CHAPTER 14

Similarity, Induction, Naming and Categorization: A Bottom-up Approach

Vladimir M. Sloutsky

People are remarkably smart: they know language, possess complex skills, and function in a complex dynamic environment. However, they do not exhibit evidence of this knowledge at birth and one of the central and most interesting issues in the study of human cognition is the question of how people acquire this knowledge in the course of development. One possibility is that sensory and cognitive systems have the capacity to extract regularities from structured input. However, this possibility has been frequently criticized as an untenable candidate for a complete story. In particular, it has been often argued that input data are severely underdetermined: the same data are compatible with multiple possibilities. Consider a visual occlusion phenomenon presented in Figure 14.1: the same event presented in Figure 14.1A is compatible with events presented in Figures 14.1B, C, D, and E, or an infinite number of other possibilities, and there is not enough information in the input to decide among these possibilities. However, under many conditions, even very early in development, people are more likely to expect the situation presented in Figure 14.1B than situations presented in Figure 14.1C and 14.1D (Kellman & Spelke, 1983; Spelke, 1990, but see Johnson, 2003).

Figure 14.1 about here

Similarly, when adults label objects for young children (e.g., “Look, look, a dax!”), a mere co-occurrence of an object and a word is not sufficient for determining that the word refers to the object. In fact, the word could refer to a variety of things, including, any part of the object, its color, texture, manner of motion, or the speaker's attitude toward the situation. Again, despite the fact that the referential situation appears to be severely under constrained, even very young children somehow manage to determine that the word (at least if it is a noun) refers to the object (see Woodward & Markman, 1998, for a review).

It has been argued that these examples indicate that infants and young children perform these tasks with some expectations about the language and the world (e.g., they may expect words to refer to objects). This hypothesis raises a number of important questions. Where does this knowledge come from? How is this knowledge represented? And how do infants and young children deploy this knowledge? If the input is under constrained, then the "cognitive" constraints imposed on the input cannot be acquired from the input, and hence these constraints have to be a priori. Therefore, according to this position, some aspects of structured knowledge (e.g., knowledge of syntax, semantics, or visual organization) are a priori and the task of learning is to recognize these structures in the input. To reflect the fact that the constraints are a precondition rather than a consequence of learning, I will refer to this position as a "top-down approach." "Top-down" here refers to the origins of knowledge and many constraints that emerge in the bottom-up manner (in the course of learning) do not have "top-down" origins.

However, there is a growing body of evidence indicating that input is substantially richer than the above examples imply and that human learning is powerful enough to extract regularities from structured, albeit somewhat noisy, input. Of course, if this is the case, then people may not need a priori constraints. To reflect the fact that regularities are extracted from data (rather than imposed on the data), I will refer to this position as a "bottom-up" approach. The primary focus of this chapter is whether learning and cognitive development are constrained by a priori knowledge (and are thus top-down processes) or by the nature of the input and by learning mechanisms (and are thus bottom-up processes).

TOP-DOWN VERSUS BOTTOM-UP LEARNING

We say that a process is *top-down* when an outcome of a process is determined by a higher-order structure. For example, the number of light bulbs produced in North Korea in a given year (as well as their price) is completely determined by a government plan. The plan is not influenced by market conditions, but is designed to affect these conditions. Therefore, the plan, rather than the market conditions, prescribes the manufacturers' output, and radically different market conditions may result in a similar or identical output. Similarly, if cognitive development is a top-down process, then acquisition of syntax, the development of perception, and the development of generalization (such as word leaning or categorization) are constrained by knowledge, beliefs, or biases that are independent of input. As a result, radically different input conditions should result in similar linguistic and perceptual development outcomes.

Alternatively, we say that the process is *bottom-up* when an outcome of the process stems from an interaction of simpler processes, none of which alone can determine the outcome. In this case, the regularity in the outcome emerges from regularities in the input. For example, although it cannot be known in advance how many light bulbs will be produced in the United States in a given year, the production (as well as the price) is a function of market conditions (e.g., the cost of energy, the amount of construction activity, etc.), and it is likely that similar market conditions will result in somewhat similar production outcomes. Similarly, acquisition of syntax, the development of perception, and the development of generalization is a joint function of the regular input and the ability of the organism to extract and exploit these regularities.

The distinction between top-down and bottom-up approaches to learning has been recurring under different names. For example, in the domain of language learning, there has been a distinction between inductive and deductive approaches to language learning (e.g., Wexler & Culicover, 1980; Xu & Tenenbaum, 2007). According to the former view, learners extract regularities from linguistic input. In contrast, according to the latter view, learners entertain a fixed set of a priori hypotheses, and the role of input is to eliminate the inappropriate hypotheses, thus enabling the learner to settle on the correct hypothesis.

In this chapter, I focus on a different domain—the ability to extend knowledge from known to novel. People deploy this ability every time they (1) extend a known word to a novel entity (i.e., in naming or label extension), (2) treat a novel entity as a member of a familiar class (i.e., categorization), and (3) extend a property from a familiar entity to a novel entity (i.e., projective induction).

This ability to extend knowledge from known to novel, or inductive generalization, and the development of this ability are interesting and controversial issues in human cognition. Some researchers propose a top-down approach to generalization, arguing that even early in development, conceptual knowledge is a critical component of inductive generalization (Gelman, 2003; Gelman & Wellman, 1991; Keil, Smith, Simons, & Levin, 1998). This position is commonly known as knowledge-based (or naïve theory) approach. Others suggest that inductive generalization may develop in a bottom-up manner, with similarity playing a critical role in early induction (Colunga & Smith, 2005; French, Mareschal, Mermillod, & Quinn, 2004; Rogers & McClelland, 2004; Sloutsky, 2003; Sloutsky & Fisher, 2004a). While both positions agree that bottom-up processes, such as similarity computation, play a role (e.g., Keil et al., 1998; Keil, 2003), there is less agreement about the role and origins of top-down knowledge. For example, there no agreement on whether these top-down constraints are necessary for inductive generalization, when these constraints come online, and where they come from.

TOP-DOWN APPROACHES TO GENERALIZATION

According to this position, even at the outset of development, infants and young children have a repertoire of "smart" conceptual assumptions about the language and the world (R. Gelman, 1990; Gelman & Markman, 1987; Gelman & Wellman, 1991; Keil, 1989; Mandler, 1997, 2004; Soja, Carey, & Spelke, 1991; Wellman & Gelman, 1992). Although conceptual knowledge plays a prominent role in top-down approaches, it is hard to pin down what exactly the top-down theorists mean by “conceptual knowledge” is. One possibility is that conceptual knowledge is knowledge that cannot be observed directly, but has to be inferred. Therefore, knowledge that birds have wings is perceptual, whereas knowledge that birds and fish share many biological properties is conceptual. Another issue that remains unclear is the origin of conceptual knowledge. Does conceptual knowledge emerge as generalization over data? Or does it exist independently of data? And if the latter is the case, where does it come from? Although these questions remain unanswered, conceptual knowledge plays a prominent role in top-down theories of generalization.

According to some accounts (see Massey & R. Gelman, 1988; Spelke, 1994) conceptual knowledge constrains (or even overrides) perceptual input, while remaining impervious to perceptual input. Conceptual effects in an inductive generalization task have been described as follows. When “trying to determine whether to draw an inference from object A to object B, a child would not simply calculate the similarity between the two objects. Rather the child would determine whether A and B belong to members of the same natural kind category that encompasses both A and B” (Gelman & Coley, 1991, p. 185). Therefore, according to this account of generalization, abstract (and not directly perceptible) category information is of greater importance than appearance information. Furthermore, identification of an abstract category is a necessary step in inductive generalization, and therefore induction in young children is category-based (e.g., Gelman, 1988).

Overall, among the putative conceptual assumptions, two are especially important for inductive generalization. The category assumption is the belief that individual entities are members of more general categories and that members of the same category share many unobserved properties. The linguistic assumption is the belief that words (especially count nouns) denote categories rather than individuals. When performing induction, people rely on these assumptions to conclude that entities sharing a label belong to the same kind, and therefore share many unobservable properties. For instance, when shown a picture of a yellow fish and told that this fish needs branchia to breathe, children are more likely to generalize this property to a red fish than to a turtle or a couch (Gelman, 1988). Presumably, children apply the two assumptions to infer that since the entities have a matching label (i.e., both are referred to as "a fish"), these entities belong to the same kind, and therefore they share many important properties.

Supporting Evidence

There are several lines of research supporting the idea of the linguistic and category assumptions in infants and children. For example, in support of the linguistic assumption, Markman and Hutchinson (1984) demonstrated that in an absence of a label, children may group things thematically (e.g., a police car and a policeman), whereas when the police car was named "a dax" and children were asked to select another "dax," they selected the passenger car, thus grouping the cars together. More recently, Gelman and Heyman (1999) demonstrated that young children were more willing to generalize properties from one person to another when both persons were referred to by a noun (i.e., “carrot-eaters”) than when both were referred to by a descriptive sentence (e.g., “both like to eat carrots”). Furthermore, infants and children may also expect words to refer to categories, although, younger infants may hold this expectation not only for count nouns, but also for other speech sounds (e.g., Balaban & Waxman, 1997; Booth & Waxman, 2002).

There is also evidence supporting the idea that young children expect things to belong to categories and they expect members of the same category to share important properties. This evidence stems primarily from the study of inductive generalization. For example, Gelman and Markman (1986) presented young children with a Target item and two Test items, with one Test item looking more like the Target and another Test item sharing the label with the Target. Participants were also told that one Test item had a hidden property (e.g., “This one has hollow bones”), whereas another Test item had a different hidden property (e.g., “This one has solid bones”), and asked to induce a hidden property to the Target. Results indicated that young children tended to induce properties from the identically labeled, but not from the similarly looking item. Similar findings were reported with infants. Specifically, when objects were not labeled, infants generalized nonobvious properties based on appearance similarity, whereas when objects were labeled participants were using labeling information (Welder & Graham, 2001).

Challenges to the Top-Down Approach

There are several challenges to the idea that early generalization is driven by top-down conceptual assumptions. Some of these challenges stem from findings indicating that conceptual knowledge is not an a priori constraint, but is rather a product of learning and development.

Other challenges stem from findings questioning the very existence of particular conceptual assumptions early in development.

Most importantly, empirical evidence does not lend unequivocal support to the idea that children's inductive generalization is driven by a priori assumptions, such as the linguistic and the category assumptions. In particular, several lines of research on categorization, word learning, and projective induction cast doubt on the existence of a priori assumptions in infants and young children (Colunga & Smith, 2005; French et al., 2004; Oakes & Madole, 2003, Rakison, 2003, Smith, Jones, & Landau, 1996; Yoshida & Smith, 2005). For example, it has been demonstrated that category learning in early infancy can be readily accounted for by the distribution of feature values within a category and across categories (French et al., 2004). It has been also found that when categorizing objects or inducing properties, the type of property infants focus on changes in the course of development (e.g., Madole, Oakes, & Cohen, 1993; Rakison, 2003). Similarly, "biases" that drive word learning have been found to come on-line in the course of learning rather than being a priori. For example, Smith, and colleagues (e.g., Smith, Jones, Landau, Gershkoff-Stowe, & Samuelson, 2002) examined the developmental course of the "shape bias"—the tendency of toddlers to extend a novel word to items that have the same shape as the originally labeled item. These researchers hypothesized that the shape bias is not a prerequisite, but a consequence of word leaning. The underlying idea is that in the course of language acquisition, babies first detect that a particular shape co-occurs with a particular label (e.g., "cup-shaped" objects are called "cups," while "ball-shaped" objects are called "balls"). As the corpus of evidence grows, they detect a more abstract regularity: similar-shaped entities have the same name, and this more abstract regularity, in turn, makes shape an important predictor of category membership. If this hypothesis is correct, then the shape bias is a function of word learning experience and it should emerge as children accumulate their vocabularies. This is exactly what Smith and colleagues found, thus presenting evidence that the “shape bias” in naming is not an a priori assumption, but is a product of learning.

Additional challenges to the top-down approach come from studies with preschoolers (i.e., 3- to 5-year-old children). For example, there are recent findings indicating that effects of labels on categorization and induction may be driven by attentional factors and do not have to stem from the linguistic assumption (Napolitano & Sloutsky, 2004; Sloutsky & Napolitano 2003). In particular, Amanda Napolitano and I demonstrated that under many conditions, auditory input (including linguistic input) automatically captures young children’s attention, thus overshadowing (or attenuating processing of) corresponding visual input. Therefore, it is possible that (due to overshadowing effects) shared labels contribute to the overall similarity of compared entities (Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999), and thus to both categorization and induction.

Furthermore, there are reasons to be skeptical whether some of the findings, which are being used as evidence for linguistic and category assumptions, do in fact constitute such evidence. For example, Anna Fisher and I (Sloutsky & Fisher, 2004a) reexamined young children’s performance on Gelman and Markman (1986) task. Similar to the original task, we presented 4- to 5-year-olds with triads of items, with each triad consisting of a target and two test items (these are presented in Figure 14.2A). Also, as in the original task, we named items, such that one test item shared the label with the target, whereas the other appeared more similar to the target. Finally, as in the original task, we told participants that each of the test items had a particular nonobvious property and asked them to guess the property of the target. The entire procedure, including pictures, labels, and properties were identical to the original study. In addition, prior to the experiment proper, we conducted a separate experiment estimating for each triad the similarity of the test items to the target.

Figure 14.2 about here

While replicating the overall mean reported by Gelman and Markman, we demonstrated that (a) children's performance varied drastically across picture triads (which should not be the case if induction was driven by a set of assumptions), (b) appearance similarity made a sizable contribution to induction, and (c) a simple model of similarity (the model is presented below) very accurately predicted how children perform on individual triads (see Figure 14.2B for a comparison of predicted and observed means). Therefore, a detailed analysis of children's performance does not support the conclusion that in the presence of shared labels, children ignore appearance information.

Another serious challenge to the idea that early generalization is constrained by a priori conceptual knowledge is that it is unclear how this knowledge (even if it existed) could be implemented and deployed in real-life situations. For example, in order for infants and young children to ignore salient (yet surface-level) perceptual information in favor of nonobservable and thus less salient conceptual information, they should have a substantial level of control over their attention. However, there is little evidence indicating that young children have sufficient control of attention, enabling them to focus on less salient (yet deep) information, while ignoring more salient (yet surface) information.

BOTTOM-UP APPROACHES TO GENERALIZATION

According to this view, young children generalize on the basis of multiple commonalities, or similarities, among presented entities. Because members of the same category often happen to be more perceptually similar to each other than they are to nonmembers (i.e., a yellow fish is more similar to a red fish than it is to a turtle or a couch), children are more likely to generalize properties to members of a category than to nonmembers. Furthermore, common labels could be features directly contributing to perceptual similarity rather than denoting a common category (Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999). Proponents of this view challenge the position that young children hold conceptual assumptions, and they argue that induction with both familiar and novel categories is similarity-based. In the remainder of this section, I will consider a recently proposed model of generalization (SINC standing for Similarity, Induction, Naming, and Categorization) that considers categorization, induction, and naming as variants of similarity-based generalization (see Sloutsky & Fisher, 2004a).

SINC Model

Overview

SINC assumes that young children consider linguistic labels as attributes of objects that contribute to similarity among compared entities. This assumption has been supported empirically (Sloutsky & Lo, 1999; Sloutsky, Lo, & Fisher, 2001). Qualitatively, SINC suggests that linguistic labels contribute to similarity of compared entities and that similarity drives induction and categorization in young children. The model is based on the product-rule model of similarity (Estes, 1994; Medin, 1975) that specifies similarity among nonlabeled feature patterns. In the product-rule model, similarity is computed using the following equation:

$$\mathrm{Sim}(i, j) = S^{N-k} \qquad (14.1)$$

where N denotes the total number of relevant attributes, k denotes the number of matches, and S ($0 \le S \le 1$) denotes values (weights) of a mismatch. For example, suppose that one is presented with two visual patterns (e.g., schematic faces A and B). Further suppose that these patterns consist of four distinct features (i.e., the shape of the face, eyes, and nose, and the size of ears) that the patterns share two of these features (i.e., the shape of the face and eyes) and differ on the other two. Assuming that S = 0.5 [the value frequently derived empirically (Estes, 1994)], similarity between A and B would be equal to 0.25 (i.e., 0.5^2). Note that similarity between entities decreases very rapidly with a decrease in the number of mismatches, approximating the exponential decay function discussed elsewhere (Nosofsky, 1984). For example, if the faces shared only one of the four features, their similarity would be equal to 0.125 (i.e., 0.5^3). On the other hand, if the faces shared all four features, they would be identical, and their similarity would be equal to 1 (i.e., 0.5^0).

According to SINC, similarity of labeled feature patterns could be calculated using the following equation:

$$\mathrm{Sim}(i, j) = W_{\mathrm{Label}}^{1-L} S_{Vis.attr}^{N-k} \begin{Bmatrix} L = 1, \text{if } L_i = L_j \\ L = 1, \text{ otherwise} \end{Bmatrix} \quad (14.2)$$

Again, N denotes the total number of visual attributes, k denotes the number of matches, $S_{\mathrm{vis.attr.}}$ denotes values (attentional weights) of a mismatch on a visual attribute, W_{Label} denotes values of label mismatches, and L denotes a label match.

When there is a label match, L = 1 and W_{Label} = 1; when there is a label mismatch, L = 0 and W_{Label} < 1. Note that S and W (their values vary between 0 and 1) denote attentional weights of mismatches and the contribution of S and W is large when these parameters are close to 0 and is small they are close to 1. This is because the closer the value of these parameters to 1, the smaller the contribution of a mismatch to the detection of difference, while the closer the value to 0, the greater the contribution to the detection of difference. When two entities are identical on all dimensions (i.e., there are no mismatches), their similarity should be equal to 1; otherwise, it is smaller than 1.

Note that, according to the model, when neither entity is labeled (i.e., $W_{\text{Label}} = 1$), similarity between the entities is determined by the number of overlapping visual attributes, thus conforming to Equation 14.1. Labels are presented as a separate term in the equation because they are expected to have larger attentional weight than most visual attributes (Sloutsky & Lo, 1999). In the case that the weight of a label does not differ from that of other attributes, the label will become one of the attributes in the computation of similarity, and Equation 14.2 becomes Equation 14.1.

Why would labels contribute to similarity? And what might be a mechanism underlying the greater weight of labels at earlier age demonstrated in previous research (e.g., Sloutsky & Lo, 1999)? One possibility is that labels have larger weights because they are presented auditorily, and auditory processing dominates the visual processing in infancy and early childhood, but this dominance decreases with age (Lewkowicz, 1988a, 1988b; Sloutsky & Napolitano, 2003). Alternatively, it is possible that larger weights of labels are grounded in a special status of sounds of human speech (Balaban & Waxman, 1997; Waxman & Markow, 1995). We discuss both possibilities in the "Why Do Labels Contribute to Similarity" section.

Finally, SINC suggests that if the child is presented with a Target feature pattern (T) and Test feature patterns (A and B) and asked which of the Test patterns is more similar to the Target, the child's choice for one of the Test items (e.g., Test B) could be predicted using a variant of the Luce's choice rule presented in the following equation:

$$P(B) = \frac{\text{Sim}(T,B)}{\text{Sim}(T,B) + \text{Sim}(T,A)} \qquad (14.3)$$

We argue that if induction and categorization in young children are indeed similarity-based, then this model that predicts similarity judgment in young children (e.g., Sloutsky & Lo, 1999) should be able to predict their induction.

However, for the majority of naturalistic visual stimuli patterns, it is impossible to individuate features and calculate feature overlap (e.g., think about photographs of two animals and the multiplicity of perceptual features that they have). At the same time, perceptually rich naturalistic stimuli constitute the most interesting and informative test of the proposed model. Because neither *N* nor *k* presented in Equation 14.1 are determinable a priori for perceptually rich naturalistic stimuli, we made several additional steps to apply the model to naturalistic stimuli. Denoting similarity of Test stimuli A and B to the Target as S^x and S^y, respectively, and performing simple derivations from Equation 14.3 allow us to get equations predicting categorization and induction performance. First, consider the case when entities are not labeled. Substituting Sim(*T*,*B*) and Sim(*T*,*A*) by S^x and S^y, we get the following equation:

$$P(B) = \frac{S^x}{S^x + S^y} = \frac{S^x}{S^x\left(1 + S^{y-x}\right)} = \frac{1}{1 + \frac{S^y}{S^x}} \qquad (14.4)$$

For the labeled entities, derivations remain essentially the same, except for the W_{Label} parameter. The parameter equals to 1, if there is a label match, otherwise it varies from 0 to 1, and the smaller the value of *W*, the greater the contribution of label mismatch. Therefore, in the case of labeled entities, the probability of selecting the item that shares the same label (say item B if it shares the label with the Target) could be derived as follows:

$$P(B) = \frac{S^x}{S^x + WS^y} = \frac{S^x}{S^x\left(1 + WS^{y-x}\right)} = \frac{1}{1 + \frac{WS^y}{S^x}} \qquad (14.5)$$

In short, Equations 14.4 and 14.5 should predict participants' induction responses in label and no-label conditions, respectively. In other words, their willingness to induce from Test B to the Target should be a function of the ratio of S^y/S^x (i.e., of similarity of A and B to the Target) when no labels are provided, and it should be a joint function of S^y/S^x and W (i.e., the attentional weight of label) when labels are provided. Note that Equation 14.5 reflects a situation when the Target and Test B have the same labels, whereas Test A has a different label. For the purpose of expository convenience, in the description of data and in figures, I will refer to the Test stimulus sharing the label with the Target as "Test B." Note that both W and S^y/S^x can be estimated from data, and therefore Equations 14.4 and 14.5 can be used for predicting specific probabilities of induction and categorization.

One important (and testable) consequence of this proposal is that because linguistic labels contribute to similarity in a quantitative manner rather than in a qualitative "all-or-nothing" manner, they should also make a quantitative contribution to induction as well. Therefore, the top-down approach and SINC make different predictions about the effect of linguistic labels on induction. If inductive generalization is made solely on the basis of linguistic label (as predicted by the linguistic assumption), then induction should be independent of appearance similarity. Alternatively, SINC predicts that labels make a quantitative contribution to similarity and thus to induction.

Empirical Findings Generated by SINC

SINC is a model implementing a theory of inductive generalization, which can predict a wide range of phenomena across a variety of tasks. These phenomena include: (1) effects of labels on similarity early in development; (2) effects of phonological similarity of labels on induction; (3) low-level attentional mechanisms underlying effects of labels on similarity and induction; (4)

flexible (yet nondeliberate) adjustment of attentional weights of different sources of information; (5) differential effects of induction on recognition memory at different points of development; (6) dissociation between label and category information; (7) integration of labeling and appearance information in the course of similarity judgment and induction; and (8) interrelationships among induction, categorization, naming, and similarity.

First, it has been demonstrated that early in development, labels contribute to similarity of compared entities (Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999). In particular, when two entities share a label, young children tend to consider these entities as looking more alike than when the same entities are presented without labels, and, as we explain below, these effects stem from attentional factors, such as auditory information overshadowing (or attenuating processing) of corresponding visual information (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003).

Second, if labels are features contributing to the overall similarity, then it is possible that labels are perceived as subjectively continuous variables, in which case, not only the identity, but also phonological similarity of labels would contribute to the overall similarity and to induction. Therefore, the theory underlying SINC suggests that phonological similarity of labels may contribute to induction. There is recent evidence supporting this prediction (Fisher & Sloutsky, 2004): Young children were more likely to generalize a property from a test item to a target item if the test and target items had a phonologically similar label (e.g., *Guma* and *Gama*) than if it had a phonologically different label (*Guma* and *Fika*). Furthermore, young children tended to extend phonologically similar words to visually similar entities.

Third, according to SINC, effects of words on similarity and thus on induction stem from low-level attentional mechanisms rather than from an understanding of the conceptual importance of labels. Based on this idea, it was predicted that words, as well as other auditory stimuli, may affect processing of corresponding visual information. This prediction was confirmed empirically: For infants and young children, auditory information overshadows corresponding visual information (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003). In particular, when discriminable visual and auditory stimuli (including human speech) were presented together, discrimination of visual (but not of auditory) stimuli decreased compared to a unimodal baseline.

Fourth, because effects of words and visual information on induction stem from attentional mechanisms rather than from understanding of the conceptual importance of labels or appearances, SINC predicts that contribution of labels or appearance to induction can be changed by changing attentional weights of labels or appearance through associative training. This prediction was supported in a set of experiments (Fisher & Sloutsky, 2006; Sloutsky & Spino, 2004). In one set of experiments (Sloutsky & Spino, 2004), attention to labels or to appearances was manipulated by varying their predictive values (when a cue is consistently nonpredictive, attention to this cue decreases, see Hall, 1991, for a review). After training, participants were presented with an induction task, which was repeated again 3.5 months after training. It was found that as a result of training, young children exhibited (depending on the training condition) either appearance-based or label-based induction, with either pattern being different from pretraining induction. Furthermore, 3.5 months after training, young children retained these effects of training.

Fifth, SINC enabled a novel prediction regarding effects of induction on recognition memory. Recall that according to SINC, early induction is driven by similarity, whereas according to the knowledge-based approach, even early in development induction is category-based and is driven by more abstract category information. To address this issue, Anna Fisher and I developed the Induction-then-Recognition (ITR) paradigm, allowing the distinction between these two possibilities (see Sloutsky & Fisher, 2004a, 2004b).

The ability of ITR to distinguish between these possibilities is based on the following reasoning. Research on false-memory phenomena showed 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., Koutstall & 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, a memory test administered after an induction task may reveal information about how items were processed during the induction task. If participants processed the items relatively abstractly as members of a category (i.e., they performed category-based induction), then they would likely have difficulty discriminating studied targets from conceptually similar critical lures. If, on the other hand, they processed items more concretely, focusing on perceptual details (i.e., they performed similarity-based induction), then they should discriminate relatively well. This is exactly what we found: After performing induction with pictures 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, 2005a; Sloutsky & Fisher, 2004a, 2005). Their accuracy, however, dropped to the level of adults when they were trained to perform induction in a category-based manner—by deciding whether the test and the target items belong to the same kind.

Unlike some top-down approaches that consider linguistic labels as category proxies, SINC clearly differentiates between labeling information and category information: According to SINC, linguistic labels contribute to induction by contributing to similarity rather than by denoting categories. Initial evidence for this distinction comes from the fact that young children's induction is not category-based (Fisher & Sloutsky, 2005a; Sloutsky & Fisher, 2004a, 2004b), yet labels contribute to early induction by affecting similarity (Sloutsky & Fisher, 2004a; Sloutsky et al., 2001).

More direct evidence for the distinction comes from a set of recent studies using the ITR paradigm (Fisher & Sloutsky, 2005b). Note that in the recognition studies using the ITR paradigm discussed above (Fisher & Sloutsky, 2005a; Sloutsky & Fisher, 2004a, 2004b), pictures were not accompanied by labels. If label and category information are the same for young children, then when presented entities are labeled, children should exhibit effects of semantic processing—low-recognition accuracy stemming from high hits and elevated false alarms. Furthermore, effects of category labels (i.e., the words *Cat* referring to each individual cat) should differ from those of individual labels (i.e., a different count noun referring to each member of a category). In particular, only the former, but not the latter should promote category-based induction. However, because SINC does not consider category labels to be category markers, the prediction is different. In particular, because labels may overshadow corresponding visual information, introduction of labels may disrupt encoding of visual information, thus resulting in a decreased proportion of hits. Furthermore, category labels and individual labels should have comparable effects on recognition memory. This is exactly what was found—individual label and category labels exerted similar effects on recognition memory (Fisher & Sloutsky, 2005b). However, when young children were trained to perform category-based induction (Fisher & Sloutsky, 2005b), not only did their induction performance increase, but participants also exhibited patterns of recognition accuracy that were similar to those of adults (i.e., high hits and elevated false alarms).

We also found that labels and visual similarity jointly contribute to induction early in development, and SINC can quantify these contributions: For example, as mentioned above, SINC accurately predicted young children's performance on individual triads with Gelman and Markman (1986) stimuli and, as predicted, children's induction was driven by the overall similarity rather than by reliance on labels. Also, as shown in Figure 14.3, SINC accurately predicts young children's performance on a similarity judgment task (Figure 14.3A) and on induction and categorizations tasks (Figure 14.3B). Finally, SINC assumes the interrelatedness of similarity, induction, and categorization, and results support this assumption, pointing to high intercorrelations among similarity, induction, and categorization (Sloutsky & Fisher, 2004a).

Figure 14.3 about here

However, it could be argued that reliance on similarity does not constitute unequivocal evidence against the idea that conceptual assumptions constrain generalization. For example, it has been argued that categorization and induction "reflect an interaction of perceptual knowledge, language, and conceptual knowledge" (Gelman & Medin, 1993, p. 159). Therefore, it is important to know whether or not young children's generalization is driven by the category and linguistic assumptions.

In an attempt to answer this question, we conducted a study, in which 4- to 5-year-olds learned a novel category and then performed an induction task (Sloutsky, Kloos, & Fisher, 2007). The category was bound by a relational inclusion rule rather than by similarity (see Figure 14.4 for examples). After successfully learning the category, participants were presented with a triad induction task, such that one test item shared category membership with the target and another was similar to the target (while being a member of a different category). Despite the fact that young children ably learned the categories and readily categorized items throughout the experiment, they did not use this knowledge when making inductions, relying instead on appearance similarity. Therefore, there is little evidence that top-down conceptual information is an important factor in early generalization, even when this information is directly given to young children.

Figure 14.4 about here

WHY DO LABELS CONTRIBUTE TO SIMILARITY?

As mentioned above, linguistic labels play an important role in the early generalization: if two items are accompanied by the same label, young children are more likely to generalize properties from one item to another than when labels are different or no labels are introduced. In an attempt to explain these effects, two classes of explanations have been proposed. According to the language-specific explanation, young children assume that (a) entities are members of categories and (b) count nouns convey category membership (Gelman & Coley, 1991). Furthermore, for 9-month-old infants, even speech sounds devoid of morphosyntactic information may communicate category membership (Balaban & Waxman, 1997). These assumptions lead young children and infants to infer that entities that are denoted by the same count noun belong to the same category (Gelman & Markman, 1986; Markman, 1989; see also Waxman & Markow, 1995 for a discussion). Therefore, according to the language-specific explanation, if entities share a label presented as a count noun, then this shared count noun suggests that entities belong to the same category (thus supporting categorization), and further, belonging to the same category indicates that the members of the category share nonobvious properties (thus supporting inductive inference about these properties).

However, it is also possible that effects of labels stem from attentional factors, such as auditory information overshadowing (or attenuating processing) of corresponding visual information (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003). As a result of overshadowing, young children may consider entities that share the label as looking more similar (Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999), with similarity affecting inductive generalization.

Initial evidence for overshadowing was presented in the Sloutsky and Napolitano (2003) study, in which 4-year-olds and adults were presented with an auditory-visual target item (AUD_TVIS_T), where the visual and the auditory components were presented in synchrony. The Target was followed by one of the four test items. Some test items were identical to the target items (i.e., AUD_TVIS_T). Other test items had either the auditory component changed ($AUD_{new}VIS_T$), the visual component changed (AUD_TVIS_{new}), or both components changed ($AUD_{new}VIS_{new}$). Participants had to respond *same* if the two compound stimuli had the same auditory and visual components, and to respond *different* if either the auditory or visual component differed between the target and test items. The auditory components consisted of unfamiliar nonlinguistic sounds and the visual components consisted of unfamiliar images (e.g., geometric shapes). If participants encode both auditory and visual stimuli, then they should accept target items as the *same* and reject items that had either new visual or new auditory components.

It was found that 4-year-olds failed to report that the visual components changed when stimuli were cross-modal, whereas they had no difficulty noticing when the same visual components changed in the unimodal condition (see Figure 14.5 for overshadowing effects in 4- and 6-year-olds). As shown in Figure 14.5, processing of visual stimuli was not difficult per se: in the absence of auditory stimuli, young children ably encoded the visual stimuli, whereas when both visual and auditory stimuli were presented simultaneously, encoding of visual (but not of auditory) stimuli decreased compared to a unimodal baseline (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003). The results presented in Figure 14.5 also indicate that there was a decrease in overshadowing effects between 4 and 6 years of age. Furthermore, we found no evidence of overshadowing in adults: adults ably processed both auditory and visual stimuli.

Figure 14.5 about here

Finding overshadowing effects in young children and the decrease of these effects in the course of development enabled us to make predictions about infants' ability to encode cross-modal stimuli. If overshadowing effects decrease with age, then infants should exhibit even stronger overshadowing effects than young children. These predictions were confirmed empirically (Robinson & Sloutsky, 2004). In a set of experiments, Chris Robinson and I (Robinson & Sloutsky, 2004) presented the same auditory and visual compounds to infants, children, and adults. In the infant task, 8-, 12-, and 16-month-olds were familiarized to an auditory-visual compound ($AUD_T VIS_T$). At test, infants were presented with four test trials where either the auditory component changed ($AUD_{new} VIS_T$), the visual component changed ($AUD_T VIS_{new}$), both components changed ($AUD_{new} VIS_{new}$), or neither component changed ($AUD_T VIS_T$). If infants encode the auditory component during familiarization, then they should increase looking when the auditory component changes at test (i.e., $AUD_{new} VIS_T - AUD_T VIS_T > 0$), and if infants encode the visual stimulus, then they should increase looking when the visual component changes at test. Even though infants ably encoded the visual stimuli when presented unimodally, they often failed to encode visual stimuli when paired with an auditory stimulus (auditory overshadowing). Furthermore, infants exhibited auditory overshadowing under a wider range of stimulus conditions than children and adults.

Although these auditory overshadowing effects should hinder forming word-object associations, it is well known that 14- and 15-month-olds can form such associations (e.g., Schafer & Plunkett, 1998; Werker, Cohen, Lloyd, Casasola, & Stager, 1998). To examine whether words also overshadow corresponding visual input, we paired the same visual stimuli that were used in Robinson and Sloutsky (2004) with nonsense words (Sloutsky & Robinson, in press). Results presented in Figure 14.6A indicate that 8- and 12-month-olds only encoded the word, whereas 16-month-olds encoded both the word and visual stimulus. Thus, at 8 and 12 months of age, both unfamiliar words and unfamiliar sounds overshadowed visual input, whereas, at 16 months, words did not overshadow corresponding visual input. Therefore, when cross-modal stimuli are presented for a protracted period of time (such as in familiarization or habituation paradigms), by 16 months of age, words stop interfering with visual processing, whereas nonlinguistic sounds continue to interfere.

Figure 14.6 about here

Why is there a difference between encoding visual stimuli accompanied by words and by sounds at 16 months of age? First, it is possible that human speech is a special class of stimuli for humans, with infants and young children having broad assumptions that words refer to categories (e.g., Waxman & Booth, 2003). Thus, according to this “language-specific hypothesis,” labels play a special role in processing of visual information by directing children’s attention to visual input (e.g., Balaban & Waxman, 1997; Baldwin & Markman, 1989; Xu, 2002). Alternatively, it is possible that the decreased interference of words in visual processing at 16 months of age stems from familiarity effects: by that age, human speech may become more familiar than many other sounds (e.g., Jusczyk, 1998; Napolitano & Sloutsky, 2004), and, under repeated presentation conditions, more familiar stimuli may be processed faster and may be less likely to interfere with processing of visual stimuli. If familiarity can account for the increased processing of visual input, then prefamiliarizing infants to the nonlinguistic sounds prior to the experiment proper should attenuate overshadowing effects.

To determine whether the effects of words at 16 months were language-specific or stemmed from familiarity, we conducted another experiment, in which 16-month-olds were presented with the same visual patterns accompanied by the same auditory input (i.e., either linguistic or non-linguistic). However, prior to the experiment proper, participants were prefamiliarized to the auditory input. Note that the presence of prefamiliarization was the only difference between this and the previous experiment. As can be seen in Figure 14.6B, across the different auditory conditions, infants ably encoded auditory input. In contrast, encoding of visual input was affected by the familiarity of the auditory input. When the auditory input was unfamiliar, the sounds overshadowed corresponding visual input and the words did not interfere with processing of visual input (compared to the unimodal baseline). However, when infants were first prefamiliarized to the auditory component (either sound or word), prefamiliarized sounds and words facilitated processing of corresponding visual input—in these conditions, infants were more likely to encode the visual stimuli than in a unimodal baseline. Therefore, under these conditions, familiar auditory stimuli tuned attention to corresponding visual input.

Taken together, these findings present evidence that early in development, auditory input may affect attention allocated to visual input and in the course of development, visual processing becomes more independent from auditory processing. If this is the case, then effects of words may stem from the dynamics of cross-modal processing rather than from conceptual assumptions. As processing of visual input becomes more independent of processing of auditory input and as children learn that under many conditions, words are reliable cues, effects of words may become more "conceptual" in nature. This account may elucidate mechanisms underlying effects of words on early generalization: many of these effects may stem from auditory input affecting processing of visual input. In addition, this account suggests that the importance of words is not "given," but is acquired in the course of learning and development.

UNRESOLVED ISSUES

In sum, SINC—a bottom-up model reviewed in this chapter—can account well for early generalization. First, the model accurately predicts children's performance on similarity judgment, induction, and categorization tasks. I presented evidence from induction, categorization, and recognition memory tasks, indicating that early generalization is driven by similarity rather than by abstract conceptual information. Second, the model quantifies effects of labels on similarity and generalization early in development, and the theory underlying SINC suggests the mechanism driving the effects of labels. I presented evidence indicating that labels affect similarity (and thus generalization) by attenuating visual processing. The reviewed findings indicate that many effects of labels may stem from dynamics of cross-modal processing rather than from conceptual assumptions about the language and the world. These ideas also explain why phonologically similar labels may affect early generalization in a manner similar to that of identical labels. Finally, the model suggests that the contribution of various predictors (e.g., labels, appearances, etc.) to early generalization is determined by attention allocated to these predictors. The fact that attention allocated to predictors (i.e., attentional weights of the predictors) is not fixed, but can be flexibly (yet nondeliberately) adjusted in a course of learning suggests how generalization can change in the course of learning and development. At the same time, there are several issues that remain unresolved. Most importantly, SINC cannot account for many aspects of mature generalization, such as some of the sophisticated strategies exhibited by adults, or for the transition from the early to mature generalization. These are interesting challenges that have to be addressed in the future.

CONCLUSIONS

In this chapter, I considered two broad theoretical approaches to cognitive development. One approach advocates the importance of a priori constraints in cognitive development (i.e., the top-down approach). Evidence for these constraints comes from a variety of studies with infants and young children, indicating that even early in development, infants and children treat some aspects of input as more "important" or central than others. Another approach argues that people have powerful learning mechanisms enabling them to extract regularities from the input (i.e., the bottom-up approach). Proponents of the latter approach argue that conceptual knowledge is not a priori, but it emerges in the course of learning and development. Therefore, whenever children treat some aspects of input as more "important," these "important" stimuli have to be natural "attention grabbers"; otherwise the importance of these stimuli has to be learned. In the course of learning, children may realize that some stimuli are regular and reliable predictors of other important events and they may start treating the reliable predictors differently from the less reliable ones.

I then reviewed a recently proposed bottom-up model of inductive generalization and several phenomena predicted by the model; some of these phenomena present challenges to the top-down approach. I specifically focused on the role of label and appearance information in induction and presented evidence that labels affect induction by contributing to the overall similarity of compared items. I also considered a mechanism that may underlie the effects of labels on generalization: I suggested that labels (and other sounds) contribute to similarity by overshadowing (or attenuating processing of) corresponding visual input and presented supporting evidence from studies with infants and young children. Finally, I discussed how effects of words may change in the course of development as a function of increasing familiarity of human speech.

I also tried to make it clear throughout the chapter that much of the developmental story remains a mystery. Despite the significant advances in the study of inductive generalization in the past 20 years, the most interesting developmental question has not been answered: How do people become so smart, acquiring extraordinary complex knowledge that shapes their striking intellectual abilities? A detailed answer to this question would likely to constitute the most significant contribution to the study of cognitive development.

ACKNOWLEDGMENTS

This research was supported by grants from the NSF (BCS-0720135) and from the Institute of Education Sciences, U.S. Department of Education (R305H050125) to VM.S. The opinions expressed are those of the authors and do not represent views of the awarding organizations. I thank Anna Fisher and Chris Robinson for helpful comments.

References

Balaban, M. T., & Waxman, S. R. (1997). Do words facilitate object categorization in 9-month-old infants? *Journal of Experimental Child Psychology*, *64*, 3–26.

Baldwin, D. A., & Markman, E. M. (1989). Establishing word–object relations: A first step. *Child Development*, *60*, 381–398.

Booth, A. E., & Waxman, S. R. (2002). Word learning is ‘smart’: Evidence that conceptual information affects preschoolers’ extension of novel words. *Cognition*, *84*, 11–22.

Colunga, E. & Smith, L. B. (2005). From the lexicon to expectations about kinds: A role for associative learning. *Psychological Review, 112*, 347–382.

Estes, W. K. (1994). *Classification and cognition*. New York: Oxford University Press.

Fisher, A. V., & Sloutsky, V. M. (2004). When mats meow: Phonological similarity of labels and induction in young children. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the XXVI Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

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

Fisher, A. V., & Sloutsky, V. M. (2005b). Effects of category labels on induction and visual processing: Support or interference? In B. G.Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the XXVII Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Fisher, A. F., & Sloutsky, V. M. (2006). Flexible attention to labels and appearances in early induction. In R. Sun and N. Miyake (Eds.). *Proceedings of the XXVIII Annual Conference of the Cognitive Science Society* (pp. 220–225). Mahwah, NJ: Erlbaum.

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.

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

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

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. London: Oxford University Press.

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

Gelman, S.A., & Heyman, G.D. (1999). Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. *Psychological Science, 10,* 489-493.

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

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

Gelman, S. A., & Medin, D. L. (1993). What's so essential about essentialism? A different perspective on the interaction of perception, language, and concrete knowledge. *Cognitive Development*, *8*, 157–167.

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

Hall, G. (1991). *Perceptual and associative learning*. New York: Oxford University Press.

Johnson, S. P. (2003). Development of fragmented versus holistic object perception. In G. Schwarzer & H. Leder, Helmut (Eds.), *The development of face processing* (pp. 3–17). Ashland, OH: Hogrefe & Huber Publishers.

Jusczyk, P. W. (1998). *The discovery of spoken language*. Cambridge, MA: MIT Press.

Keil, F.C. (2003). Categories, cognitive development and cognitive science. In D. H. Rakison & L. M. Oakes (Eds.), *Early category and concept development: Making sense of the blooming buzzing confusion* (pp. vi-xi). New York: Oxford University Press.

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

Keil, F. C., Smith, W. C., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition*, *65*, 103–135.

Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology, 15*, 483–524.

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

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

Lewkowicz, D. J. (1988a). Sensory dominance in infants: 1. Six-month-old infants' response to auditory-visual compounds. *Developmental Psychology*, *24*, 155–171.

Lewkowicz, D. J. (1988b). Sensory dominance in infants: 2. Ten-month-old infants' response to auditory-visual compounds. *Developmental Psychology*, *24*, 172–182.

Mandler, J. M. (1997). Development of categorisation: Perceptual and conceptual categories. In G. Bremner, A. Slater & G. Butterworth (Eds.), *Infant development: Recent advances* (pp. 163–189). Hove, England: Psychology Press.

Mandler, J. (2004). *The foundations of mind: Origins of conceptual thought*. New York, NY: Oxford University Press.

Madole, K. L., Oakes, L. M., & Cohen, L. B. (1993). Developmental changes in infants' attention to function and form–function correlations. *Cognitive Development*, *8*, 189–209

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

Markman, E.M., and Hutchinson, J.E. (1984). Children’s sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology*, 16: 1–27.

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

Massey, C. M., & Gelman, R. (1988). Preschooler's ability to decide whether a photographed unfamiliar object can move itself. *Developmental Psychology*, *24*, 307–317.

Medin, D. (1975). A theory of context in discrimination learning. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 9, pp. 263–314). New York: Academic Press.

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.

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

Oakes, L. M., & Madole, K. L. (2003). Principles of developmental changes in infants' category formation. In D. H. Rakison & L. M. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp. 132–158). New York: Oxford University Press.

Rakison, D. H. (2003). Parts, motion, and the development of the animate–inanimate distinction in infancy. In Rakison & L. M. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp. 159–192). New York, NY: Oxford University Press.

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

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

Schafer, G., & Plunkett, K. (1998). Rapid word learning by fifteen-month-olds under tightly controlled conditions. *Child Development*, *69*, 309–320.

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. (2005). Similarity, Induction, Naming, and Categorization (SINC): Generalization or Inductive Reasoning? Response to Heit and Hayes. *Journal of Experimental Psychology: General*, *134*, 606–611.

Sloutsky, V. M., Kloos. H., & Fisher, A. V. (2007). What's beyond looks? Reply to Gelman and Waxman. *Psychological Science*, *18*, 556–557.

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., & 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., & 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.

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.

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

Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological Science*, *13*, 13–19.

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

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29–56.

Spelke, E. S. (1994). Initial knowledge: Six suggestions. *Cognition*, *50*, 431–445.

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 & Cognition*, *29*, 424–432.

Werker, J. F., Cohen, L. B., Lloyd, V. L., Casasola, M., & Stager, C. L. (1998). Acquisition of word–object associations by 14-month-old infants. *Developmental Psychology*, *34*, 1289–1309.

Waxman, S. R., & Booth, A. E. (2003). The origins and evolution of links between word learning and conceptual organization: New evidence from 11-month-olds. *Developmental Science*, *6*, 130–137.

Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology*, *29*, 557–302.

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

Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, *43*, 337–375.

Wexler, K., & Culicover, P. W. (1980). *Formal principles of language acquisition*. Cambridge MA: MIT Press.

Woodward, A. L., & Markman, E. M. (1998). Early word learning. In W. Damon, D. Kuhn & R. Siegler (Eds.), *Handbook of child psychology, Volume 2: Cognition, perception and language* (pp. 371–420). New York: John Wiley and Sons.

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

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*, 245–272.

Yoshida, H., & Smith, L. B. (2005) Linguistic cues enhance the learning of perceptual cues. *Psychological Science*, *16*, 90–95.

Figure 14.1 An example of visual occlusion, B–E. Possible states of the world compatible with A.

Figure 14.2 (Originally presented in Sloutsky & Fisher, 2004a). (A) Stimuli used in Gelman and Markman (1986). (B) Predicted and observed proportions of B-choices by stimuli triads in induction task. Error bars represent standard errors of the mean.

Figure 14.3 (Originally presented in Sloutsky & Fisher, 2004a). (A) Predicted and observed probabilities of B-choices as a function of similarity ratio and labeling in the similarity judgment task. (B) Predicted and observed probabilities of B-choices as a function of similarity ratio and labeling in induction and categorization tasks. Error bars represent standard errors of the mean.

Figure 14.4 Examples of stimuli used by Sloutsky, Kloos, and Fisher (2007).

Figure 14.5 Recognition accuracy by modality, presentation condition and age. Error bars represent standard errors of the mean.

Figure 14.6 (A) Encoding of auditory and visual stimuli at 8-, 12-, and 16-months of age. (B) 16-month-olds' encoding of auditory and visual components across different auditory conditions.