently (Kloos & Sloutsky, 2008). The focal idea of this proposal is that category structure can be measured by statistical density of a category. Statistical density is a function of within-category compactness and between-category distinctiveness, and it may have profound effects on category learning. In what follows, I will elaborate this idea.

3.1.1. Statistical density as a measure of category structure

Any set of items can have a number of possible dimensions (e.g., color, shape, size), some of which might vary and some of which might not. Categories that are statistically dense have multiple intercorrelated (or covarying) features relevant for category membership, with only a few features being irrelevant. Good examples of statistically dense

categories are basic-level animal categories such as cat or dog. Category members have particular distributions of values on a number of dimensions (e.g., shape, size, color, texture, number of parts, type of locomotion, type of sounds they produce, etc.). These distributions are jointly predictive, thus yielding a dense (albeit probabilistic) category. Categories that are statistically sparse have very few relevant features, with the rest of the features varying independently. Good examples of sparse categories are dimensional groupings (e.g., ''round things''), relational concepts (e.g., ''more''), scientific concepts (e.g., ''accelerated motion''), or role-governed concepts (e.g., cardinal number; see Markman & Stilwell, 2001, for a discussion of role-governed categories).

Conceptually, statistical density is a ratio of variance relevant for category membership to the total variance across members and nonmembers of the category. Therefore, density is a measure of statistical redundancy (Shannon & Weaver, 1948), which is an inverse function of relative entropy.

Density can be expressed as

$$D = 1 - \frac{H_{\text{within}}}{H_{\text{between}}}, \tag{1}$$

where H_{within} is the entropy observed within the target category, and H_{between} is the entropy observed between target and contrasting categories.

A detailed treatment of statistical density and ways of calculating it is presented elsewhere (Kloos & Sloutsky, 2008); thus, only a brief overview of statistical density is presented below. Three aspects of stimuli are important for calculating statistical density: variation in stimulus dimensions, variation in relations among dimensions, and attentional weights of stimulus dimensions.

First, a stimulus dimension may vary either within a category (e.g., members of a target category are either black or white) or between categories (e.g., all members of a target category are black, whereas all members of a contrasting category are white). Within-category variance decreases the density, whereas between-category variance increases the density.

Second, dimensions of variation may be related (e.g., all items are black circles), or they may vary independently of each other (e.g., items can be black circles, black squares, white circles, or white squares). Covarying dimensions result in smaller entropy than dimensions that vary independently. It is not unreasonable to assume that only dyadic relations (i.e., relations between two dimensions) are detected spontaneously, whereas relations of higher arity (e.g., a relation among color, shape, and size) are not (Whitman & Garner, 1962). Therefore, only dyadic relations are included in the calculation of entropy.

The total entropy is the sum of the entropy due to varying dimensions (H^{dim}), and the entropy due to varying relations among the dimensions (H^{rel}). More specifically,

$$H_{\text{within}} = H_{\text{within}}^{\text{dim}} + H_{\text{within}}^{rel}, \text{and} \tag{2a}$$

$$H_{\text{between}} = H_{\text{between}}^{\text{dim}} + H_{\text{between}}^{\text{rel}}. \tag{2b}$$

The concept of entropy was formalized by the information theory (Shannon & Weaver, 1948), and we use these formalisms here. First consider the entropy due to dimensions. This within- and between-category entropy is presented in Eqs. 3a and 3b, respectively.

$$H_{\text{within}}^{\text{dim}} = -\sum_{i=1}^{M} w_i \left[\sum_{j=0,1} {}_{\text{within}}(p_j \log_2 p_j) \right] \tag{3a}$$

$$H_{\text{between}}^{\text{dim}} = -\sum_{i=1}^{M} w_i \left[\sum_{j=0,1} {}_{\text{between}}(p_j \log_2 p_j) \right] \tag{3b}$$

where M is the total number of varying dimensions, w_i is the attentional weight of a particular dimension (the sum of attentional weights equals to a constant), and p_j is the probability of value j on dimension i (e.g., the probability of a color being white). The probabilities could be calculated within a category or between categories.

The attentional-weight parameter is of critical importance—without this parameter, it would be impossible to account for learning of sparse categories. In particular, when a category is dense, even relatively small attentional weights of individual dimensions add up across many dimensions. This makes it possible to learn the category without supervision. Conversely, when a category is sparse, only few dimensions are relevant. If attentional weights of each dimension are too small, supervision could be needed to direct attention to these relevant dimensions.

Next, consider entropy that is due to a relation between dimensions. To express this entropy, we need to consider the co-occurrences of dimensional values. If dimensions are binary, with each value coded as 0 or 1 (e.g., white = 0, black = 1, circle = 0, and square = 1), then the following four co-occurrence outcomes are possible: 00 (i.e., white circle), 01 (i.e., white square), 10 (i.e., black circle), and 11 (i.e., black square). The within- and between-category entropy that is due to relations is presented in Eqs. 4a and 4b, respectively.

$$H_{\text{within}}^{\text{rel}} = -\sum_{k=1}^{o} w_k \left[\sum_{\substack{m=0,1 \\ n=0,1}} {}_{\text{within}}(p_{mn} \log_2 p_{mn}) \right], \tag{4a}$$

$$H_{\text{between}}^{\text{rel}} = -\sum_{k=1}^{o} w_k \left[\sum_{\substack{m=0,1 \\ n=0,1}} {}_{\text{between}}(p_{mn} \log_2 p_{mn}) \right], \tag{4b}$$

where o is the total number of possible dyadic relations among the varying dimensions, w_k is the attentional weight of a particular relation (again, the sum of attentional weights equals to a constant), and p_{mn} is the probability of a co-occurrence of values m and n on a binary relation k (which conjoins two dimensions of variation).

3.1.2. Density, salience, and similarity

The concept of density is closely related to the ideas of salience and similarity, and it is necessary to clarify these relations. First, density is a function of *weighted* entropy, with attentional weight corresponding closely to the salience of a feature. Therefore, feature salience can affect density by affecting the attentional weight of the feature in question. Of course, as mentioned earlier, attentional weights are not fixed and they can change as a result of learning. Second, perceptual similarity is a sufficient, but not necessary condition of density—all categories bound by similarity are dense, but not all dense categories are bound by similarity. For example, some categories could have multiple overlapping relations rather than overlapping features (e.g., members of a category have short legs and short neck or long legs and long neck). It is conceivable that such nonlinearly separable (NLS) categories could be relatively dense, yet not bound by similarity.

3.1.3. Category structure and early learning

Although it is difficult to precisely calculate the density of categories surrounding young infants, some estimates can be made. It seems that many of these categories, while exhibiting within-category variability in color (and sometime in size), have similar within-category shape, material, and texture (*ball*, *cup*, *bottle*, *shoe*, *book*, or *apple* are good examples of such categories); these categories should be relatively dense. As I show next, dense categories can be learned implicitly, without supervision. Therefore, it is possible that prelinguistic infants implicitly learn many of the categories surrounding them. Incidentally, the very first noun words that infants learn denote these dense categories (see Dale & Fenson, 1996; Nelson, 1973). Therefore, it is possible that some of early word learning consists of learning lexical entries for already known dense categories. This possibility, however, is yet to be tested empirically.

3.1.3.1. Characteristics of the learning system: Multiple competing systems of category learning: The role of category structure in category learning has been a focus of the neuroscience of category learning. Recent advances in that field suggest that there might be multiple systems of category learning (e.g., Ashby et al., 1998; Cincotta & Seger, 2007; Nomura & Reber, 2008; Seger, 2008; Seger & Cincotta, 2002) and an analysis of these systems may elucidate how category structure interacts with category learning. I consider these systems in this section.

There is an emerging body of research on brain mechanisms underlying category learning (see Ashby & Maddox, 2005; Seger, 2008, for reviews). Although the anatomical localization and the involvement of specific circuits remain a matter of considerable debate, there is substantial agreement that ''wholistic'' or ''similarity-based'' categories (which are typically dense) and ''dimensional'' or ''rule-based'' categories (which are typically sparse) could be learned by different systems in the brain.

There are several specific proposals identifying brain structures that comprise each system of category learning (Ashby et al., 1998; Cincotta & Seger, 2007; Nomura & Reber, 2008; Seger, 2008; Seger & Cincotta, 2002). Most of the proposals involve three major hierarchical structures: cortex, basal ganglia, and thalamus. There is also evidence for the

involvement of the medial temporal lobe (MTL) in category learning (e.g., Nomura et al., 2007; see also Love & Gureckis, 2007). However, because the maturational time course of the MTL is not well understood (Alvarado & Bachevalier, 2000), I will not focus here on this area of the brain.

One influential proposal (e.g., Ashby et al., 1998) posited two cortical–striatal–pallidal–thalamic–cortical loops, which define two acting in parallel circuits. The circuit responsible for learning of similarity-based categories originates in extrastriate visual areas of the cortex (such as inferotemporal [IT] cortex) and includes the posterior body and tail of the caudate nucleus. In contrast, the circuit responsible for the learning of rule-based categories originates in the prefrontal and anterior-cingulated cortices (ACCs) and includes the head of the caudate (Lombardi et al., 1999; Rao et al., 1997; Rogers, Andrews, Grasby, Brooks, & Robbins, 2000).

In a similar vein, Seger and Cincotta (2002) proposed the *visual* loop, which originates in the inferior temporal areas and passes through the tail of the caudate nucleus in the striatum, and the *cognitive* loop, which passes through the prefrontal cortex (PFC) and the head of the caudate nucleus. The visual loop has been shown to be involved in visual pattern discrimination in nonhuman animals (Buffalo et al., 1999; Fernandez-Ruiz, Wang, Aigner, & Mishkin, 2001; Teng, Stefanacci, Squire, & Zola, 2000), and Seger and Cincotta (2002) have proposed that this loop may subserve learning of similarity-based visual categories. The cognitive loop has been shown to be involved in learning of rule-based categories (e.g., Rao et al., 1997; Seger & Cincotta, 2002; see also Seger, 2008).

There is also evidence that category learning is achieved differently in the two systems. The critical feature of the visual system is the reduction of information or *compression*, with only some but not all stimulus features being encoded. Therefore, I will refer to this system as the *compression-based* system of category learning. A schematic representation of processing in this system is depicted in Fig. 1A. The feature map in the top layer gets compressed in the bottom layer, with only some features of the top layer represented in the bottom layer.

This compression is achieved by many-to-one projections of the visual cortical neurons in the IT cortex onto the neurons of the tail of the caudate (Bar-Gad, Morris, & Bergman, 2003; Wilson, 1995). In other words, many cortical neurons converge on an individual caudate neuron. As a result of this convergence, information is compressed to a more basic form, with redundant and highly probable features likely to be encoded (and thus learned) and idiosyncratic and rare features likely to be filtered out.

Category learning in this system results in a reduced (or compressed) yet fundamentally perceptual representation of stimuli. If every stimulus is compressed, then those features and feature relations that are frequent in category members should survive the compression, whereas rare or unique features/relations should not. Because compression does not require selectivity, compression-based learning could be achieved implicitly, without supervision (such as feedback or even more explicit forms of training), and it should be particularly successful in the learning of dense categories.

In short, there is a critical feature of the compression-based system—it can learn dense categories without supervision. Under some conditions, the compression-based system may
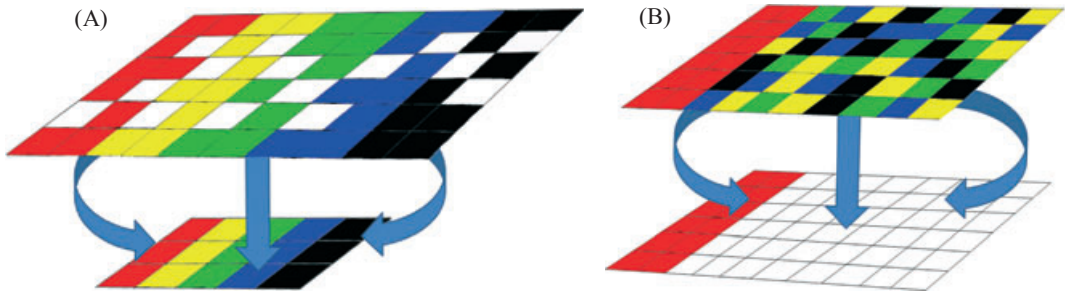
Fig. 1. (A) Schematic depiction of the compression-based system. The top layer represents stimulus encoding in inferotemporal cortex. This rich encoding gets compressed to a more basic form in the striatum represented by the bottom layer. Although some of the features are left out, much perceptual information present in the top layer is retained in the bottom layer. (B) Schematic depiction of the selection-based system. The top layer represents selective encoding in the prefrontal cortex. The selected dimension is then projected to the striatum represented by the bottom layer. Only the selected information is retained in the bottom layer.

also learn structures defined by a single dimension of variation (e.g., color or shape). For example, when there is a small number of dimensions of variation (e.g., color and shape, with shape distinguishing among categories), compression may be sufficient for learning a category-relevant dimension. However, if categories are sparse, with only few relevant dimensions and multiple irrelevant dimensions, learning of the relevant dimensions by compression could be exceedingly long or not possible at all.

The critical aspect of the second system of category learning is the cognitive loop, which involves (in addition to the striatum) the dorsolateral PFC and the ACC—the cortical areas that subserve attentional selectivity, working memory, and other aspects of executive function (Posner & Petersen, 1990). I will therefore refer to this system as *selection-based*. The selection-based system enables attentional learning—allocation of attention to some stimulus dimensions and ignoring others (e.g., Kruschke, 1992, 2001; Mackintosh, 1975; Nosofsky, 1986). Unlike the compression-based system where learning is driven by reduction and filtering of idiosyncratic features (while retaining features and feature correlations that recur across instances), learning in the selection-based system could be driven by error reduction. As schematically depicted in Fig. 1B, attention is shifted to those dimensions that predict error reduction and away from those that do not (e.g., Kruschke, 2001; but see Blair et al., 2009).

Given that attention has to be shifted to a relevant dimension, the task of category learning within the selection-based system should be easier when there are fewer relevant dimensions (see Kruschke, 1993, 2001, for related arguments). This is because it is easier to shift attention to a single dimension than to allocate it to multiple dimensions. Therefore, the selection-based system is better suited to learn sparse categories (recall that the compression-based system is better suited to learn dense categories). For example, Kruschke (1993) describes an experiment where participants learned a category in a supervised manner, with feedback presented on every trial. For some categories, a single dimension was relevant, whereas for other categories, two related dimensions were relevant for categorization.

Participants were shown to learn better in the former than in the latter condition. Given that learning was supervised (i.e., category learning and stimulus dimensions that might be relevant for categorization were mentioned explicitly, and feedback was given on every trial), it is likely that the selection-based system was engaged.

The selection-based system depends critically on prefrontal circuits because these circuits enable the selection of a relevant stimulus dimension (or rule), while inhibiting irrelevant dimensions. The selected (and perhaps amplified) dimension is likely to survive the compression in the striatum, whereas the nonselected (and perhaps weakened) dimensions may not. Therefore, there is little surprise that young children (whose selection-based system is still immature) tend to exhibit successful categorization performance when categories are based on multiple dimensions than when they are based on a single dimension (e.g., Smith, 1989).

How are the systems deployed? Although the precise mechanism remains unknown, several ideas have been proposed. For example, Ashby et al. (1998) posited competition between the systems, with the selection-based system being deployed by default. This idea stems from evidence that participants exhibited more able learning when categories were based on a single dimension than when categories are based on multiple dimensions (e.g., Ashby et al., 1998; Kruschke, 1993). However, it is possible that the selection-based system was triggered by feedback and explicit learning regime, whereas in the absence of supervision the compression-based system is a default (Kloos & Sloutsky, 2008). Furthermore, it seems unlikely that the idea of the default deployment of the selection-based system describes accurately what happens early in development. As I argue in the next section, because some critical cortical components of the selection-based system mature relatively late, it is likely that early in development the competition is weakened (or even absent), thus making the compression-based system default.

If the compression-based system is deployed by default early in development (and, when supervision is absent, it is deployed by default in adults as well), this default deployment may have consequences for category learning. In particular, if a category is sparse, the compression-based system may fail to learn it due to a low signal-to-noise ratio in the sparse category. In contrast, the selection-based system may have the ability to increase the signal-to-noise ratio by shifting attention to the signal, thus either amplifying the signal or by inhibiting noise.

The idea of multiple systems of category learning has been supported by both fMRI and neuropsychological evidence. In one neuroimaging study reported by Nomura et al. (2007), participants were scanned while learning two categories of sine wave gratings. The gratings varied on two dimensions: spatial frequency and orientation of the lines. In the rule-based condition, category membership was defined only by the spatial frequency of the lines (see Fig. 2A), whereas in the ''wholistic'' condition, both frequency and orientation determined category membership (see Fig. 2B). Note that each point in Fig. 2 represents an item and the colors represent distinct categories. Rule-based categorization showed greater differential activation in the hippocampus, the ACC, and medial frontal gyrus. At the same time, the wholistic categorization exhibited greater differential activation in the head and tail of the caudate.
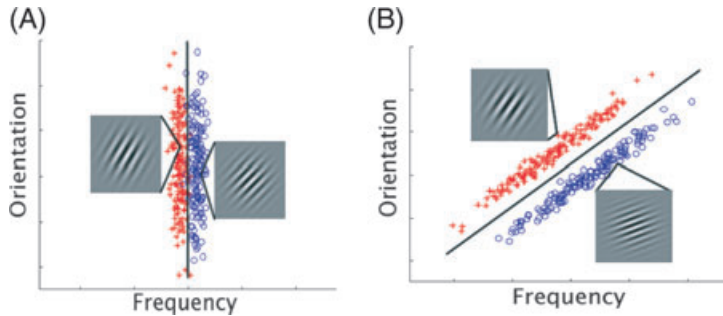
Fig. 2. RB (A) and II stimuli (B) (after Nomura & Reber, 2008). Each point represents a distinct Gabor patch (sine-wave) stimulus defined by orientation (tilt) and frequency (thickness of lines). In both stimulus sets, there are two categories (red and blue points). RB categories are defined by a vertical boundary (only frequency is relevant for categorization), whereas II categories are defined by a diagonal boundary (both orientation and frequency are relevant). In both RB and II stimuli, there are examples of a stimulus from each category. II, information integration; RB, rule based.

Some evidence for the possibility of the two systems of category learning stem from neuropsychological research. One of the most frequently studied populations are patients with Parkinson's disease (PD), because the disease often affects frontal cortical areas in addition to striatal areas (e.g., van Domburg & ten Donkelaar, 1991). As a result, these patients often exhibit impairments in both the compression-based and the selection-based systems of category learning. Therefore, this group provides only indirect rather than clear-cut evidence for the dissociation between the systems. For example, impairments of the compression-based system in PD were demonstrated in a study by Knowlton, Mangels, and Squire (1996), in which patients with PD (which affects the release of dopamine in the striatum) had difficulty in learning probabilistic categories that were determined by co-occurrence of multiple perceptual cues. Impairments of the selection-based learning system have been demonstrated in patients with damage to the PFC (which also often include PD patients). Specifically, in the multiple studies using the Wisconsin Card Sorting Test (WCST: Berg, 1948; Brown & Marsden, 1988; Cools, van den Bercken, Horstink, van Spaendonck, & Berger, 1984), it was found that the patients often exhibit impaired learning of categories based on verbal rules, as well as impairments in shifting attention from successfully learned rules to new rules (see Ashby et al., 1998, for a review).

In the WCST, participants have to discover an experimenter-defined matching rule (e.g., ''objects with the same shape go together'') and respond according to the rule. In the middle of the task, the rule may change and participants must sort according to the new rule. Two aspects of the task are of interest, rule learning and response shifting, with both being likely to be subserved by the selection-based system (see Ashby et al., 1998, for a discussion). There are several types of shifts, with two being of particular interest for understanding of the selection-based system—the reversal shift and the extradimensional shift.

The reversal shift consists of a reassignment of a dimension to a response. For example, a participant could initially learn that ''if Category A (say the color is green), then press

button 1, and if Category B (say the color is red), then press button 2.'' The reversal shift requires a participant to change the pattern of responding, such that ''if Category A, then press button 2, and if Category B, then press button 1.'' In contrast, the extradimensional shift consists of change in which dimension is relevant. For example, if a participant initially learned that ''if Category A (say the color is green), then press button 1, and if Category B (say the color is red), then press button 2,'' the extradimensional shift would require a different pattern of responding: ''if Category K (say the size is small), then press button 1, and if Category M (say the size is large), then press button 2.'' Findings indicate that patients with lesions to PFCs had substantial difficulties with extradimensional, but not with the reversal shifts on the WCST (e.g., Rogers et al., 2000). Therefore, these patients did not have a difficulty in inhibiting the previously learned pattern of responding but rather had difficulty in shifting attention to a formerly irrelevant dimension, which is indicative of a selection-based system impairment.

In sum, there is evidence that the compression-based and the selection-based system may be dissociated in the brain. Furthermore, although both systems involve parts of the striatum, they differ with respect to other areas of the brain. Whereas the selection-based system relies critically on the PFC and the ACC, the compression-based system relies on IT cortex. As I argue in the next section, the IT and the PFCs may exhibit differential maturational time course. The relative immaturity of PFCs early in development coupled with a relative maturity of the IT cortex and the striatum should result in young children having a more mature compression-based than selection-based system and thus being more efficient in learning dense than sparse categories (Smith, 1989; Smith & Kemier-Nelson, 1984).

3.2. Characteristics of the learner: Differential maturational course of brain systems underlying category learning

Many vertebrates have a brain structure analogous to the IT cortex and the striatum, whereas only mammals have a developed PFC (Striedter, 2005). Studies of normal brain maturation (Jernigan, Trauner, Hesselink, & Tallal, 1991; Pfefferbaum et al., 1994; Caviness, Kennedy, Richelme, Rademacher, & Filipek, 1996; Giedd et al., 1996a, 1996b; Sowell & Jernigan, 1999; Sowell, Thompson, Holmes, Batth, Jernigan, and Toga, 1999, Sowell, Thompson, Holmes, Jernigan, and Toga, 1999) have indicated that brain morphology continues to change well into adulthood. As noted by Sowell, Thompson, Holmes, Batth, et al. (1999), maturation progresses in a programmed way, with phylogenetically more primitive regions of the brain (e.g., brain stem and cerebellum) maturing earlier, and more advanced regions of the brain (e.g., the association circuits of the frontal lobes) maturing later. In addition to the study of brain development focused on the anatomy, physiology, and chemistry of the changing brain, researchers have studied the development of function that is subserved by particular brain areas.

Given that the two learning systems differ primarily with respect to the cortical structures involved (the basal ganglia structures are involved in both systems), I will focus primarily on the maturational course of these cortical systems. I will first review data pertaining to the

maturational course of IT and associated visual recognition functions and then pertaining to the PFC and associated executive function.

3.2.1. Maturation of the IT cortex

Maturation of the IT cortex has been extensively studied in monkeys using single-cell recording techniques. As demonstrated by several researchers (e.g., Rodman, 1994; Rodman, Skelly, & Gross, 1991), many fundamental properties of IT emerge quite early. Most importantly, as early as 6 weeks, neurons in this cortical area exhibit adult-like patterns of responsiveness. In particular, researchers presented subjects with different images (e.g., monkey faces and objects varying in spatial frequency), while recording electrical activity of IT neurons. They found that, in both infant and adult monkeys, IT neurons exhibited a pronounced form of tuning, with different neurons responding selectively to different types of stimuli. These and similar findings led researchers to conclude that the IT cortex is predisposed to rapidly develop major neural circuitry necessary for basic visual processing. Therefore, although some aspects of the IT circuitry may exhibit a more prolonged development, the basic components develop relatively early. These findings contrast sharply with findings indicating a lengthy developmental time course of PFCs (e.g., Bunge & Zelazo, 2006).

3.2.2. Maturation of the PFC

There is a wide range of anatomical, neuroimaging, neurophysiological, and neurochemical evidence indicating that the development of the PFC continues well into adolescence (e.g., Sowell, Thompson, Holmes, Jernigan, et al., 1999; see also Luciana & Nelson, 1998; Rueda et al., 2004; Davidson et al., 2006, for extensive reviews).

The maturational course of the PFC has been studied in conjunction with research on executive function—the cognitive function that depends critically on the maturity of the PFC (Davidson et al., 2006; Diamond & Goldman-Rakic, 1989; Fan, McCandliss, Sommer, Raz, & Posner, 2002; Goldman-Rakic, 1987; Posner & Petersen, 1990). Executive function comprises a cluster of abilities such as holding information in mind while performing a task, switching between tasks or between different demands of a task, inhibiting a dominant response, deliberate selection of some information and ignoring other information, selection among different responses, and resolving conflicts between competing stimulus properties and competing responses.

There is a large body of behavioral evidence that, early in development, children exhibit difficulties in deliberately focusing on relevant stimuli, inhibiting irrelevant stimuli, and switching attention between stimuli and stimulus dimensions (Diamond, 2002; Kirkham, Cruess, & Diamond, 2003; Napolitano & Sloutsky, 2004; Shepp & Swartz, 1976; Zelazo, Frye, & Rapus, 1996; Zelazo, Muller, Frye, & Marcovitch, 2003; see also Fisher, 2007, for a more recent review).

Maturation of the prefrontal structures in the course of individual development results in progressively greater efficiency of executive function, including the ability to deliberately focus on what is relevant while ignoring what is irrelevant. This is a critical step in acquiring the ability to form abstract, similarity-free representations of categories and use these

representations in both category and property induction. Therefore, the development of relatively abstract category-based generalization may hinge on the development of executive function. As suggested above, while the selection-based system could be deployed by default in adults when learning is supervised (e.g., Ashby et al., 1998), it could be that, early in development, it is the compression-based system that is deployed by default.

Therefore, there are reasons to believe that the cortical circuits that subserve the compression-based learning system (i.e., IT) come online earlier than the cortical circuits that subserve the selection-based learning system (i.e., PFC). Thus, it seems likely that, early in development, children would be more efficient in learning dense, similarity-bound categories (as these could be efficiently learned by the compression-based system) than sparse, similarity-free ones (as these require the involvement of the selection-based system).

In sum, understanding category learning requires understanding an interaction of at least three components: (a) the structure of the input, (b) the learning system that evolved to process this input, and (c) the characteristics of the learner in terms of the availability and maturity of each of the system. Understanding the interaction among these components leads to several important predictions. First, dense categories should be learned more efficiently by the nondeliberate, compression-based system, whereas sparse categories should be learned more efficiently by the more deliberate selection-based system. Second, because the critical components of the selection-based system develop late (both phylo- and ontogenetically) relative to the compression-based system, learning of dense categories should be more universal, whereas learning of sparse categories should be limited to those organisms that have a developed PFC. Third, because the selection-based system of category learning undergoes a more radical developmental transformation, learning of sparse categories should exhibit greater developmental change than learning of dense categories. Fourth, young children can spontaneously learn dense categories that are based on multiple overlapping features, whereas they should have difficulty in spontaneously learning sparse categories that have few relevant features or dimensions and multiple irrelevant features. Note that the critical aspect here is not whether a category is defined by a single dimension or by multiple dimensions, but whether the category is dense or sparse. For example, it should be less difficult to learn a color-based categorization if color is the only dimension that varies across the categories, whereas it should be very difficult to learn a color-based categorization if items vary on multiple irrelevant dimensions. And finally, given the immaturity of the selection-based system of category learning and of executive function, it seems implausible that, early in development, children can spontaneously use a single predictor as a category marker overriding all other predictors. In particular, this immaturity casts doubt on the ability of babies or even young children to spontaneously use linguistic labels as category markers in category representation. Because the issue of the role of category labels in category representation is of critical importance for understanding of conceptual development, I will focus on it in one of the sections below.

In what follows, I review empirical evidence that has been accumulated over the years, with particular focus on research generated in my lab. Although many issues remain unknown, I will present two lines of evidence supporting these predictions. First, I present evidence that category structure, learning system, and developmental characteristics of the

learner interact in category learning and category representation. In particular, early in development, the compression-based system exhibits greater efficiency than the selection-based system. In addition, early in development, categories are represented perceptually, and only later do participants form more abstract, dimensional, rule-based, or lexicalized representations of categories. And second, the role of words in category learning is not fixed; rather, it undergoes developmental change: Words initially affect processing of visual input, and only gradually they become category markers.

4. Interaction among category structure, learning system, and characteristics of the learner: Evidence from category learning and category representation

Recall that I hypothesized an interaction among (a) the structure of the category (in particular, its density), (b) the learning system that evolved to process this input, and (3) the characteristics of the learner in terms of the availability and maturity of each system. In what follows, I consider components of this interaction with respect to category learning and category representation.

4.1. Category learning

As discussed earlier, there are reasons to believe that, in the course of individual development, the compression-based system comes online earlier than the selection-based system (i.e., due to the protracted immaturity of the executive function that subserves the selection-based system). Therefore, it seems plausible that, at least early in development, the compression-based system is deployed by default, whereas the selection-based system has to be triggered explicitly (see Ashby et al., 1998 for arguments that this may not be the case in adults). It is also possible that there are experimental manipulations that could trigger the nondefault system. In particular, the selection-based system could be triggered by explicit supervision or an error signal.

If the systems are dissociated, then sparse categories that depend critically on selective attention (as they require focusing on a few relevant dimensions, while ignoring irrelevant dimensions) may be learned better under the conditions triggering the selection-based system. At the same time, dense categories that have much redundancy may be learned better under the conditions of implicit learning. Finally, because dense categories could be efficiently learned by the compression-based system, which is more primary, both phylo- and ontogenetically, learning of dense categories should be more universal than learning of sparse categories. In what follows, I review evidence exemplifying these points.

4.1.1. Interactions between category structure and the learning system

In a recent study (Kloos & Sloutsky, 2008), we demonstrated that category structure interacts with the learning system as well as with characteristics of the learner. In this study, 5-year-olds and adults were presented with a category learning task where they learned

either dense or sparse categories. These categories consisted of artificial bug-like creatures that had a number of varying features: sizes of tail, wings, and fingers; the shadings of body, antenna, and buttons; and the numbers of fingers and buttons (see Fig. 3, for examples of categories). Category learning was administered under either an unsupervised, spontaneous learning condition (i.e., participants were merely shown the items) or under a supervised, deliberate learning condition (i.e., participants were told the category inclusion rule). Recall that the former learning condition was expected to trigger the compression-based system of category learning, whereas the latter was expected to trigger the selection-based system. If category structure interacts with the learning system, then implicit, unsupervised learning should be more optimal for learning dense categories, whereas explicit, supervised learning should be more optimal for learning sparse categories. This is exactly what was found: For both children and adults, dense categories were learned better under the unsupervised, spontaneous learning regime, whereas sparse categories were learned more efficiently under the supervised learning regime. Critical data from this study are presented in Fig. 4. The figure presents categorization accuracy (i.e., the proportion of hits, or correct identification of category members minus the proportion of false alarms, or confusion of nonmembers for members) after the category learning phase.

These findings dovetail with results reported by Yamauchi, Love, and Markman (2002) and Yamauchi and Markman (1998) in adults. In these studies, participants completed a category learning task that had two learning conditions: classification and inference. In the classification condition, participants learned categories by predicting category membership of each study item. In the inference condition, participants learned categories by predicting a feature shared by category members. Across the conditions, results revealed a category structure by learning condition interaction. In particular, NLS categories (which are
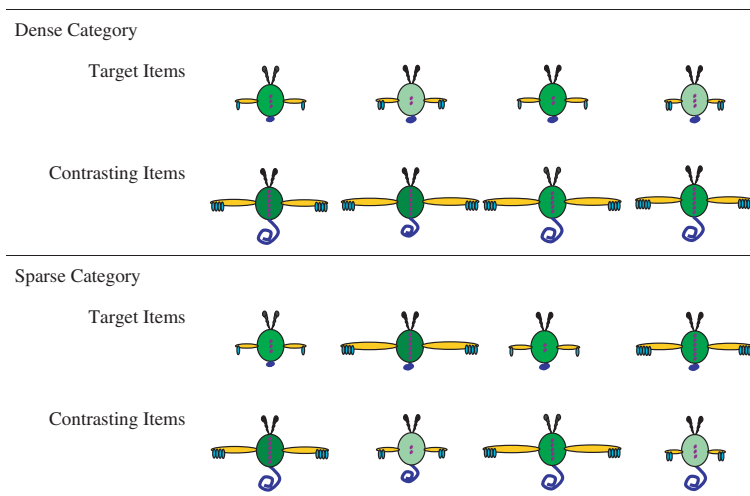

Fig. 3. Examples of items used from Kloos and Sloutsky (2008), Experiment 1. In the dense category, items are bound by similarity, whereas in the sparse category, the length of the tale is the predictor of the category membership.
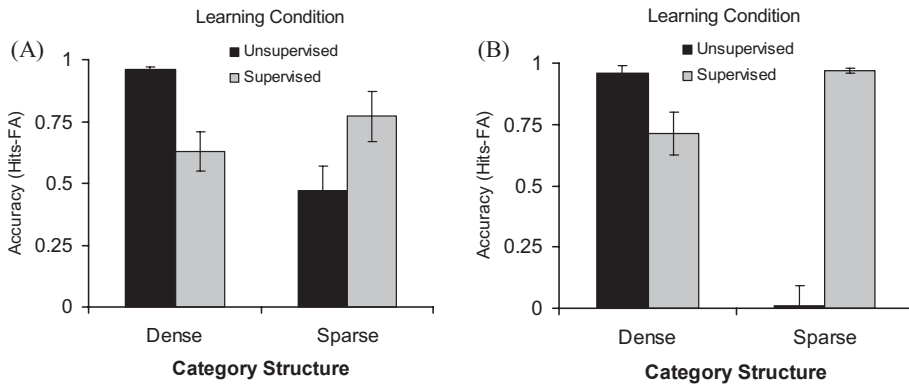
Fig. 4. Mean accuracy scores by category type and learning condition in adults (A) and in children (B). In this and all other figures, error bars represent standard errors of the mean. For the dense category, $D = 1$, and for the sparse category, $D = 0.17$.

typically sparser) were learned better in the classification condition, whereas prototype-based categories (which are typically denser) were learned better in the inference condition.

The interaction between the category structure and the learning system has been recently demonstrated by Hoffman and Rehder (2010), with respect to the cost of selectivity in category learning. Similar to Yamauchi and Markman (1998), participants learned categories either by classification or by feature inference. In the classification condition, participants were presented with two categories (e.g., A and B). On each trial, they saw an item and their task was to predict whether the item in question is a member of A or B. In the inference condition, participants were also presented with categories A and B. On each trial, they saw an item that had one missing feature and their task was to predict whether it was a feature common to A or common to B. In both conditions, upon responding, participants received feedback.

Each category had three binary dimensions whose values were designated as 0 or 1. There were two learning phases. In Phase 1, participants learned two categories A and B, with Dimensions 1 and 2 distinguishing between the categories and Dimension 3 being fixed across the categories (e.g., all items had a value of 0 on the fixed Dimension 3). In Phase 2, participants learned two other Categories C and D, with Dimensions 1 and 2 again distinguishing between the categories and Dimension 3 being fixed again (e.g., now items had a value of 1 on the fixed Dimension 3). After the two training phases, participants were given categorization trials involving contrasts between categories that were not paired during training (e.g., A vs. C). Note that correct responding on these novel contrasts required attending to Dimension 3, which had been previously irrelevant during training. If participants attend selectively to dimensions, their attention should have been allocated to Dimensions 1 and 2 during learning, which should have attenuated attention to Dimension 3. This attenuated attention represents the cost of selectivity. Alternatively, if no selectivity is involved, there should be little or no attenuation, and therefore, little or no cost. It was found that the cost was higher for classification learners than for inference learners, thus

suggesting that classification learning, but not inference learning, engages the selection-based system.

4.1.2. Developmental primacy of the compression-based system

Zentall et al. (2008) present an extensive literature review indicating that although birds, monkeys, apes, and humans are capable of learning categories consisting of highly similar yet discriminable items (i.e., dense categories), only some apes and humans could learn sparse relational categories, such as ''sameness'' when an equivalence class consisted of dissimilar items (e.g., a pair of red squares and a pair of blue circles are members of the same sparse category). However, even here it is not clear that subjects were learning a sparse category. As shown by Wasserman, Young, and Cook (2004), nonhuman animals readily distinguish situations with no variability in the input (i.e., zero entropy) from situations where input has stimulus variability (i.e., nonzero entropy). Therefore, it is possible that learning was based on the distinction between zero entropy in each of the ''same'' displays and nonzero entropy in each of the ''different'' displays.

The idea of the developmental primacy of the compression-based system is supported by data from Kloos and Sloutsky (2008) reviewed earlier. In particular, data presented in Fig. 4 clearly indicate that, for both children and adults, sparse categories were learned better under the explicit, supervised condition, whereas dense categories were learned better under the implicit, unsupervised condition. Also note that adults learned the sparse category even in the unsupervised condition, whereas young children exhibited no evidence of learning. These data support the contention that the compression-based system is the default in young children.

In addition, data from Kloos and Sloutsky (2008) indicate that although both children and adults exhibited able spontaneous learning of a dense category, there were marked developmental differences in spontaneous learning of sparse categories. Categorization accuracy in the spontaneous condition by category density and age is presented in Fig. 5. Two aspects of these data are worth noting. First, there was no developmental difference in spontaneous
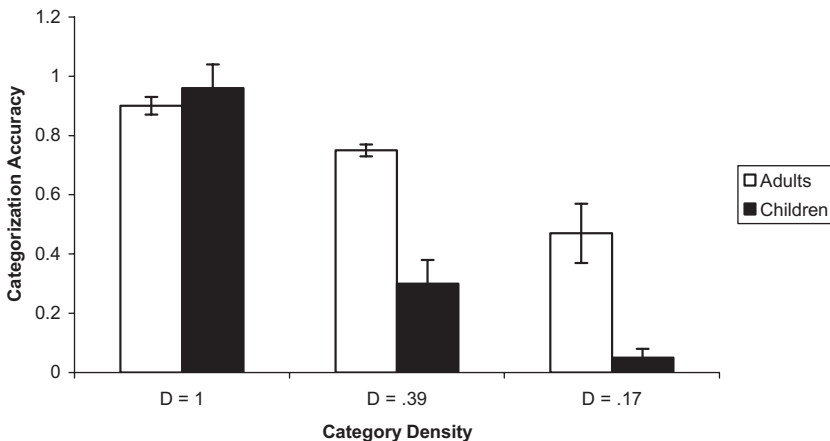

Fig. 5. Unsupervised category learning by density and age group from Kloos and Sloutsky (2008).

learning of the very dense category, which suggests that the compression-based system of category learning exhibits the adult level of functioning in 4- to 5-year-olds. And second, there were substantial developmental differences in spontaneous learning of sparser categories, which suggests that adults, but not young children, may spontaneously deploy the selection-based system of category learning. Therefore, the marked developmental differences pertain mainly to the deployment and functioning of the selection-based system, but not of the compression-based system (see also Hammer, Diesendruck, Weinshall, & Hochstein, 2009, for related findings).

Additional evidence for the developmental primacy of the compression-based learning system stems from research demonstrating that young children can learn complex contingencies implicitly, but not explicitly (Sloutsky & Fisher, 2008). The main idea behind the Sloutsky and Fisher (2008) experiments was that implicit (and perhaps compression-based) learning of complex contingencies might underlie seemingly selective generalization behaviors of young children. There is much evidence suggesting that, even early in development, people's generalization could be selective—depending on the situation, people may rely on different kinds of information. This selectivity has been found in a variety of generalization tasks, including lexical extension, categorization, and property induction. For example, in a lexical extension task (Jones et al., 1991), 2- and 3-year-olds were presented with a named target (i.e., ''this is a dax''), and then were asked to find another dax among test items. Children extended the label by shape alone when the target and test objects were presented without eyes. However, they extended the label by shape and texture when the objects were presented with eyes.

Similarly, in a categorization task, 3- and 4-year-olds were more likely to group items on the basis of color if the items were introduced as food, but group on the basis of shape if the items were introduced as toys (Macario, 1991). More recently, Opfer and Bulloch (2007) examined flexibility in lexical extension, categorization, and property induction tasks. It was found that across these tasks, 4- to 5-year-olds relied on one set of perceptual predictors when the items were introduced as ''parents and offspring,'' whereas they relied on another set of perceptual predictors when items were introduced as ''predators and prey.'' These findings pose an interesting problem—is this putative selectivity subserved by the selection-based system or by the compression-based system? Given critical immaturities of the selection-based system early in development, the latter possibility seems more plausible. Sloutsky and Fisher's (2008) study supported this possibility.

A key idea is that many stimulus properties intercorrelate, such that some clusters of properties co-occur with particular outcomes, and other clusters co-occur with different outcomes, thus resulting in a dense ''context–outcome'' structures (cf., with the idea of ''coherent covariation'' presented in Rogers & McClelland, 2004). Learning these correlations may result in differential allocation of attention to different stimulus properties in different situations or contexts, with flexible generalizations being a result of this learning. In particular, participants could learn the following set of contingencies: In Context 1, Dimension 1 (say, color) was predictive, but Dimension 2 (say, shape) was not, whereas the reverse is true in Context 2. If, as argued earlier, the system of implicit compression-based learning is fully functioning even early in development, then the greater the number of contextual

variables correlating with the relevant dimension (i.e., the greater the density), the greater the likelihood of learning. However, if learning is selection-based the reverse may be the case. This is because the larger the number of relevant dimensions, the more difficult it could be to formulate a contingency as a simple rule.

These possibilities have been tested in multiple experiments reported in Sloutsky and Fisher (2008). In these experiments, 5-year-olds were presented with triads of geometric objects differing in color and shape. Each triad consisted of a Target and two Test items. Participants were told that a prize was hidden behind the Target and their task was to determine the Test item that had a prize behind it. Children were trained that, in Context 1, shape of an item was predictive of an outcome, whereas in Context 2 color was predictive. Context was defined as the color of the background on which stimuli appeared and the location of the stimuli on the screen. Therefore, in Context 1, training stimuli appeared on a yellow background in the upper-right corner of the computer screen, and on a green background in the bottom-left corner of the computer screen in Context 2. Training stimuli were triads each consisting of a target and two test items. Participants were given information about a target item and they had to generalize this information to one of the test items. Each participant was given three training blocks. In one training block, only color was predictive, in another training block, only shape was predictive, whereas the third block was a mixture of the former two blocks. Participants were then presented with testing triads that had an important difference from training triads. Whereas training triads were ''unambiguous'' in that only one dimension of variation (either color or shape) was predictive and only one test item matched the target on the predictive dimension, this was not the case for testing triads. In particular, testing triads were ''ambiguous'' in that one test item matched the target on one dimension and the other test item matched the target on the other dimension. The only disambiguating factor was the context.

It was found that participants had no difficulty in learning the contingency between the context and the predictive dimension when there were multiple contextual variables correlating with the predictive dimension. In particular, children tested in Context 1 primarily relied on shape and those in Context 2 primarily relied on color. Learning, however, attenuated markedly when the number of contextual variables was reduced, which should not have happened if learning was selection based. And finally, when presented with testing triads and explicitly given a simple rule (e.g., children were asked to make choices by focusing either on color or on shape), they were unable to focus on the required dimension. These findings present further evidence for the developmental asynchrony of the two learning systems: Although 5-year-old children could readily perform the task when relying on the compression-based learning system, they were unable to perform the task when they had to rely on the selection-based system.

In sum, there is emerging body of evidence from category learning suggesting an interaction between the category structure and the learning system, pointing to developmental asynchronies in the two systems. Future research should reexamine category structure and category learning in infancy. In particular, given the critical immaturity of the selection-based system, most (if not all) of category learning in infancy should be accomplished by the compression-based system.

4.2. Category representation

In the previous section, I reviewed evidence indicating that category learning is affected by an interaction among category structure, the learning systems processing this structure, and the characteristics of the learner. In this section, I will review evidence demonstrating components of this interaction for category representation. Most of the evidence reviewed in this section pertains to developmental asynchronies between the learning systems. Two interrelated lines of evidence will be presented: (a) the development of selection-based category representation and (b) the changing role of linguistic label in category representation.

4.2.1. The development of selection-based category representation

If the compression-based and the selection-based learning systems mature asynchronously, such that early in development the former system exhibits greater maturity than the latter, then it is likely that most of the spontaneously acquired categories are learned implicitly by the compression-based learning system. If this is the case, it is unlikely that young children form abstract rule-based representations of spontaneously acquired categories, whereas they are likely to form perceptually rich representations. A representation of a category is abstract if category items are represented by either a category inclusion rule or by a lexical entry. A representation of a category is perceptually rich if category representation retains (more or less fully) perceptual detail of individual exemplars.

One way of examining category representation is focusing on what people remember about category members. For example, Kloos and Sloutsky (2008, Experiment 4B) presented 5-year-olds and adults with a category learning task. Similar to the above-described experiment by Kloos and Sloutsky (2008), there were two between-subjects conditions, with some participants learning a dense category and some learning a sparse category. Both categories consisted of the described above artificial bug-like creatures that had a number of varying features: sizes of tail, wings, and fingers; the shadings of body, antenna, and buttons; and the numbers of fingers and buttons. The relation between the two latter features defined the arbitrary rule: Members of the target category had either many buttons and many fingers or few buttons and few fingers. All the other features constituted the appearance features. Members of the target category had a long tail, long wings, short fingers, dark antennas, a dark body, and light buttons (target appearance A_T), whereas members of the contrasting category had a short tail, short wings, long fingers, light antennas, a light body, and dark buttons (contrasting appearance A_C). All participants were presented with the same set of items; however, in the sparse condition participants' attention was focused on the inclusion rule, whereas in the dense condition it was focused on appearance information. This was achieved by varying the description of items across the conditions. In the sparse-category condition, the description was as follows: ''Ziblets with many aqua fingers on each yellow wing have many buttons, and Ziblets with few aqua fingers on each yellow wing have few buttons.'' In the dense-category condition, in addition to the above-described rule, the appearance of exemplars was described. In both conditions, appearance features were probabilistically related to

category membership, whereas the rule was fully predictive. After training, participants were tested on their category learning and then presented with a surprise recognition task. During the recognition phase, they were presented with four types of recognition items: $A_T R_T$ (the items that had both the appearance and the rule of the Target category), $A_C R_C$ (the items that had both the appearance and the rule of the Contrast category), $A_T R_C$ (the items that had the appearance of Target category and the rule of the Contrast category), and $A_C R_T$ (the items that had the appearance of the Contrast category and the rule of the Target category). If participants learned the category, they should accept $A_T R_T$ items and reject $A_C R_C$ items. In addition, if participants' representation of the category is based on the rule, they may false alarm on $A_C R_T$, but not on $A_T R_C$ items. However, if participants' representation of the category is based on the appearance, they should false alarm on $A_T R_C$, but not on $A_C R_T$ items.

False alarm rates by age and test item type are presented in Fig. 6. As can be seen in the figure, adults were more likely to false alarm on same appearance items (i.e., $A_T R_C$) in the dense condition and on same rule items (i.e., $A_C R_T$) in the sparse condition. In contrast, young children were likely to false alarm on same appearance items (i.e., $A_T R_C$) in both conditions. These results suggest that, in adults, dense and sparse categories could be represented differently: The former are represented perceptually, whereas the latter are represented more abstractly. At the same time, 5-year-old children are likely to represent perceptually both dense and sparse categories. These data suggest that the representation of sparse (but not dense) categories changes in the course of development.

These findings, however, were limited to newly learned categories that were not lexicalized. What about the representation of lexicalized dense categories? One possibility is that lexicalized dense categories are also represented perceptually, similar to newly learned dense categories. In this case, there should be no developmental differences in the representation of lexicalized dense categories. However, representations of lexicalized dense categories may include the linguistic label (which could be the most reliable guide to category membership). In particular, it is possible that lexicalization of a perceptual grouping eventually results in an abstract label-based representation (in the limit, a member of a category
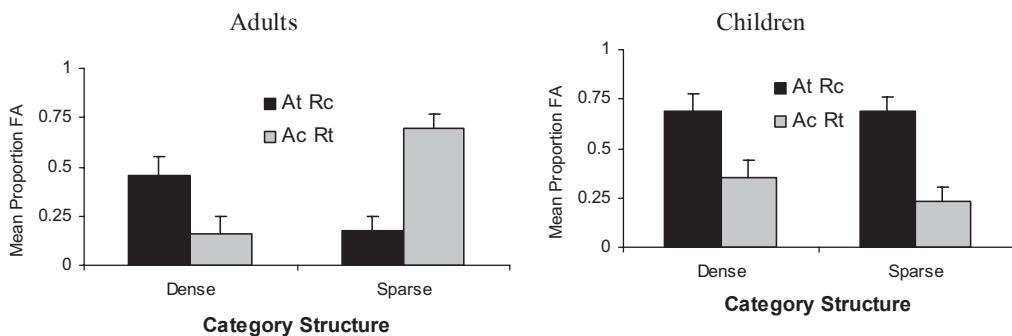

Fig. 6. False alarm rate by category structure and foil type in adults and children from Kloos and Sloutsky (2008), Experiment 4.

could be represented just by its label). If this is the case, then there should be substantial developmental differences in the representation of lexicalized dense categories. Furthermore, in this case, adults should differently represent highly familiar lexicalized dense categories (e.g., cat) and newly learned nonlexicalized dense categories (e.g., categories consisting of bug-like creatures). In particular, they should form an abstract representation of the former, but not the later.

These possibilities have been examined in a set of recognition memory experiments (e.g., Fisher & Sloutsky, 2005; Sloutsky & Fisher, 2004a, 2004b). If participants form abstract representation of category items, then a task that prompts categorization of items may result in attenuated memory for appearance information. This reasoning is based on a long tradition of false memory research demonstrating that deep semantic processing of studied items (including grouping of items into categories) often increases memory intrusions—false recognition and recall of nonpresented ''critical lures'' or items semantically associated with studied items (e.g., Koutstaal & Schacter, 1997; Thapar & McDermott, 2001). Thus, ''deeper'' processing can lead to lower recognition accuracy when critical lures are semantically similar to studied items. In contrast to deep processing, focusing on perceptual details of pictorially presented information results in accurate recognition (Marks, 1991).

Therefore, if a recognition memory task is presented after a task that encourages access to the abstract representation of familiar categories, patterns of recognition errors may reveal information about how categories are represented. If participants processed items relatively abstractly as members of a category, then they would be more likely to have difficulty in discriminating studied targets from conceptually similar critical lures. If, on the other hand, they processed items more concretely, focusing on perceptual details, then they should discriminate relatively well.

In a set of experiments, Fisher and Sloutsky (2005) presented adults with one of two tasks. In the Baseline condition, the task was to remember items as accurately as possible, whereas in the Induction condition, the task was to generalize a property from a target item to each presented item. In both conditions, study phase items consisted of several categories, with multiple items per category. Following this study phase, participants in both conditions were presented with a surprise recognition task. Recognition items included Old Items (those presented during the Study phase), Critical Lures (novel items from studied categories), and Unrelated Items (novel items from new categories). If participants accept Old Items and Critical Lures, but reject Unrelated Items, then it is likely that they represented only abstract category information, not appearance information. However, if they accept only Old Items, but reject Critical Lures and Unrelated Items, then it is likely that they represented appearance information.

In one experiment reported by Fisher and Sloutsky (2005), adults were presented with familiar lexicalized dense categories (e.g., cats, bears, etc.), whereas in another condition, dense categories included artificial bug-like creatures, similar to those used by Kloos and Sloutsky (2008). Memory accuracy (which is a function of hits and false alarms on Critical Lures) by condition and category type in adults is presented in Fig. 7. Note that the dependent variable is A-prime (A-prime is a nonparametric analog of the signal-detection d-prime statistic), and the value of 0.5 represents no discrimination between Old Items and Critical
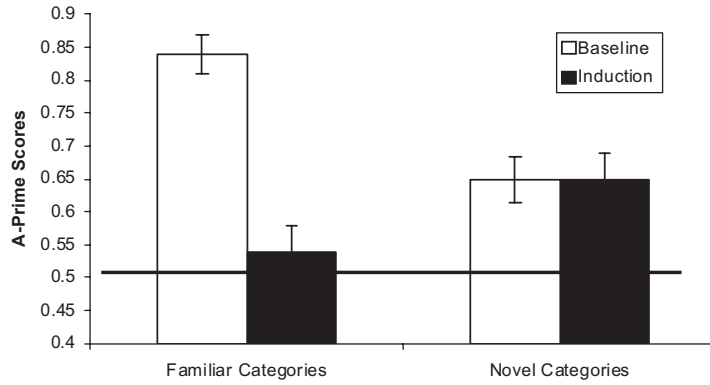
Fig. 7. Recognition accuracy in adults by category familiarity and study phase condition from Fisher and Sloutsky (2005).

Lures. When categories were familiar, adults were accurate in the Baseline condition, whereas they did not distinguish between Old Items and Critical Lures in the Induction condition. This *category processing effect* indicates that adults form a relatively abstract representation of familiar (and lexicalized) dense categories. It is also possible that category label plays an important role in such a representation (cf., findings reported by Tipper & Driver, 2000 on priming between pictures of objects and their labels in adults). At the same time, when categories were novel, adults were accurate in both the Baseline and Induction condition. Therefore, perceptual information plays an important role in the representation of novel dense categories in adults.

In contrast to adults, young children do not exhibit evidence of abstract representation of even familiar dense categories. As shown in Fig. 8, after performing induction with pictures
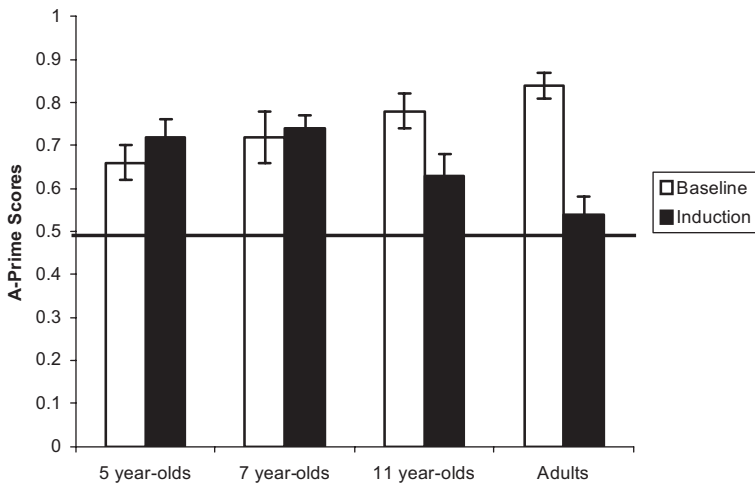

Fig. 8. Recognition accuracy by age and study phase condition from Fisher and Sloutsky (2005).

of members of familiar categories (e.g., cats), young children exhibited greater recognition accuracy than did adults, with recognition gradually decreasing with increasing age (Fisher & Sloutsky, 2005; Sloutsky & Fisher, 2004a, 2004b). The figure depicts A-prime scores across the conditions and the difference in A-prime score between the Baseline and Induction conditions reflects the ''category processing effect''—a decreased recognition of categorized items compared with the baseline. As shown in the figure, there is no evidence of the category processing effect early in development, and even in preadolescence the magnitude of the effect is smaller than that in adults. Recall that when adults were given the same task with novel items for which they did not have compressed category representation, their recognition accuracy increased to the level of young children (see Fig. 7).

These findings in conjunction with the relative immaturity of the executive function in 4- and 5-year-olds suggest that these participants, even if they learn a sparse rule-based category, would be unable to use this learned category in other tasks. It has been often argued that one of the most important roles of categories is to support inductive generalization. If one learns that an individual has a particular property (e.g., a particular dog likes bones), one could generalize this property to other members of this category. Although most transient properties (e.g., is awake) cannot be generalized, many stable properties can. Therefore, examining the pattern of inductive generalization could elucidate how categories are represented. If participants do not form an abstract representation of a sparse category, they would be unable to use the category in induction.

One way of addressing this issue is to teach participants a novel sparse rule-based category. Once participants learn the category, they could be presented with a property induction task, in which they could rely either on the rule or on appearance information, which is irrelevant for category membership. If young children represent the category by an abstract rule, they should use this representation when performing inductive generalization. Conversely, if they represent appearance of the items, then young children (even when they successfully learn the category) should rely on appearance information, while disregarding category membership information. These possibilities were tested in a set of experiments reported by Sloutsky, Kloos, and Fisher (2007). In these experiments, participants were first presented with a category learning task during which they learned two categories of artificial animals. Category membership was determined by a rule, whereas perceptual similarity was not predictive of category membership. Children were then given a categorization task with items that differed from those used during training. Participants readily acquired these categories and accurately sorted the items according to their category information. Then participants were presented with a triad induction task. Each triad consisted of a target and two test items, with one test item sharing the target's category membership, and the other test item being similar to the target (without sharing category membership). Participants were familiarized with a quasi-biological property of the target and asked to generalize this property to one of the test items. Finally, participants were given a final (i.e., postinduction) categorization task using the same items as the induction task. The results indicate that, although participants learned the category-inclusion rule, they did not use it in the course of induction, rather basing their induction on perceptual information.

In sum, early in development, similarity plays an important role in the representation of even sparse categories, whereas later in development categories may be represented in a more abstract manner. One possibility is that, later in development, labels begin to play a more central role in category representation.

4.2.1.1 The developing role of linguistic labels in category representation: In the previous section, I reviewed evidence that in young children (in contrast to adults) a category label does not figure prominently in category representation. This developmental change in the role of category labels represents another source of evidence for the developmental asynchronies between the two systems of category learning. In this section, I focus on the changing role of category labels in greater detail.

To examine the role of linguistic labels in category representation of adults, Yamauchi and colleagues conducted a series of studies supporting the idea that for adults a label is a symbol that represents a category (Yamauchi & Markman, 2000; Yamauchi & Yu, 2008). The overall reasoning behind this work is that if labels are category markers, they should be treated differently from the rest of features (such shape, color, size, etc.). However, this may not be the case if labels are features. Therefore, inferring a label when features are given (i.e., a classification task) should elicit different performance from a task of inferring a feature when the label is given (i.e., a feature induction task).

To test these ideas, Yamauchi and Markman (2000) used the above-described category learning task that was presented under either classification or feature induction learning condition. There were two categories, C_1 and C_2 denoted by two labels, L_1 and L_2. Stimuli were bug-like artificial creatures that varied on several dimensions, with one range of values determining C_1 and another range of values determining C_2. In the feature induction task, participants were shown a creature with one missing feature and were given a category label. Their task was to predict the missing feature. In the classification task, they were presented with a creature that was not labeled, and the task was to predict the category label. The critical condition was the case when an item was a member of C_1, but was similar to C_2, with the dependent variable being the proportion of C_1 responses. The results indicated that there were significantly more category-based responses in the induction condition (where participants could rely on the category label) than in the categorization condition (where participants had to infer the category label). It was concluded therefore that category labels differed from other features in that participants treated labels as category markers. These findings have been replicated in a series of follow-up studies (Yamauchi, Kohn, & Yu, 2007; Yamauchi & Yu, 2008; see also Markman & Ross, 2003, for a review). For example, Yamauchi et al. (2007) examined patterns of mouse-tracking (a procedure that is similar to eye tracking) to examine attention allocated to labels when labels were introduced as category markers (e.g., ''This is a dax'') or as denoting category features (e.g., ''This one has a dax''). Results indicated that participants viewed these visually presented labels more often in the former condition than in the latter condition. In sum, there is a body of evidence indicating that adults tend to treat the category label as a category marker rather than as a category feature.

However, the reliance on category labels in category representation requires the involvement of the selection-based system. At the same time, if the selection-based system exhibits

a slow developmental course, the ability to use category labels as category markers should be limited early in development. Furthermore, simultaneous processing of auditory and visual input (e.g., an object and corresponding sound) requires the ability to integrate information coming from different modalities. This ability also exhibits a relatively slow maturational course (see Robinson & Sloutsky, 2010, for a review) and is unlikely to be fully functional in infancy. In part, this slow maturational course in the ability to integrate cross-modal information could be related to a slow maturational course of neurons processing multisensory information. For example, there is evidence from animal models indicating that multisensory neurons located in the superior colliculus and at various cortical locations do not mature until the sufficient visual experience is accumulated (see Wallace, 2004, for a review).

If the contribution of labels to categorization and category learning hinges on (a) the ability to process cross-modal information and (b) the ability to attend selectively, with both abilities undergoing substantial developmental change, then the role linguistic labels play in categorization and category learning may change across development. In what follows, I review evidence indicating the changing role of category labels and consider possible mechanisms underlying these developmental changes.

As my colleagues and I have argued elsewhere, auditory input may affect attention allocated to corresponding visual input (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003; Sloutsky & Robinson, 2008), and these effects may change in the course of learning and development. In particular, linguistic labels may strongly interfere with visual processing in prelinguisitic children, but these interference effects may weaken when children start acquiring language (Sloutsky & Robinson, 2008; see also Robinson & Sloutsky, 2007a, 2007b).

In one experiment, Sloutsky and Robinson (2008) familiarized 10- and 16-month-olds with auditory–visual compounds. The familiarization compound consisted of a three-shape pattern and a word presented at the same time (both the word and the three-shape pattern were ably processed by infants of these age groups when presented unimodally). The familiarization phase was followed by the test phase, in which participants were presented with four different auditory–visual test items. One test item was the familiarization compound ($AUD_{Target}VIS_{Target}$), one had a changed visual component ($AUD_{Target}VIS_{New}$), one had a changed auditory component ($AUD_{New}VIS_{Target}$), and one had both components changed ($AUD_{New}VIS_{New}$).

The dependent variable was looking time at each test item. If participants considered a test item to be different from the familiarization item, looking time to this item should increase compared with the end of familiarization. Because the $AUD_{Target}VIS_{Target}$ is the familiarization item, it should elicit looking that is comparable with looking at the end of familiarization phase. Because the $AUD_{New}VIS_{New}$ is a novel item, it should elicit longer looking. At the same time, looking at $AUD_{Target}VIS_{New}$ and $AUD_{New}VIS_{Target}$ items should depend on whether participants processed auditory and visual components of the familiarization compound. If infants did, they should increase looking to both test items. If infants processed only the auditory component, they should increase looking only to $AUD_{New}VIS_{Target}$ item, whereas if they processed only the visual component, they should

increase looking only to $AUD_{Target}VIS_{New}$ item. Looking times to $AUD_{Target}VIS_{New}$, $AUD_{New}VIS_{Target}$, and $AUD_{New}VIS_{New}$ items compared with the $AUD_{New}VIS_{Target}$ item are presented in Fig. 9. These results clearly indicate that although 10-month-old infants failed to process the visual component, 16-month-old infants processed both components. It was concluded therefore that linguistic input interfered with processing of visual input at 10 months of age, but these interference effects weakened by 16 months of age.

In another experiment, Robinson and Sloutsky (2007a) presented 8- and 12-month-olds with a categorization task. Participants were familiarized with category exemplars under one of the three conditions: (a) all items were accompanied by the same label, (b) all items were accompanied by the same sound, or (c) all items were presented in silence. At test, participants were presented with two types of test trials: (a) recognition trials (i.e., a studied item was paired with a new item) and (b) categorization trials (i.e., a novel in-category exemplar was paired with a novel out-of-category exemplar). If participants recognize the studied item, they should prefer looking to the novel item, and if they learned the category, they should prefer looking to an out-of-category item. Results indicated that performance was significantly better in the silent condition, thus suggesting that both sounds and labels interfered with the categorization task. Similar results were reported for individuation tasks (Robinson & Sloutsky, 2008).

By the onset of word learning, children should start acquiring the ability to integrate linguistic and visual input (Robinson & Sloutsky, 2007b; Sloutsky & Robinson, 2008). However, even then cross-modal processing may not reach the full level of maturity and therefore linguistic labels may attenuate the processing of corresponding visual input. As
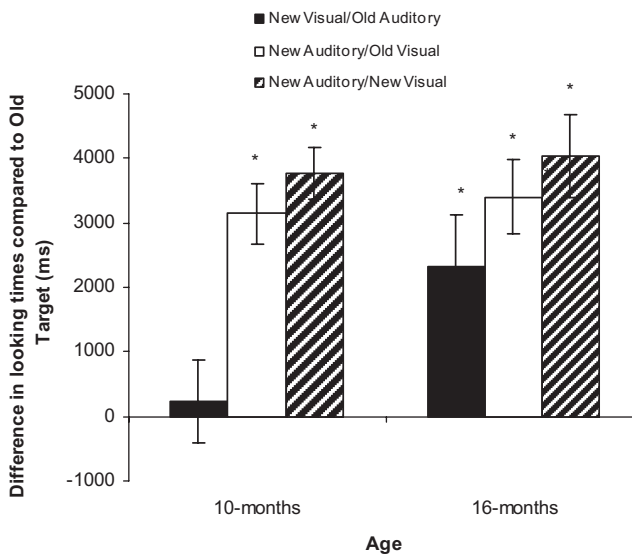

Fig. 9. Differences in looking times by Age and Test item type from Sloutsky and Robinson (2008). *Difference scores >0, $p < .05$.

discussed below, this attenuated processing may result in an increased similarity of entities that have the same label and thus in an increased tendency to group them together (e.g., Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999; Sloutsky, Lo, & Fisher, 2001).

Although interference effects attenuate with development, they do not disappear completely. This issue has been examined in depth in a series of recognition experiments (e.g., Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003).

In these recognition experiments, 4-year-olds and adults were presented with a compound Target stimulus, consisting of simultaneously presented auditory and visual components ($AUD_{Target}VIS_{Target}$). These experiments were similar to the above-described experiment, except that no learning was involved. Participants were presented with a Target, which was followed immediately by a Test item and the task was to determine whether the Target and Test items were exactly the same.

There were four types of test items: (a) $AUD_{Target}VIS_{Target}$, which was the Old Target item; (b) $AUD_{Target}VIS_{New}$, which had the target auditory component and a new visual component; (c) $AUD_{New}VIS_{Target}$, which had the target visual component and a new auditory component; or (d) $AUD_{New}VIS_{New}$, which had a new visual component and a new auditory component. The task was to determine whether each presented test item was exactly the same as the Target (i.e., both the same auditory and visual components) or a new item (i.e., differed on one or both components).

Similar to the experiment with infants (Robinson & Sloutsky, 2004), it was reasoned that if participants process both auditory and visual stimuli, they should correctly respond to all items by accepting Old Target items and rejecting all other test items. Alternatively, if they fail to process the visual component, they should falsely accept $AUD_{Target}VIS_{New}$ items, while correctly responding to other items. Finally, if they fail to process the auditory component, they should falsely accept $AUD_{New}VIS_{Target}$ items, while correctly responding to other items. In one experiment (Napolitano & Sloutsky, 2004), speech sounds were paired with either geometric shapes or pictures of unfamiliar animals. Results indicated that although children ably processed either stimulus in the unimodal condition, they failed to process visual input in the cross-modal condition. Furthermore, a yet unpublished study by Napolitano and Sloutsky indicates that interference effects attenuate gradually in the course of development, with very little evidence of interference in adults.

There is also evidence that this dominance of auditory input is not under strategic control: Even when instructed to focus on visual input young children had difficulties doing so (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004). In one of the experiments described in Napolitano and Sloutsky (2004), 4-year-olds were explicitly instructed to attend to visual stimuli, with instructions repeated before each trial. However, despite the repeated explicit instruction to attend to visual stimuli, 4-year-olds continued to exhibit auditory dominance. These results suggest that auditory dominance is unlikely to stem from deliberate selective attention to a particular modality, but it is more likely to stem from automatic pulls on attention.

If linguistic labels attenuate visual processing, such that children ably process a label, but they do so to a lesser extent the corresponding visual input, then these findings can explain the role of labels in categorization tasks. In particular, items that share a label may appear

more similar than the same items presented without a label. In other words, early in development, labels may function as features contributing to similarity, and their role may change in the course of development. In fact, there is evidence supporting this possibility (e.g., Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999).

The key idea behind these experiments is if two items have a particular degree of visual similarity, then adding a common label would increase this similarity due to the above-described attenuated visual processing. These effects have been demonstrated with a frequently used forced choice task, where participants are expected to make either a similarity judgment (i.e., which one of the several test items looks more like the target) or a categorization judgment (i.e., which one of the several test items belongs to the same kind as the target).

In this case, the probability of selecting a particular test item is a function of a *ratio* of the similarity of a given test item to the Target to the summed similarity of other test items to the Target. In this case, the common label affects the similarity ratio. These ideas have been implemented in model SINC (for Similarity, Induction, Naming, and Categorization; Sloutsky & Lo, 1999; Sloutsky & Fisher, 2004a) that accurately predicted similarity and categorization judgment in young children when labels were and were not introduced.

In these experiments, young children were presented with triads of items (a Target and two Test items) and were asked which of the Test items looked more similar to the Target. One of the test items (e.g., Test A) was very similar to the Target, whereas similarity of the other test item (say Test B) varied across trials from very similar to very different. In the Baseline condition, labels were not provided, whereas in the Label condition, one of the Test items shared the label with the Target, whereas the other Test item did not. The labels were artificial bisyllabic count nouns. Proportions of selecting Test B as more similar to the Target by condition and similarity ratio (Test B–Target/Test A–Target) are presented in Fig. 10A. As can be seen in the figure, the presence of labels increased similarity for all levels of similarity. However, when the same task was given to adults (Fig. 10B), labels had no effect on similarity judgment.

Therefore, it seems that labels function differently across development: Whereas labels are likely to contribute to similarity of compared items in children (e.g., Sloutsky & Fisher,
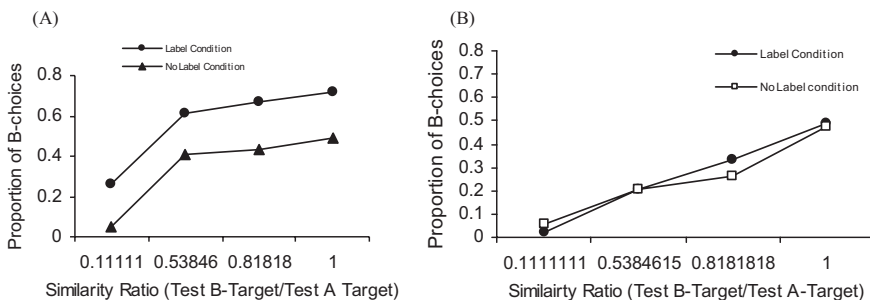

Fig. 10. Similarity judgment by similarity ratio and labeling condition in (A) children and (B) adults from Sloutsky and Fisher (2004a).

2004a; Sloutsky & Lo, 1999), they are not likely to do so in adults (Yamauchi & Markman, 2000).

There is also evidence that labels have similar effects on categorization—these effects are also graded rather than rule-like, with labels affecting, but not overriding perceptual similarity (e.g., Sloutsky & Fisher, 2004a). In several experiments conducted by Sloutsky and Fisher, 4- and 5-year-olds performed a match-to-sample categorization task. On each trial, they were presented with a triad of pictures, a target, and two test items. All items were labeled and only one of the test items shared the label with the target. Participants were asked to decide which of the test items belongs to the same kind as the Target. Strikingly similar patterns were observed for categorization and feature induction tasks in young children: Again, participants' categorization and induction responses were affected by the similarity ratio, with labels contributing to these effects of similarity rather than overriding them.

In yet another experiment, Sloutsky and Fisher (2004a, 2004b) used items that had been previously used by Gelman and Markman (1986), which turned out to vary widely in terms of appearance similarity. Again, there was little evidence that, in their induction responses, 4- and 5-year-olds relied exclusively on linguistic labels.

In short, the reviewed evidence supports the idea that young children treat labels as perceptual features that contribute to similarity of compared entities. It seems that these effects of labels stem from critical immaturities of cross-modal processing coupled with immaturities of selective attention. Further development of cross-modal processing and the selection-based system, coupled with acquired knowledge that a category label is highly predictive of category membership, may result in category labels becoming category markers in adults (e.g., Yamauchi & Markman, 2000; Yamauchi & Yu, 2008; see also Markman & Ross, 2003). However, additional research is needed to establish a detailed understanding of the changing role of linguistic labels in category representation.

4.3. Summary

In this section, I considered interactions among category structure, the learning system, and characteristics of the learner in category learning and category representation. First, I reviewed evidence demonstrating that dense categories could be learned efficiently by the compression-based system, whereas sparse categories require the involvement of the selection-based system. Second, although the compression-based system exhibits able functioning even early in development, the selection-based system undergoes developmental transformations. As a result, early in development learning subserved by the compression-based system exhibits greater efficiency than learning subserved by the selection-based system. Third, representation of sparse categories changes in the course of development: Although adults form an abstract representation of sparse categories, young children form similarity-based representations of sparse categories. Fourth, there are developmental differences in the representation of dense lexicalized categories: Adults, but not young children, can represent these categories abstractly. And finally, there is evidence that the role of category labels in category representation changes in the course of development; not until late

in development do labels become category markers (although see Waxman & Markow, 1995; Welder & Graham, 2001; Xu, 2002).

5. Conceptual development: From perceptual categories to abstract concepts

On the basis of the formulated characteristics of the input, of the learning systems, and of the learner, we can propose a rough sketch of how conceptual development proceeds. The early functioning of the compression-based system suggests that even young infants should ably learn dense perceptual categories. The ability to learn perceptual categories from relatively dense input has been demonstrated in nonhuman animals as well as in 3- and 4-month-old human infants (Cook & Smith, 2006; Quinn et al., 1993; Smith et al., 2008; Zentall et al., 2008). Although some of these perceptual categories (e.g., cats, dogs, or food) will undergo lexicalization, others (e.g., some categories of speech sounds) will not.

The next critical step is the development of the ability to integrate cross-modal information that may subserve word learning and learning of dense cross-modal categories. There is evidence that very young infants have difficulty in integrating input coming from different modalities, unless both modalities express the same amodal relation (e.g., when the same amodal relation [such as rhythm or rate] is presented cross-modally, cross-modal presentation is likely to facilitate processing of the amodal relation [see Lewkowicz, 2000; Lickliter & Bahrick, 2000, for reviews]). Initially the sensory systems are separated from one another, with multisensory integration being a product of development and learning. There is much recent neuroscience evidence pointing to slow postnatal maturation of multisensory neurons, coupled with slow maturation of functional corticotectal connections (see Wallace, 2004, for a review). Cross-modal integration is at the heart of the ability to learn cross-modal perceptual categories, which permeate early experience (e.g., dogs bark, cats meow, and humans speak).

Once the ability to integrate cross-modal information is somewhat functional, infants can start learning words, which requires binding auditory and visual input. However, given the immaturity of cross-modal processing, it is easier to learn words that denote perceptual categories that the child already knows. Furthermore, infants may spontaneously learn categories of items that are frequent in their environment and these categories would be the first to be labeled by parents. There is evidence (e.g., Nelson, 1973) that the most frequent type of words among the first 100 words produced by babies is a count noun, with most of these count nouns denoting perceptual categories of entities in the child's environment. Therefore, learning the first words could be a way of lexicalizing those perceptual categories that the child already learned. Lexicalization also opens the possibility of acquiring knowledge of unobservable properties about category members, as well as generalizing this knowledge. Unobservable information includes properties that one does not typically observe (e.g., that one's pet dog has a heart) as well as properties that cannot be observed in principle, but have to be inferred from the observed properties (e.g. ''that another person has thoughts and feelings''). Once acquired, these unobservable properties can be entered into the computation of similarity, thus enabling the development of more abstract superordinate categories.

Therefore, lexicalization is a critical step in the transition from perceptual groupings to concepts. The ability to process cross-modal input also enables children to use a combination of perceptual and linguistic cues in acquiring broad ontological distinctions (Jones & Smith, 2002; Samuelson & Smith, 1999; Yoshida & Smith, 2003).

The next important step is learning of dimensional words, denoting dimensional values (e.g., ''green'' or ''square''). Learning of these words coupled with further maturation of the PFC and the development of executive function may result in lexicalization of some stimulus dimensions (such as color, shape, or size). As argued by many researchers (Carey, 1982; Gasser & Smith, 1998), learning of dimensional words follows learning of count nouns. One explanation is that perceptual groupings, such as ''dog'' or ''cup,'' denoted by count nouns are dense—they are based on an intercorrelated set of features and feature dimensions. In contrast, dimensional groupings (e.g., ''red things'') are sparse. Therefore, the later, but not the former, requires selective attention, which appears later in development than the ability to learn perceptual groupings and to integrate cross-modal information.

Further development of the PFC coupled with learning of abstract words lays the foundation for the development of abstract concepts. However, unlike their concrete counterparts (such as ''dog'' or cup'') where category learning may precede word learning, there are reasons to believe that words denoting abstract concepts are learned prior to the concept itself (e.g., Vygotsky, 1964). For example, according to the MacArthur Lexical Development Norms (Dale & Fenson, 1996) a 30-month-old toddler may produce words, such as *love*, *time*, and *same*; however, it is unlikely that these children have concepts of LOVE, TIME, or EQUIVALENCE. Furthermore, because these abstract concepts refer to exceedingly sparse categories, it is likely that the acquisition of these categories requires supervision. The relative maturity of the PFC is of critical importance because learners need to focus on a small set of category-relevant features, while ignoring irrelevant features. The ability to lexicalize categories and the ability to acquire abstract concepts paves the way to acquisition of abstract mathematical and scientific concepts. However, some of these concepts are so sparse and put so much demand on selectivity that supervision alone may not be sufficient—and sophisticated explicit instruction is needed—for successful learning of these concepts (e.g., Kaminski, Sloutsky, & Heckler, 2008).

In sum, the proposal presented here attempts to connect conceptual development with the structure of input and the availability of the learning system necessary for processing of this input. This rough sketch, however, is just a first step in uncovering the great mystery of conceptual development—a progression from a newborn who has difficulty in perceiving the world to an adult who has the ability of changing the world.

6. Concluding comments

In this study, I considered the possibility of conceptual development progressing from simple perceptual grouping to highly abstract scientific concepts. I reviewed evidence suggesting that conceptual development is a product of an interaction of the structure of input,

the category learning system that processes this input, and maturational characteristics of the learner.

I also considered three steps that are critical for conceptual development. First, the development of the selection-based system of category learning that depends critically on the maturation of cortical regions subserving executive function. The second critical step is the ability to integrate cross-modal information. This ability is critical for word learning and lexicalization of spontaneously acquired perceptual groupings, as well as for forming broad ontological classes. And the third critical step, depending on the former two, is the ability to learn and use abstract categories. Unlike their concrete counterparts that can be acquired by perceptual means and lexicalized later, for learning of some abstract categories lexicalization might be a prerequisite.

The proposal presented here considers a complex developmental picture that depends on a combination of maturational and experience factors in conceptual development. Under this view, learning of perceptual categories, cross-modal integration, lexicalization, learning of conceptual properties, the ability to focus and shift attention, and the development of lexicalized concepts are logical steps in conceptual development. This proposal offers a theoretical alternative to the idea of innate knowledge structures specific to various knowledge domains. However, much research is needed to move from a rough sketch to detailed understanding of conceptual development.

Note

1. For the moment, I will ignore a relatively small class of abstract concepts—''electron'' would be a good example—that start out as a lexical entry. However, I will return to this issue later in the study.

Acknowledgments

References

Alvarado, M. C., & Bachevalier, J. (2000). Revisiting the maturation of medial temporal lobe memory functions in primates. *Learning & Memory*, *7*, 244–256.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.

Ashby, F. G., & Maddox, T. W. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology*, *71*, 439–473.

Berg, E. A. (1948). A simple objective test for measuring flexibility and thinking. *Journal of General Psychology*, *39*, 15–22.

Blair, M. R., Watson, M. R., & Meier, K. M. (2009). Errors, efficiency, and the interplay between attention and category learning. *Cognition*, *112*, 330–336.

Brown, R. G., & Marsden, C. D. (1988). Internal versus external cues and the control of attention in Parkinson's disease. *Brain*, *111*, 323–345.

Buffalo, E. A., Ramus, S. J., Clark, R. E., Teng, E., Squire, L. R., & Zola, S. M. (1999). Dissociation between the effects of damage to perirhinal cortex and area TE. *Learning & Memory*, *6*, 572–599.

Bunge, S. A., & Zelazo, P. D. (2006). Brain-based account of the development of rule use in childhood. *Current Directions in Psychological Science*, *15*, 118–121.

Carey, S. (1982). Semantic development: State of the art. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 347–389). Cambridge, England: Cambridge University Press.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Carey, S., & Spelke, E. (1994). Domain specific knowledge and conceptual change. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 169–200). Cambridge, MA: Cambridge University Press.

Carey, S., & Spelke, E. (1996). Science and core knowledge. *Philosophy of Science*, *63*, 515–533.

Caviness, V. S., Kennedy, D. N., Richelme, C., Rademacher, J., & Filipek, P. A. (1996). The human brain age 7–11 years: A volumetric analysis based on magnetic resonance images. *Cerebral Cortex*, *6*, 726–736.

Chater, N., & Christiansen, M. H. (2010). Language acquisition meets language evolution. *Cognitive Science*, *34*(7), 1131–1157.

Chomsky, N. (1980). *Rules and representations*. Oxford, England: Blackwell.

Cincotta, C. M., & Seger, C. A. (2007). Dissociation between striatal regions while learning to categorize via observation and via feedback. *Journal of Cognitive Neuroscience*, *19*, 249–265.

Cook, R. G., & Smith, J. D. (2006). Stages of abstraction and exemplar memorization in pigeon category learning. *Psychological Science*, *17*, 1059–1067.

Cools, A. R., van den Bercken, J. H. L., Horstink, M. W. I., van Spaendonck, K. P. M., & Berger, H. J. C. (1984). Cognitive and motor shifting aptitude disorder in Parkinson's disease. *Journal of Neurology, Neurosurgery and Psychiatry*, *47*, 443–453.

Dale, P. S., & Fenson, L. (1996). Lexical development norms for young children. *Behavior Research Methods, Instruments, & Computers*, *28*, 125–127.

Davidson, M. C., Amso, D., Anderson, L. C., & Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia*, *44*, 2037–2078.

Diamond, A. (2002). Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In D. T. Stuss & R. T. Knight (Eds.), *Principles of frontal lobe function* (pp. 466–503). London, UK: Oxford University Press.

Diamond, A., & Goldman-Rakic, P. S. (1989). Comparison of human infants and rhesus monkeys on Piaget's AB task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research*, *44*, 24–40.

van Domburg, P. H. M. E., & ten Donkelaar, H. J. (1991). *The human substantia nigra and ventral tegmental area*. Berlin: Springer-Verlag.

Fan, J., McCandliss, B. D., Sommer, T., Raz, A., & Posner, M. I. (2002). Testing the efficiency and independence of attentional networks. *Journal of Cognitive Neuroscience*, *14*, 340–347.

Fernandez-Ruiz, J., Wang, J., Aigner, T. G., & Mishkin, M. (2001). Visual habit formation in monkeys with neurotoxic lesions of the ventrocaudal neostriatum. *Proceedings of the National Academy of Sciences*, *98*, 4196–4201.

Fisher, A. V. (2007). Are developmental theories of learning paying attention to attention? *Cognition, Brain, and Behavior*, *11*, 635–646.

Fisher, A. V., & Sloutsky, V. M. (2005). When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. *Child Development*, *76*, 583–597.

French, R. M., Mareschal, D., Mermillod, M., & Quinn, P. C. (2004). The role of bottom-up processing in perceptual categorization by 3- to 4-month-old infants: Simulations and data. *Journal of Experimental Psychology: General*, *133*, 382–397.

Gasser, M., & Smith, L. B. (1998). Learning nouns and adjectives: A connectionist account. *Language and Cognitive Processes*, *13*, 269–306.

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

Gelman, R. (1990). Structural constraints on cognitive development: Introduction to a special issue of Cognitive Science. *Cognitive Science*, *14*, 3–10.

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

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

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj (Ed.), *Language development: Vol. 2. Language, thought and culture* (pp. 301–334). Hillsdale, NJ: Erlbaum.

Giedd, J. N., Snell, J. W., Lange, N., Rajapakse, J. C., Casey, B. J., Kozuch, P. L., Vaituzis, A. C., Vauss, Y. C., Hamburger, S. D., Kaysen, D., & Rapoport, J. L. (1996a). Quantitative magnetic resonance imaging of human brain development: Ages 4–18. *Cerebral Cortex*, *6*, 551–560.

Giedd, J. N., Vaituzis, A. C., Hamburger, S. D., Lange, N., Rajapakse, J. C., Kaysen, D., Vauss, Y. C., & Rapoport, J. L. (1996b). Quantitative MRI of the temporal lobe, amygdala, and hippocampus in normal human development: Ages 4–18 years. *Journal of Comparative Neurology*, *366*, 223–230.

Goldman-Rakic, P. S. (1987). Development of cortical circuitry and cognitive function. *Child Development*, *58*, 601–622.

Golinkoff, R. M., Mervis, C. B., & Hirsh-Pasek, K. (1994). Early object labels: The case for a developmental lexical principles framework. *Journal of Child Language*, *21*, 125–155.

Gureckis, T. M., & Love, B. C. (2004). Common mechanisms in infant and adult category learning. *Infancy*, *5*, 173–198.

Hammer, R., & Diesendruck, G. (2005). The role of dimensional distinctiveness in children's and adults' artifact categorization. *Psychological Science*, *16*, 137–144.

Hammer, R., Diesendruck, G., Weinshall, D., & Hochstein, S. (2009). The development of category learning strategies: What makes the difference? *Cognition*, *112*, 105–119.

Hoffman, A. B., & Rehder, B. (2010). The costs of supervised classification: The effect of learning task on conceptual flexibility. *Journal of Experimental Psychology: General*, *139*, 319–340.

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

Jernigan, T. L., Trauner, D. A., Hesselink, J. R., & Tallal, P. A. (1991). Maturation of the human cerebrum observed in vivo during adolescence. *Brain*, *114*, 2037–2049.

Jones, S. S., & Smith, L. B. (2002). How children know the relevant properties for generalizing object names. *Developmental Science*, *5*, 219–232.

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

Kaminski, J. A., Sloutsky, V. M., & Heckler, A. F. (2008). The advantage of abstract examples in learning math. *Science*, *230*, 454–455.

Keil, F. C. (1979). *Semantic and conceptual development: An ontological perspective*. Cambridge, MA: Harvard University Press.

Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science*, *6*, 449–476.

Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General*, *137*, 52–72.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399–1402.

Koutstaal, W., & Schacter, D. L. (1997). Gist-based false recognition of pictures in older and younger adults. *Journal of Memory & Language*, *37*, 555–583.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.

Kruschke, J. K. (1993). Human category learning: Implications for back propagation models. *Connection Science*, *5*, 3–36.

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*, 812–863.

Lewkowicz, D. J. (2000). Development of intersensory temporal perception: An epigenetic systems/limitations view. *Psychological Bulletin*, *126*, 281–308.

Lickliter, R., & Bahrick, L. E. (2000). The development of infant intersensory perception: Advantages of a comparative convergent-operations approach. *Psychological Bulletin*, *126*, 260–280.

Lombardi, W. J., Andreason, P. J., Sirocco, K. Y., Rio, D. E., Gross, R. E., Umhau, J. C., & Hommer, D. W. (1999). Wisconsin Card Sorting Test performance following head injury: Dorsolateral fronto-striatal circuit activity predicts perseveration. *Journal of Clinical and Experimental Neuropsychology*, *21*, 2–16.

Love, B. C., & Gureckis, T. M. (2007). Models in search of a brain. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 90–108.

Luciana, M., & Nelson, C. A. (1998). The functional emergence of prefrontally-guided working memory systems in four- to eight-year-old children. *Neuropsychologia*, *36*, 273–293.

Macario, J. F. (1991). Young children's use of color in classification: Foods and canonically colored objects. *Cognitive Development*, *6*, 17–46.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298.

Mandler, J. B., Bauer, P. J., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology*, *23*, 263–298.

Mareschal, D., Quinn, P. C., & French, R. M. (2002). Asymmetric interference in 3- to 4-month-olds' sequential category learning. *Cognitive Science*, *26*, 377–389.

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

Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592–613.

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, *13*, 329–358.

Marks, W. (1991). Effects of encoding the perceptual features of pictures on memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 566–577.

Napolitano, A. C., & Sloutsky, V. M. (2004). Is a picture worth a thousand words? The flexible nature of modality dominance in young children. *Child Development*, *75*, 1850–1870.

Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development*, *38*(1–2), 1–137.

Nomura, E. M., Maddox, W. T., Filoteo, J. V., Ing, A. D., Gitelman, D. R., Parrish, T. B., Mesulam, M. M., & Reber, P. J. (2007). Neural correlates of rule-based and information-integration visual category learning. *Cerebral Cortex*, *17*, 37–43.

Nomura, E. M., & Reber, P. J. (2008). A review of medial temporal lobe and caudate contributions to visual category learning. *Neuroscience and Biobehavioral Reviews*, *32*, 279–291.

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

Opfer, J. E., & Bulloch, M. J. (2007). Causal relations drive young children's induction, naming, and categorization. *Cognition*, *105*, 206–217.

Opfer, J. E., & Siegler, R. S. (2004). Revisiting preschoolers' living things concept: A microgenetic analysis of conceptual change in basic biology. *Cognitive Psychology*, *49*, 301–332.

Pfefferbaum, A., Mathalon, D. H., Sullivan, E. V., Rawles, J. M., Zipursky, R. B., & Lim, K. O. (1994). A quantitative magnetic resonance imaging study of changes in brain morphology from infancy to late adulthood. *Archives of Neurology*, *51*, 874–887.

Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42.

Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, *22*, 463–475.

Rakison, D. H., & Poulin-Dubois, D. (2001). Developmental origin of the animate-inanimate distinction. *Psychological Bulletin*, *127*, 209–228.

Rao, S. M., Bobholz, J. A., Hammeke, T. A., Rosen, A. C., Woodley, S. J., Cunningham, J. M., Cox, R. W., Stein, E. A., & Binder, J. R. (1997). Functional MRI evidence for subcortical participation in conceptual reasoning skills. *NeuroReport*, *8*, 1987–1993.

Robinson, C. W., & Sloutsky, V. M. (2004). Auditory dominance and its change in the course of development. *Child Development*, *75*, 1387–1401.

Robinson, C. W., & Sloutsky, V. M. (2007a). Linguistic labels and categorization in infancy: Do labels facilitate or hinder? *Infancy*, *11*, 233–253.

Robinson, C. W., & Sloutsky, V. M. (2007b). Visual processing speed: Effects of auditory input on visual processing. *Developmental Science*, *10*, 734–740.

Robinson, C. W., & Sloutsky, V. M. (2008). Effects of auditory input in individuation tasks. *Developmental Science*, *11*, 869–881.

Robinson, C. W., & Sloutsky, V. M. (2010). Development of cross-modal processing. *WIRES: Cognitive Science*, *1*, 1–7.

Rodman, H. R. (1994). Development of inferior temporal cortex in the monkey. *Cerebral Cortex*, *4*, 484–498.

Rodman, H. R., Skelly, J. P., & Gross, C. G. (1991). Stimulus selectivity and state dependence of activity in inferior temporal cortex of infant monkeys. *Proceedings of the National Academy of Sciences*, *88*, 7572–7575.

Rogers, R. D., Andrews, T. C., Grasby, P. M., Brooks, D. J., & Robbins, T. W. (2000). Contrasting cortical and subcortical activations produced by attentional-set shifting and reversal learning in humans. *Journal of Cognitive Neuroscience*, *12*, 142–162.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.

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

Rueda, M., Fan, J., McCandliss, B. D., Halparin, J., Gruber, D., Lercari, L., & Posner, M. I. (2004). Development of attentional networks in childhood. *Neuropsychologia*, *42*, 1029–1040.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*, 27–52.

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

Seger, C. A. (2008). How do the basal ganglia contribute to categorization? Their roles in generalization, response selection, and learning via feedback. *Neuroscience and Biobehavioral Reviews*, *32*, 265–278.

Seger, C. A., & Cincotta, C. M. (2002). Striatal activation in concept learning. *Cognitive, Affective, & Behavioral Neuroscience*, *2*, 149–161.

Shannon, C. E., & Weaver, W. (1948). *The mathematical theory of communication*. Chicago: University of Illinois Press.

Shepp, B. E., & Swartz, K. B. (1976). Selective attention and the processing of integral and nonintegral dimensions: A developmental study. *Journal of Experimental Child Psychology*, *22*, 73–85.

Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*, *7*, 246–251.

Sloutsky, V. M., & Fisher, A. V. (2004a). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, *133*, 166–188.

Sloutsky, V. M., & Fisher, A. V. (2004b). When development and learning decrease memory: Evidence against category-based induction in children. *Psychological Science*, *15*, 553–558.

Sloutsky, V. M., & Fisher, A. V. (2008). Attentional learning and flexible induction: How mundane mechanisms give rise to smart behaviors. *Child Development*, *79*, 639–651.

Sloutsky, V. M., Kloos, H., & Fisher, A. V. (2007). When looks are everything: Appearance similarity versus kind information in early induction. *Psychological Science*, *18*, 179–185.

Sloutsky, V. M., & Lo, Y.-F. (1999). How much does a shared name make things similar? Part 1: Linguistic labels and the development of similarity judgment. *Developmental Psychology*, *35*, 1478–1492.

Sloutsky, V. M., Lo, Y.-F., & Fisher, A. (2001). How much does a shared name make things similar? Linguistic labels, similarity and the development of inductive inference. *Child Development*, *72*, 1695–1709.

Sloutsky, V. M., & Napolitano, A. C. (2003). Is a picture worth a thousand words? Preference for auditory modality in young children. *Child Development*, *74*, 822–833.

Sloutsky, V. M., & Robinson, C. W. (2008). The role of words and sounds in visual processing: From overshadowing to attentional tuning. *Cognitive Science*, *32*, 354–377.

Sloutsky, V. M., & Spino, M. A. (2004). Naive theory and transfer of learning: When less is more and more is less. *Psychonomic Bulletin and Review*, *11*, 528–535.

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

Smith, L. B., Jones, S. S., & Landau, B. (1996). Naming in young children: A dumb attentional mechanism? *Cognition*, *60*, 143–171.

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

Smith, J. D., Redford, J. S., & Haas, S. M. (2008). Prototype abstraction by monkeys (*Macaca mulatta*). *Journal of Experimental Psychology: General*, *137*, 390–401.

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

Sowell, E. R., & Jernigan, T. L. (1999). Further MRI evidence of late brain maturation: Limbic volume increases and changing asymmetries during childhood and adolescence. *Developmental Neuropsychology*, *14*, 599–617.

Sowell, E. R., Thompson, P. M., Holmes, C. J., Batth, R., Jernigan, T. L., & Toga, A. W. (1999). Localizing age-related changes in brain structure between childhood and adolescence using statistical parametric mapping. *NeuroImage*, *9*, 587–597.

Sowell, E. R., Thompson, P. M., Holmes, C. J., Jernigan, T. L., & Toga, A. W. (1999). In vivo evidence for post-adolescent brain maturation in frontal and striatal regions. *Nature Neuroscience*, *2*, 859–861.

Spelke, E. S. (2000). Core knowledge. *American Psychologist*, *55*, 1233–1243.

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*, 89–96.

Striedter, G. F. (2005). *Principles of brain evolution*. Sunderland, MA: Sinauer.

Teng, E., Stefanacci, L., Squire, L. R., & Zola, S. M. (2000). Contrasting effects on discrimination learning after hippocampal lesions and conjoint hippocampal-caudate lesions in monkeys. *Journal of Neuroscience*, *20*, 3853–3863.

Thapar, A., & McDermott, K. B. (2001). False recall and false recognition induced by presentation of associated words: Effects of retention interval and level of processing. *Memory and Cognition*, *29*, 424–432.

Tipper, S. P., & Driver, J. (2000). Negative priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. In M. S. Gazzaniga (Ed.), *Cognitive neuroscience: A reader* (pp. 176–187). Malden, MA: Blackwell Publishing.

Twyan, A. D., & Newcombe, N. S. (2010). Five reasons to doubt the existence of a geometric module. *Cognitive Science*, *34*(7), 1315–1356.

Vygotsky, L. S. (1964). *Thought and language*. Cambridge, MA: MIT Press. (Original work published in 1934)

Wallace, M. T. (2004). The development of multisensory processes. *Cognitive Processing*, *5*, 69–83.

Wasserman, E. A., Young, M. E., & Cook, R. G. (2004). Variability discrimination in humans and animals: Implications for adaptive action. *American Psychologist*, *59*, 879–890.

Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12-13-month-old infants. *Cognitive Psychology*, *29*, 257–302.

Welder, A. N., & Graham, S. A. (2001). The influence of shape similarity and shared labels on infants' inductive inferences about nonobvious object properties. *Child Development*, *72*, 1653–1673.

Whitman, J. R., & Garner, W. R. (1962). Free recall learning of visual figures as a function of form of internal structure. *Journal of Experimental Psychology*, *64*, 558–564.

Wilson, C. (1995). *The contribution of cortical neurons to the firing pattern of striatal spiny neurons*. Cambridge, MA: Bradford.

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, *85*, 223–250.

Yamauchi, T., Kohn, N., & Yu, N.-Y. (2007). Tracking mouse movement in feature inference: Category labels are different from feature labels. *Memory & Cognition*, *35*, 852–863.

Yamauchi, T., Love, B. C., & Markman, A. B. (2002). Learning nonlinearly separable categories by inference and classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 585–593.

Yamauchi, T., & Markman, A. B. (1998). Category learning by inference and classification. *Journal of Memory and Language*, *39*, 124–148.

Yamauchi, T., & Markman, A. B. (2000). Inference using categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 776–795.

Yamauchi, T., & Yu, N.-Y. (2008). Category labels versus feature labels: Category labels polarize inferential predictions. *Memory & Cognition*, *36*, 544–553.

Yoshida, H., & Smith, L. B. (2003). Shifting ontological boundaries: How Japanese- and English speaking children generalize names for animals and artifacts. *Developmental Science*, *6*, 1–34.

Zelazo, P. D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development*, *11*, 37–63.

Zelazo, P. D., Muller, U., Frye, D., & Marcovitch, S. (2003). The development of executive function in early childhood. *Monographs of the Society for Research on Child Development*, *68*, vii–137.

Zentall, T. R., Wasserman, E. A., Lazareva, O. F., Thompson, R. K. R., & Rattermann, M. J. (2008). Concept learning in animals. *Comparative Cognition & Behavior Reviews*, *3*, 13–45.