Child Development, May/June 2008, Volume 79, Number 3, Pages 639 – 651

Attentional Learning and Flexible Induction: How Mundane Mechanisms Give Rise to Smart Behaviors

Vladimir M. Sloutsky
The Ohio State University

Anna V. Fisher
Carnegie Mellon University

Young children often exhibit flexible behaviors relying on different kinds of information in different situations. This flexibility has been traditionally attributed to conceptual knowledge. Reported research demonstrates that flexibility can be acquired implicitly and it does not require conceptual knowledge. In Experiment 1, 4- to 5-year-olds successfully learned different context-predictor contingencies and subsequently flexibly relied on different predictors in different contexts. Experiments 2A and 2B indicated that flexible generalization stems from implicit attentional learning rather than from rule discovery, and Experiment 3 pointed to very limited strategic control over generalization behaviors in 4- to 5-year-olds. These findings indicate that mundane mechanisms grounded in associative and attentional learning may give rise to smart flexible behaviors.

Even early in development, people's generalization is remarkably flexible—depending on a situation, people may rely on different kinds of information. This flexibility has been found in a variety of generalization tasks, including lexical extension, categorization, and property induction. For example, in a lexical extension task (Jones, Smith, & Landau, 1991), 2- to 3-year-olds were presented with a target, which was named (i.e., "this is a dax"), and 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.

Children exhibit similar flexibility in categorization and induction tasks. For example, in a categorization task, 3- to 4-year-olds were more likely to group items on the basis of color if the items were introduced as food, but on the basis of shape if the items were introduced as toys (Macario, 1991). In another task, 4- to 5-year-olds were presented with a target and two test items, such that one test item shared the label with the target and the other looked similar to the target. Participants were then told that the target had a particular property and asked which of the test items had the same property. Participants were more likely to rely on linguistic labels when inferring a biological property than when inferring a physical property (Gelman & Markman, 1986; see also Heit & Rubinstein, 1994, for similar findings in adults). 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." Although such flexibility is found across multiple generalization tasks and is present early in development, the underlying mechanisms remain unclear. There are several theoretical proposals attempting to explicate these mechanisms. Some of these proposals ground this flexibility in associative learning, whereas others argue that this flexibility requires some conceptual knowledge and it cannot be explained by a purely associationist account.

Flexibility Is a Function of Conceptual Knowledge

There is a prominent theoretical argument that flexible generalization requires conceptual knowledge: Without domain-specific conceptual knowledge (or theories), people will not know what attributes to rely on and when (Keil, 1991; Keil, Smith, Simons, & Levin, 1998; Murphy & Medin, 1985). The

This research was supported by grants from the National Science Foundation (BCS-0720135) and from the Institute of Education Sciences, U.S. Department of Education (R305H050125) to V.M.S. The opinions expressed are those of the authors and do not represent views of the awarding organizations. We thank Heidi Kloos, John Opfer, and Chris Robinson for helpful comments.

Correspondence concerning this article should be addressed to Vladimir M. Sloutsky, Center for Cognitive Science, Ohio State University, 208C Ohio Stadium East, 1961 Tuttle Park Place, Columbus, OH 43210. Electronic mail may be sent to sloutsky.1@osu.edu.

0009-3920/2008/7903-0010

argument has three parts: (a) raw associations are too unconstrained to underlie this flexibility, (b) conceptual knowledge is necessary to determine when a feature is relevant and when it is not (i.e., one has to know what foods and toys are to determine that color is important for foods and shape is important for toys), and (c) feature relevance may depend on deliberate and strategic process of feature weighting.

The latter part of the argument was mentioned as a possibility (Gelman & Medin, 1993, p. 164), whereas the two former parts have been argued more forcefully. For example, R. Gelman and Williams (1998) suggested that flexible generalization presents a problem for associative learning for the following reasons. To exhibit the reported flexibility, an associative learner has to assign different weights to the same attribute dimensions under different conditions (e.g., color has a greater weight for cars, whereas shape has a greater weight for foods). However, "the weighting method assumes that we use different concepts to decide which weights to assign to seemingly *same* attributes" (R. Gelman & Williams, 1998, p. 602). Therefore, the associative learning account becomes dangerously circular: Differential weighting of features that allegedly explains acquisition of concepts cannot be done without having concepts in the first place. Other researchers have also argued that such flexibility (or property by domain interactions) is "difficult to explain using models that make reference only to the featural similarity of the inductive base and target" (Hayes & Thompson, 2007, p. 471).

Similar arguments were offered to explain flexibility found in lexical extension tasks (Booth & Waxman, 2002; Booth, Waxman, & Huang, 2005). For example, in one experiment (Booth & Waxman, 2002), participants were presented with a lexical extension task under different cover story conditions: for some participants, items were introduced as animate entities, whereas for others as inanimate entities. Despite the fact that both conditions used identical sets of items, 3-year-olds exhibited different patterns of label extension. In the inanimate condition, labels were extended by shape alone, whereas in the animate condition, they were extended by shape and texture. It was concluded that "because the objects presented in each condition were precisely the same perceptually, the results cannot be explained by a purely perceptual account" (Booth et al., 2005, p. 493).

In short, it has been argued that flexible generalization presents a challenge to associative accounts: Associations alone are too unconstrained to explain this flexibility, whereas domain-specific conceptual knowledge is necessary to direct generalization in each domain (i.e., to guide what attributes to rely on and when). Furthermore, it is possible that this conceptual knowledge is deployed in a deliberate and strategic manner in the course of a reasoning-type process (cf. Gelman & Medin, 1993). However, although proponents of the knowledge-based account argue that conceptual knowledge constrains otherwise unconstrained associations, it remains unclear as to where this knowledge comes from, under what conditions it gets deployed, and how it interacts with associative mechanisms (see Sloutsky, Kloos, & Fisher, 2007, for a discussion). Perhaps some proponents of this position believe that this knowledge is innate; however, at present, this issue remains largely unaddressed. We return to this issue again in the General Discussion section.

Flexibility Is a Function of Associative Learning

Proponents of another position argue that early generalization is driven by automatically detected perceptual similarity (Colunga & Smith, 2005; French, Mareschal, Mermillod, & Quinn, 2004; Jones & Smith, 2002; Rogers & McClelland, 2004; Sloutsky, 2003; Sloutsky & Fisher, 2004, 2005). However, if generalization is driven by automatically detected similarity, how can similar stimuli (depending on a situation) result in different patterns of generalization? And how could this flexibility be acquired by means of associative learning?

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. Learning of these correlations may result in differential allocation of attention to different stimulus properties in different situations or contexts, with flexibility being a result of this learning.

In particular, there is evidence that allocation of attention to different stimulus dimensions does change in the course of learning, with more attention allocated to predictive stimulus dimensions and less attention allocated to nonpredictive dimensions (e.g., Kersten, Goldstone, & Schaffert, 1998; Kruschke, 1992; Nosofsky, 1986). There is also evidence that correlated cues mutually reinforce each other and such clusters are more likely to be detected than isolated cues (Billman & Knutson, 1996; Yoshida & Smith, 2005). These clusters of correlated cues may represent situation or "context" variables, and when these variables correlate with a predictive dimension, these correlations may give rise to the observed flexibility. More specifically, if dimension X is predictive in context A and dimension Y is predictive in context B, then participants should learn to attend to dimension X in context A and to attend to dimension Y in context B. Furthermore, the likelihood of successful learning could be greater if context A and context B consist of multiple features that correlate with each other and with the predictive dimensions.

This process of associative learning is presented schematically in Figure 1. Suppose that the task is to generalize a property from a target item to a test item. Further suppose (as shown in Figure 1A) that Test 1 shares attribute X with the target (e.g., both have the same color), whereas Test 2 shares attribute Y with the target (e.g., both have the same shape). At the beginning of learning (see Figure 1B), both dimensions X and Y (e.g., shape and color) have comparable weights and a participant can rely on either dimension. Also note that there is a cluster of other features that determine a context, some of which may correlate with the predictive dimension; however, at the beginning of learning, it is not known which context features correlate with which dimension.

In the course of learning (Figure 1C), the system learns that a particular set of context variables (i.e., Context 1) co-occurs with shape being predictive, whereas another set of variables (i.e., Context 2) co-occurs with color being predictive. As a result, the system learns to allocate attention to shape in Context 1 and to color in Context 2, with shape becoming more important in Context 1 and color becoming more important in Context 2. Therefore, following learning, participants will rely on shape when asked to generalize in Context 1 and on color when asked to generalize in Context 2, thus exhibiting flexible

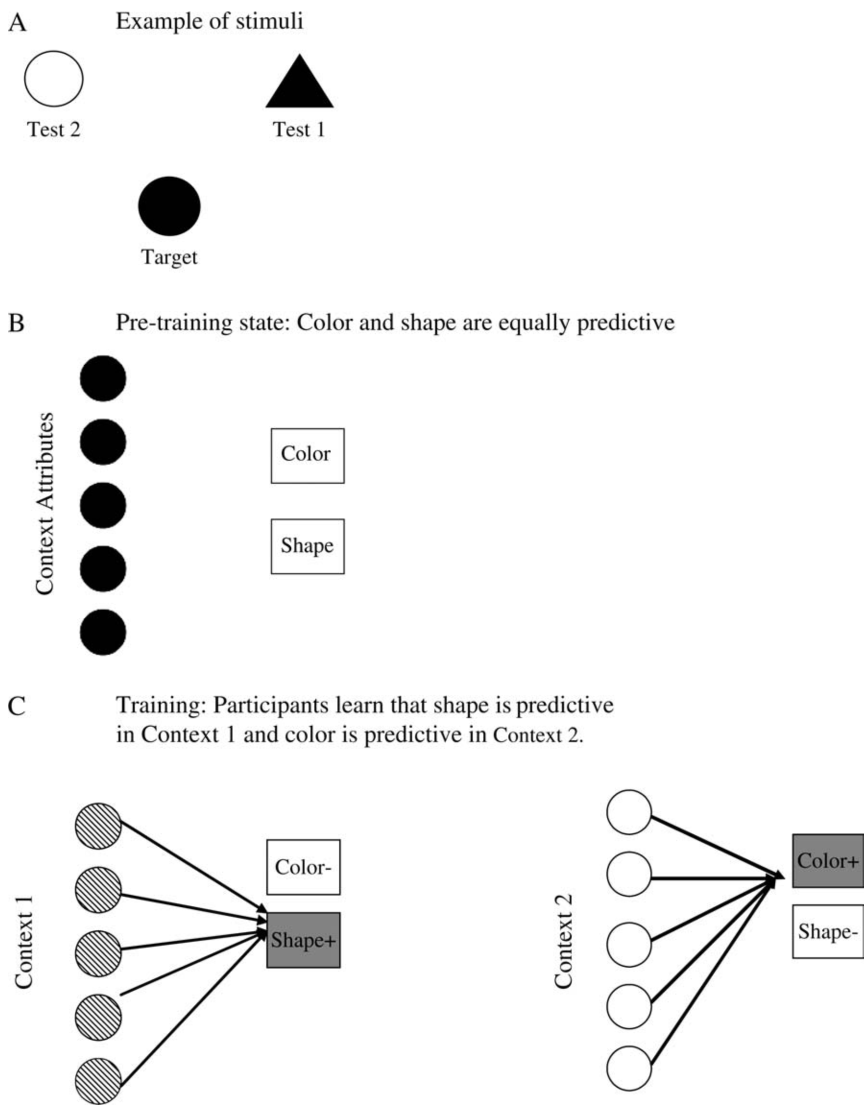

Figure 1. An overview of the proposed learning account.

generalization. Note that according to this account, a system can acquire and exhibit flexible generalizations by using purely associative information. Furthermore, there is no need for explicit knowledge or awareness of context-predictor co-occurrences: Given that animals can ably learn contingencies between clusters of correlated features and predictors (see Young & Wasserman, 2004, for a review), it seems likely that these contingencies can be learned implicitly.

Although this account of flexible generalization appears plausible, there is no direct empirical evidence that such flexibility can be acquired in the course of associative learning. Therefore, the goal of the current research is to test the proposed account and present direct evidence that flexible generalization may emerge as a result of associative learning.

In sum, we argue that (a) flexible reliance on different sources of information can be acquired through associative learning, (b) associative learning is likely to be implicit, and (c) no conceptual knowledge or the ability to strategically control attention is needed for such learning. The reported experiments were designed to test these hypotheses.

To foreshadow, in Experiment 1, we examine the ability of preschoolers to acquire flexible generalization through associative learning. In Experiments 2A and 2B, we examine the extent to which learning is implicit and whether participants exhibit awareness of what they learned. Finally, in Experiment 3, we examine the extent of children's strategic control over their generalization.

Experiment 1

The goal of Experiment 1 was to examine whether participants can learn to flexibly rely on different (arbitrarily chosen) predictors in different (arbitrarily chosen) contexts. To achieve this goal, participants were given a simple induction task: They were presented with a target and two test items and told that there was a smiley face hidden behind the target and were asked to find a test item that also had a smiley face behind it. As shown in Figure 1A, one test item had the same shape as the target, whereas another had the same color. Therefore, when performing induction, participants could rely either on color or on shape. In the Baseline condition, participants were tested using this induction task, whereas in the Experimental condition, testing was preceded by training.

During training, participants were presented with triads where they could only rely on shape (i.e., all items had the same color, but only one item had the same shape as the target) or with triads where they could only rely on color (i.e., all items had the same shape, but only one item had the same color as the target). The former kinds of triads were always presented in one context (context variables are explained in detail in the Method section), whereas the latter kinds of triads were always presented in another context. All participants received training with both kinds of triads. During testing, with ambiguous triads (those where they could rely either on color or shape, similar to the one presented in Figure 1A) participants were given an induction task. Some participants were given the induction task in Context 1 and others in Context 2. It was expected that as a result of training, participants will rely on shape when tested in Context 1 and on color when tested in Context 2, with performance in both conditions being different from no-training baseline performance.

Method

Participants

In this and all subsequent experiments, participants were recruited from suburbs of Columbus, Ohio, on the basis of returned consent forms. The majority of participants were Caucasian from middle-class families. There were 74 children participating in the experiment ($M = 5.23$ years, $SD = 0.29$ years; 30 girls, 44 boys), with 32 participants in the baseline condition and 21 in each of the two experimental conditions (see Design and Procedure section).

Materials

Materials consisted of 16 training triads and 16 testing triads, with each triad including a target item and two test items located above the target item (see Figure 2 for examples of training and testing triads). All stimuli were triangles and circles colored blue or red. The triads were presented either in Context 1 or Context 2, with context variables being the color of the background on which the triads appeared and the location of the triads on the screen. In Context 1, triads appeared on a yellow background in the upper right corner of the screen, and in Context 2, triads appeared on a green background in the bottom left corner of the screen.

Design and Procedure

There were three between-subjects conditions—baseline, Context 1 testing, and Context 2 testing—and the experiment included two phases—training and testing (note that the baseline condition included

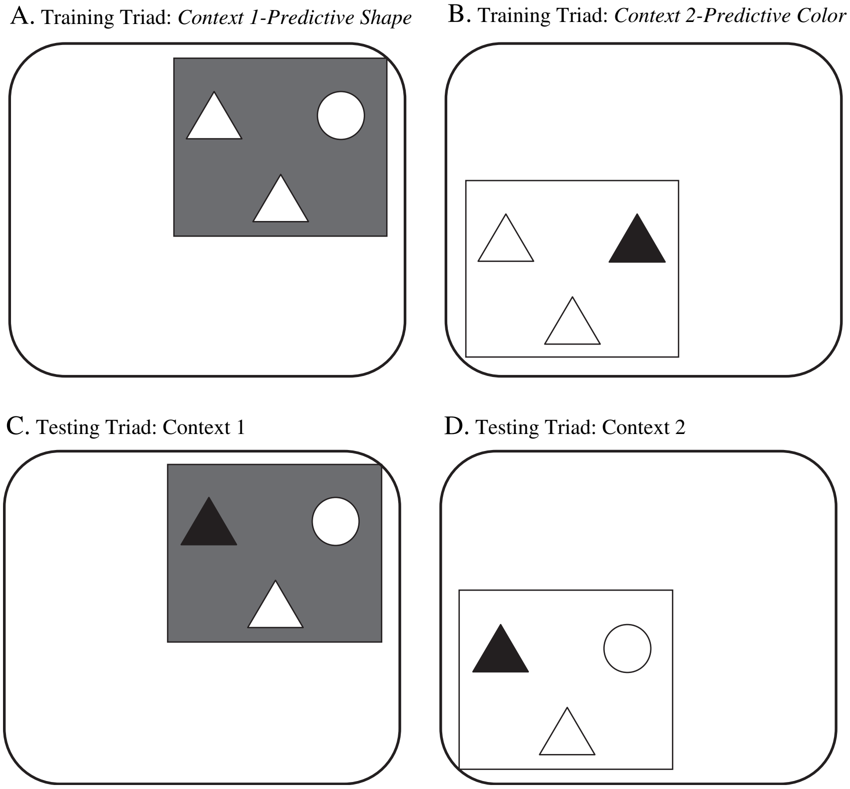

Figure 2. Examples of training and testing stimuli used in Experiments 1 and 2. (A) Context 1-Predictive Shape training triad, (B) Context 2-Predictive Color training triad, (C) Context 1 testing triad, and (D) Context 2 testing triad.

testing phase only). Table 1 presents an overview of the design, and Figures 3A – 3B present details of training and testing phases, respectively.

Participants were randomly assigned to one of the three conditions. Participants in Context 1 testing and Context 2 testing conditions received identical training; however, one group of participants was tested only in Context 1 and the other group was tested only in Context 2. There was no training phase in the baseline condition; participants were presented only with 16 testing trials, with half of the trials being presented in Context 1 and half being presented in Context 2.

Participants were tested individually in their child care centers by female hypotheses-blind experimenters and all stimuli were presented to them on the screen of a laptop computer. Participants were told that they will play a game, in which they will need to find where a smiley face was hiding. On each trial, they were presented with a triad and told that the target had a smiley face hiding behind it. They were then told that one of the test items also had a smiley face hiding behind it, and their task was to determine which of the test items had a smiley face behind it. During training, participants were provided with feedback, such that if they responded correctly, their response was followed by a smiley face appearing on the screen, whereas incorrect responses were followed by a frowning face. At the conclusion of training, participants were told that that they will continue playing the game, but this time they will have to guess where the smiley face is hidden without actually seeing the face (i.e., no feedback was provided during the testing phase). Participants were then presented with 16 no-feedback testing trials.

Table 1
Overview of the Design in Experiments 1 and 2

Conditions	Phases of the experiment	
Experimental conditions	Training	Testing
Baseline condition[a]	No training	Testing

[a]There was no baseline condition in Experiment 2B.

Training phase. During training, participants were presented with 16 training triads presented in three blocks, with a total of 48 training trials (see Figure 3A for a schematic presentation of the training phase). There were two types of training triads: In half of the training triads, all three items had the same color,

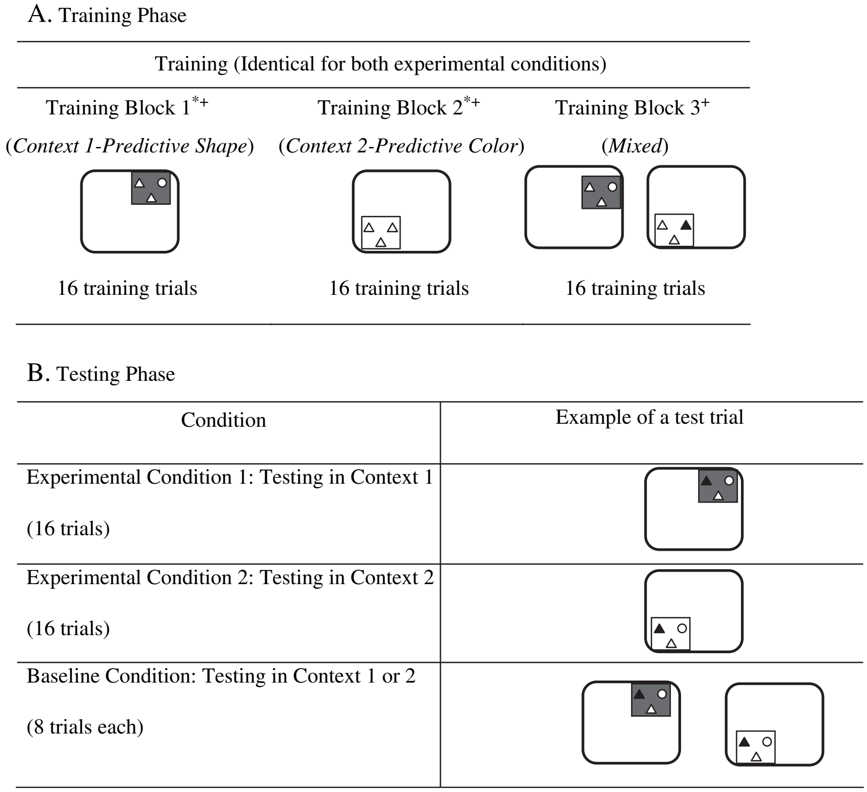

Figure 3. Training and testing phases in Experiments 1 and 2. (A) Training and (B) Testing.
Note. *The order of training Blocks 1 and 2 was counterbalanced. +The order of trials was randomized for each participant.

whereas only one test item matched the shape of the target (hereafter, *predictive shape* triads), and in another half, all three items had the same shape, whereas only one test item matched the color of the target (hereafter, *predictive color* triads). Predictive shape and predictive color triads were presented in different contexts. Examples of both types of triads are presented in Figure 2.

The three training blocks were presented seamlessly, without breaks. In the Context 1-Predictive Shape training block, all items within a triad had the same color (half of the triads included red items and half included blue items), whereas only one test item matched the shape of the target. Therefore, in this block, only shape but not color differentiated the test items. All triads in this block appeared in Context 1. There were eight Context 1-Predictive Shape triads, with each triad presented twice.

In the Context 2-Predictive Color training block, all items within a triad had the same shape (half of the triads included three circles and half included three triangles), whereas only one test item matched the color of the target. Therefore, in this block, only color but not shape differentiated the test items. All triads in this condition appeared in Context 2. There were eight Context 2-Predictive Color triads, with each triad presented twice.

The Mixed training block consisted of eight Context 1-Predictive Shape triads intermixed with eight Context 2-Predictive Color triads. The order of the first two blocks was counterbalanced across participants, and the Mixed training block was always presented last. The order of trials within each block was randomized for each participant.

Testing phase. Training was followed by testing, with participants presented with 16 testing triads (see Figure 3B for a schematic presentation of the testing phase). For half of the trained participants, the testing triads were presented in Context 1, and for another half of the trained participants, the testing triads were presented in Context 2 (examples of both types of triads are presented in Figure 2). Note that except for context variables (the color of the background and the screen location), testing triads were identical in the two testing conditions. Each testing triad consisted of a target and two test items, such that one of the test items matched the color of the target and another matched the shape. Therefore, testing triads differed from training triads in that participants could rely either on color or on shape when matching a test item with the target.

Results and Discussion

Across the three training blocks, participants were exceedingly accurate during training, $M = 0.98$ in Context 1 and $M = 0.95$ in Context 2, above chance, one-sample ts > 22.6, ps $< .0001$. To examine effects of training on performance during testing, we compared proportions of shape-based responding in each testing condition with its respective no-training baseline. These proportions are presented in Figure 4. As can be seen in the figure, performance in the training conditions differed from that in the no-training baseline: When tested in Context 1, participants were more likely to rely on shape than those in the baseline condition, and in Context 2, they were more likely to rely on color than in the baseline condition, both independent-sample ts > 2.4, ps $< .02$, ds > 0.61. There were also marked differences between the training conditions: Participants were significantly more likely to rely on shape when tested in Context 1 than when tested in Context 2, independent sample $t(40) = 4.56$, $p < 0001$, $d = 1.24$, whereas no such differences were observed in the baseline condition, $t < 1$.

These findings indicate that whereas there were no baseline differences across the contexts, after training, participants were more likely to rely on shape in Context 1 and on color in Context 2. Therefore, training resulted in flexible generalization performance. Given that no instructions or explanations were given to children, it seems likely that flexible generalization was acquired in the course of associative learning.

It could be argued, however, that flexibility acquired in the course of training in Experiment 1 did not stem from implicit associative learning but rather from participants discovering the "rules of the game" and deliberately following the rules. There were at least two rules to be discovered during learning: (a) rely on shape when shapes appear here (i.e., in the upper right corner and/or on yellow background) and (b) rely on color when shapes appear there (i.e., in the lower left corner and/or on green background). Therefore, for this explanation to be correct, participants had to (a) discover both rules during training, (b) determine the context during testing, and (c) follow the first rule in Context 1 and the second rule in Context 2. The fact that young children may have difficulty following the rule, even if the rule is given to them (cf. Fisher & Sloutsky, 2006; Napolitano & Sloutsky, 2004), casts doubt on the possibility of rule discovery. However, we deemed it necessary to directly examine the issue of whether flexibility in Experiment 1 was achieved through implicit associative learning. This issue was addressed in Experiment 2.

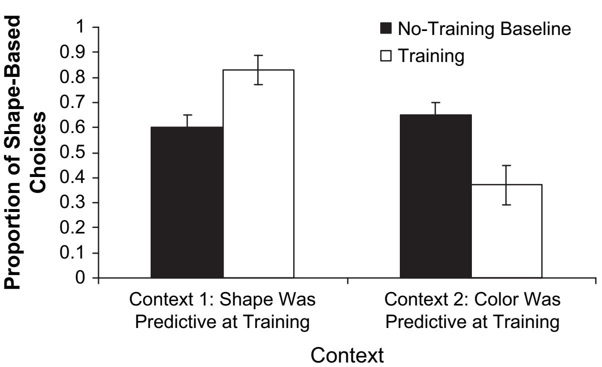

Figure 4. Proportions of shape-based choices during testing by context and condition in Experiment 1.

Experiment 2A

Experiment 2A was a replication of Experiment 1, with one major modification: After testing, participants were presented with a rule-checking procedure. The goal of the rule-checking procedure was to examine whether participants discovered the rule or not. If participants discovered the rule, such rule discovery may manifest itself in some explicit knowledge of rules. Alternatively, implicit learning should manifest itself in successful performance and little or no awareness of what was learned.

Method

Participants

There were 33 children participating in the two experimental conditions ($M = 5.31$ years, $SD = 0.23$ years; 10 girls, 23 boys), with 15 participants in Context 1 testing and 18 participants in Context 2 testing. There were also 14 children ($M = 5.19$ years, $SD = 0.23$ years; 6 girls, 8 boys) participating in the baseline condition. Participant recruitment and demographics were similar to those in Experiment 1.

Materials, Design, and Procedure

Materials were similar to those of Experiment 1, with the following exception: The two experimental conditions of Experiment 2A consisted of three phases: training, testing, and rule checking.

Training and testing. Training and testing phases were identical to those in Experiment 1: Participants were presented with 48 training trials accompanied by feedback (presented over three blocks) and with 16 no-feedback testing trials.

Rule checking. An outline of the rule-checking phase is presented in Figure 5. Participants were first presented with a triad where only color was predictive (the triad was presented on a white background in the center of the screen) and were told that this was a ''color game.'' Their attention was attracted to the fact that it was called the color game because in this game only color, but not shape, was important: All items had the same shape but some had different color. They were then presented with a triad where only shape was predictive (the triad was presented on a white background in the center of the screen) and told that this was a ''shape game.'' Their attention was attracted to the fact that it was called the shape game because in this game only shape, but not color, was important: All items had the same color but some had different shape. The order in which the two games were introduced was counterbalanced. Participants were then presented with six comprehension trials, with three shape and three color triads (these trials were presented in a random order). On each comprehension trial, they were asked whether the presented trial was the color game or the shape game. Finally, they were presented with 12 rule-checking trials. Half of these rule-checking trials consisted of a green background presented in the bottom left corner of the screen (Context 2) and half consisted of a yellow background presented in the upper right corner of the screen (Context 1). On each trial, participants were

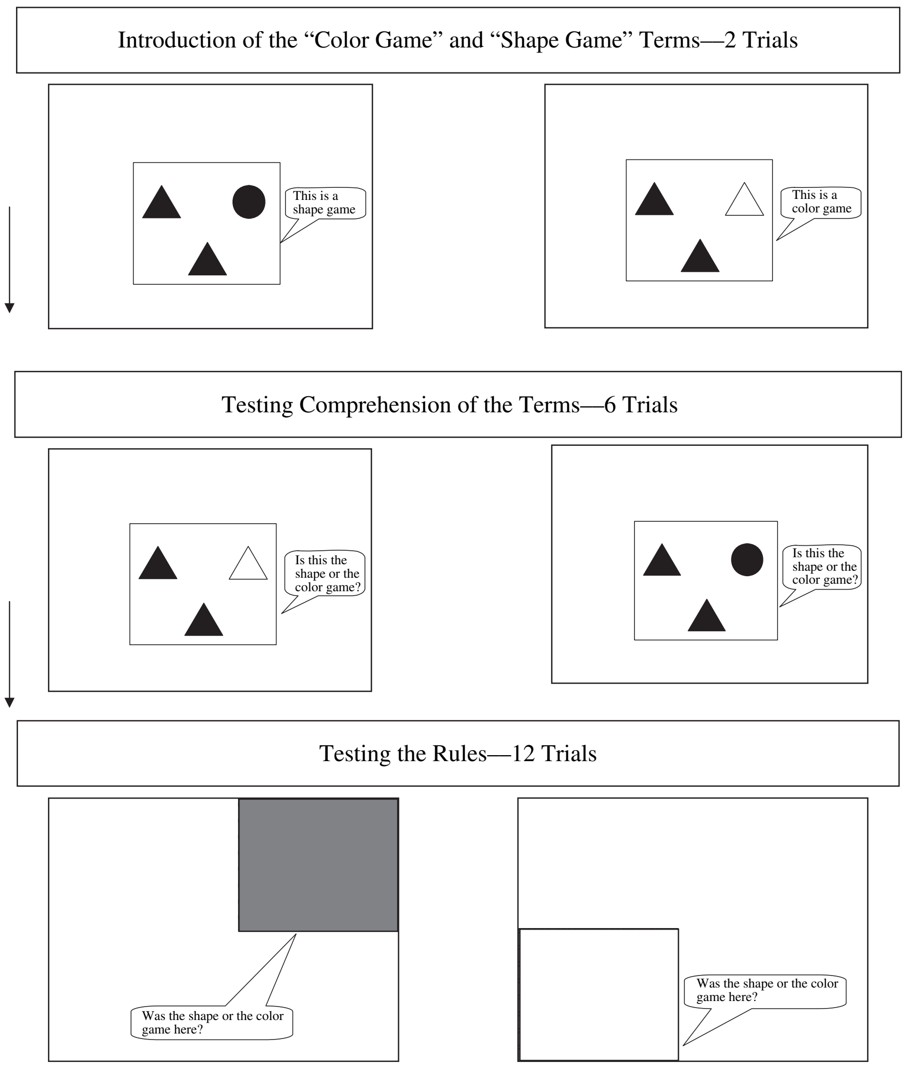

Figure 5. Overview of the rule-checking phase in Experiment 2.

pointed to the presented background and asked if they played the color or the shape game on that background in this location.

Results and Discussion

Training and Testing

Across the three training blocks, participants were accurate during training, $M = 0.85$ in Context 1 and $M = 0.80$ in Context 2, above chance, one-sample ts > 5.3, ps $< .0001$. To examine effects of training on performance during testing, we compared proportions of shape-based responding in each testing condition to the baseline performance (see Figure 6). Most important, there were significant differences between experimental and baseline conditions: When tested in Context 1 (mean shape-based responding was 82%), participants were more likely to rely on shape than in the baseline condition, whereas in Context 2 (mean shape-based responding was 33%), they were more likely to rely on color than in the baseline condition, both independent-sample ts > 2.18, ps $< .05$. The difference between experimental conditions (i.e., mean shape-based responding 82% vs. 33%) was also significant, independent sample $t(31) = 5.03$, $p < .0001$, $d = 1.33$. These findings replicate those of Experiment 1 pointing to successful training and indicating that after training participants were more likely to rely on shape in Context 1 and on color in Context 2.

Rule Checking

Recall that the rule-checking phase consisted of comprehension trials and rule-testing trials. Data from the rule-checking phase are presented in Figure 7. As can be seen in the figure, in both conditions, participants exhibited above chance comprehension accuracy (i.e., they had no difficulty identifying "the color game" and "the shape game"), both ts > 3.7, Bonferroni-adjusted ps $< .005$. At the same time, there was no evidence of rule learning, with accuracy on rule-testing trials not different from chance, both ts $<$ 1.5, Bonferroni-adjusted ps $> .3$. Furthermore, there was no evidence that accuracy on rule testing was a predictor of accuracy during the testing phase: For either testing condition, the regression between rule accuracy and testing accuracy did not approach significance, both Fs < 1.5, ps $> .25$. Therefore, there was little evidence that participants discovered the rule in the course of learning or used the rule in the course of testing. Instead, the results indicate that learning resulted in implicit rather than in explicit knowledge.

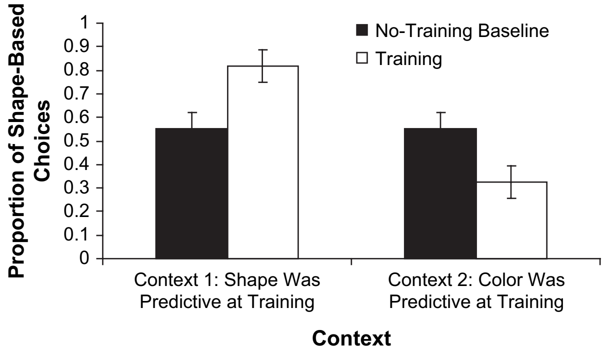

Figure 6. Proportions of shape-based choices during testing by context and condition in Experiment 2A.

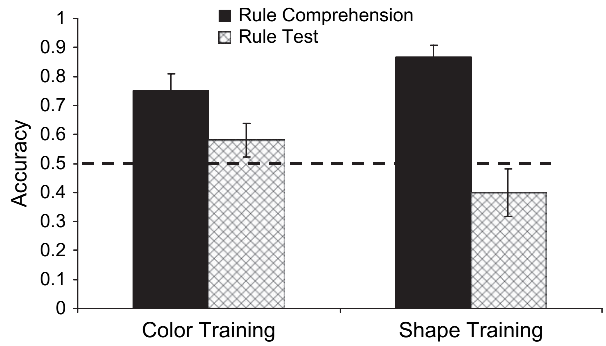

Figure 7. Accuracy of comprehending and reporting the rules by condition in Experiment 2A.
Note. The dashed line represents the chance level.

One can argue, however, that because the rule-checking phase was administered after the testing phase, the low rule-checking performance could stem from participants' forgetting the rule by the time the knowledge of the rule was checked. To address this issue, we conducted Experiment 2B, which was similar to Experiment 2A, except that the rule-checking phase was followed by a second testing phase. If participants accurately perform the task in the second testing phase, then the forgetting explanation could be eliminated.

Experiment 2B

Method

Participants

There were 10 children participating in the experiment ($M = 5.26$ years, $SD = 0.28$ years; 4 girls, 6 boys), with all children participating in Context 1 testing condition. An additional participant was tested but not included in the sample due to chance performance during the training phase.

Materials, Design, and Procedure

Materials, design, and procedure were similar to those in Experiment 2B, with the following exceptions. First, given that Experiment 2A fully replicated Experiment 1, we deemed it sufficient to include only one condition (Context 1 testing) in Experiment 2B. Second, following the rule-checking phase, participants were tested again, with Testing 2 phase being identical to Testing 1 phase. Therefore, the experiment included the following phases: training (48 trials over three blocks), Testing 1 (16 trials), rule checking (6 comprehension and 12 rule testing trials), and Testing 2 (16 trials).

Results and Discussion

Similar to previous experiments, participants were exceedingly accurate during training, $M = 0.92$, above chance, one-sample $t(9) = 10.08$, $p < .0001$. To examine effects of training on performance during testing, we compared proportions of shape-based responding to the average of baseline performances in Experiments 1 and 2 (mean baseline shape-based responding was 59%), which was used as a theoretical mean in subsequent statistical analyses. First, participants were accurate at both Testing 1 and Testing 2 (84% vs. 95% of shape-based responding, respectively), with percent of shape-based choices in both testing phases being significantly above the baseline, both one-sample ts > 3.6, ps < .005. In addition, there was no evidence that performance dropped during Testing 2 (if anything, performance during Testing 2 was numerically higher than performance during Testing 1, paired-sample $t(9) = 1.87$, $p = .094$. Finally, rule-checking data replicated those of Experiment 2A: Participants exhibited above chance comprehension accuracy (75% correct), one-sample $t > 3.31$, ps < .01, whereas there was no evidence of rule learning (55% correct), not different from chance, $t < 1$.

These results replicate and further extend those of Experiment 2A: Participants were accurate when tested before and after the rule-checking phase, whereas their accuracy on rule-testing trials was not different from chance. The results indicate that low performance on rule-testing trials did not stem from forgetting and support the conclusion that accurate testing performance stemmed from implicit rather than explicit knowledge.

Experiments 1 and 2 indicated that flexibility can be acquired in the course of implicit learning. However these experiments left an important question unanswered: Can this flexibility be achieved by deliberate weighing of different attributes in different contexts? In particular, if young children can deliberately weigh attributes (cf. Gelman & Medin, 1993), they should have no difficulty doing so when told which attribute is important. However, there is evidence that preschoolers often have limited strategic control of attention, thus casting doubt on the possibility of deliberate attribute weighting. In particular, Napolitano and Sloutsky (2004) presented 4- to 5-year-olds with auditory-visual compound and asked participants to remember the picture. Although participants could remember the pictures when presented unimodally, they failed to do so under the cross-modal presentation, despite the fact that instructions were repeated on every trial. Given this difficulty in deliberately controlling attention (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004), it seems unlikely that participants can deliberately weigh attributes, thus controlling their generalization.

This issue was addressed directly in Experiment 3, in which participants were told explicitly the rule of the game and were tested on their ability to follow the rule. The rule, however, was substantially simplified compared to contingencies used in Experiment 1: Unlike the complex rule of Experiment 1, in Experiment 3, some participants were instructed to rely only on color and others were instructed to rely only on shape. Therefore, instead of deliberate weighting of different attributes in different situations, participants simply needed to focus on one predictor (e.g., shape) and not the other (e.g., color).

Experiment 3

Method

Participants

There were 18 children ($M = 5.11$ years, $SD = 0.30$ years; 8 girls, 10 boys) participating in the focus on shape condition and 23 children ($M = 5.24$ years, $SD = 0.31$ years; 16 girls, 7 boys) participating in the focus on color condition (see the Materials, Design, and Procedure section).

Materials, Design, and Procedure

The experiment had two between-subjects conditions: focus on shape and focus on color. In both conditions, participants were told that they will play a game in which they could win by guessing correctly where a smiley face was hidden. They were also told that the game had a secret to it: In the former condition, the secret was that the smiley face was always hiding behind the picture that had the same shape as the target. In the latter condition, they were told that the secret was that the smiley face was hiding behind the picture that had the same color as the target. In both conditions, participants were asked to remember the secret and were reminded that to win the game, they had to pay attention to the shape or color, respectively. Both conditions had identical materials and procedures and these were (except for the instructions) identical to the baseline conditions of Experiment 1. Participants were given no training and after receiving instructions were presented with 16 test trials. At the end of the experiment, they were asked to recall the "secret."

Results and Discussion

The results of the memory check at the end of the experiment indicated that all but 1 participant remembered the rule presented as the "secret." Despite their accurate memory for the rule, participants failed to follow instructions: Across the focus on shape and focus on color conditions, participants exhibited equivalent reliance on shape ($M = 0.59$ and $M = 0.63$, not different from each other, $t\text{s} < 1$). To compare the ability to follow the rule across the conditions with the no-instruction baseline, the baseline performance in Experiments 1 and 2 was treated as the theoretical mean. The analyses indicated that the proportions of shape-based responding in the focus on shape and focus on color conditions did not differ from the baseline (where the mean shape-based responding was 59%), both one-sample $t\text{s} < 1$. Therefore, while remembering the rule throughout the experiment, participants failed to follow the rule. These findings indicate that participants have very limited (if any) deliberate control over their generalization performance.

As mentioned above, alternative interpretation of the results of Experiment 1 presupposes that participants could (a) discover a rule and (b) follow a rule. Results of Experiments 2A, 2B, and 3 suggest that 4- to 5-year-old children have trouble both discovering and following a rule, thus making it doubtful that participants' flexible generalization in Experiment 1 stemmed from rule discovery rather than associative learning. Taken together, results of the Experiments 1–3 indicate that flexible generalization can be readily acquired in the course of implicit associative learning and that flexible generalization was not under strategic control.

General Discussion

Summary of Findings

Results of the reported experiments present direct evidence that participants can learn to flexibly rely on different (arbitrarily chosen) predictors when there are multiple (arbitrarily chosen) context variables correlating with the predictor. Furthermore, as shown in Experiments 2A and 2B, this learning results in implicit rather than explicit knowledge: Participants exhibited little evidence that they were aware of the context-predictor correspondences. Finally, when explicitly instructed to rely on a given predictor, participants failed to follow instructions (although they ably remembered the instructions throughout the experiment). These results support our hypotheses, indicating that flexible generalization (a) can be acquired in the course of associative learning, (b) stems from implicit rather than explicit knowledge, and (c) does not require conceptual knowledge. The results also provide evidence that participants have limited strategic control over these generalization behaviors.

These findings support the proposed account of how attentional learning may result in highly flexible behaviors that have been traditionally taken as evidence of "smart" mechanisms underlying inductive generalization. In particular, during training, the shape of stimuli was consistently predictive in Context 1 (i.e., yellow background and upper right corner location of stimuli), whereas color was consistently predictive in Context 2 (i.e., green background and lower left location of stimuli). In the course of training, participants learned to automatically attend to shape in Context 1 and to color in Context 2. As a result, when presented with a test induction task, participants generalized on the basis of shape when stimuli appeared in Context 1 and on the basic of color when stimuli appeared in Context 2. Furthermore, this learning was implicit and it was deployed automatically (when triggered by the appropriate context). Of course, one can argue that in real life, dimensions are not neatly separated during learning as they were in our experiments. However, it is important to note that in real life, children have many more learning opportunities than they had in the reported experiments. It is also important that real-life contexts typically have significantly larger sets of intercorrelated variables. Overall, the reported findings present novel evidence for powerful learning mechanisms that may enable acquisition of conceptual knowledge and that do not require constraints stemming from conceptual knowledge.

Conceptual and Associative Accounts of Flexible Generalization Early in Development

The reported findings indicate that flexible generalization can be acquired in the course of implicit associative learning, with little evidence of participants' awareness or conscious control. These results support the associative account of flexible generalization elucidating how implicit associative and attentional learning can result in striking flexibility.

The possibility that powerful learning mechanisms may result in flexible behaviors has been raised previously (e.g., Landau, Jones, & Smith, 1992; Rogers & McClelland, 2004; Smith, Jones, & Landau, 1992, 1996). For example, Rogers and McClelland (2004) successfully simulated Macario's (1991) finding that young children differentially weigh different predictors in different contexts and Gelman and Markman's (1986) findings that young children rely on different predictors when generalizing biological and physical properties. Rogers and McClelland demonstrated that a simple (and thus not very knowledgeable) Rumelhart feed-forward network was capable of learning some of the flexible behaviors that have been often referred to as evidence of smart conceptual behaviors (e.g., Gelman & Markman, 1986; Keil, 1991; Keil et al., 1998; Murphy & Medin, 1985; see also Jaswal, 2004, for an overview). However, there has been no direct evidence indicating that children can acquire this flexibility in a course of associative learning and research reported here presents such direct evidence.

There is also indirect evidence suggesting that associative learning may underlie flexibility in naming or lexical extension tasks (e.g., Smith et al., 1996; Yoshida & Smith, 2005); however, it is unclear whether the same associative learning account can explain flexibility in naming and in other generalization tasks. If lexical extension, categorization, and induction are all variants of the same similarity-based generalization process (Sloutsky, 2003; Sloutsky & Fisher, 2004, 2005), then it is reasonable to expect that similar associative learning mechanisms may underlie flexible generalization across these tasks. Because the present study was limited to an inductive generalization task, it cannot address this issue, and additional research is needed to examine this possibility.

Findings presented here challenge the conceptual account of flexible generalization, at least in its strong version. The strong version of the conceptual account is that (a) raw associations are not sufficiently constrained to give rise to flexible generalization (cf. R. Gelman & Williams, 1998; Keil et al., 1998), (b) conceptual knowledge is necessary for exhibiting such flexibility, and (c) flexible generalization may depend on deliberate feature weighting. Recall that one weakness of the strong version is that it does not specify where conceptual knowledge comes from, under what conditions it gets deployed, and how it interacts with associative mechanisms. Findings reported here expose other potential weaknesses of the strong version of the conceptual account: Conceptual knowledge is not necessary for flexible generalization and young children have limited strategic control over their generalization behaviors.

At the same time, the reported findings are compatible with a weaker version: Although flexibility may stem from smart mechanisms, it does not have to, and it is possible to, acquire conceptual knowledge by associative means. This weaker position seems to offer a reasonable developmental proposal: People may initially acquire all knowledge by associative means, with newly acquired conceptual knowledge affecting subsequent learning. Although this possibility was dismissed by Keil et al. (1998) as an implausible "dogma" of conceptual empiricism, there is a growing body of research indicating that this possibility may represent a plausible account of conceptual development (see Rogers & McClelland, 2004; Sloutsky, 2003, for reviews). However, much theoretical and empirical work is needed to flesh out specific details of this developmental proposal.

Overall, research reported here supports the proposed account of implicit associative learning underlying flexible generalization. These findings support the idea that early in development, smart flexible behaviors stem from mundane mechanisms grounded in associative and attentional learning.

References

Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *22*, 458–475.

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.

Booth, A., Waxman, S. R., & Huang, Y. T. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology*, *41*, 491–505.

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

Fisher, A. F., & Sloutsky, V. M. (2006). Flexible attention to labels and appearances in early induction. In R. Sun & 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., & Williams, E. (1998). Enabling constraints for cognitive development and learning: Domain specificity and epigenesis. In D. Kuhn & R. Siegler, (Eds.), *Handbook*

of child psychology: Vol 2. Cognition, perception and language (5th ed., pp. 575 – 630). New York: Wiley.
Gelman, S. A., & Markman, E. (1986). Categories and induction in young children. *Cognition*, *23*, 183 – 209.
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.
Hayes, B. K., & Thompson, S. P. (2007). Causal relations and feature similarity in children's inductive reasoning. *Journal of Experimental Psychology: General*, *136*, 470 – 484.
Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411 – 422.
Jaswal, V. K. (2004). Don't believe everything you hear: Preschoolers' sensitivity to speaker intent in category induction. *Child Development*, *3*, 279 – 300.
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.
Keil, F. C. (1991). The emergence of theoretical beliefs as constraints on concepts. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and cognition* (pp. 237 – 256). Hillsdale, NJ: Erlbaum.
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.
Kersten, A. W., Goldstone, R. L., & Schaffert, A. (1998). Two competing attentional mechanisms in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1437 – 1458.
Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22 – 44.
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.
Macario, J. F. (1991). Young children's use of color in classification: Foods and canonically colored objects. *Cognitive Development*, *6*, 17 – 46.
Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289 – 316.
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. (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.
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., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
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. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, *133*, 166 – 188.
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.
Smith, L. B., Jones, S. S., & Landau, B. (1992). Count nouns, adjectives, and perceptual properties in children's novel word interpretations. *Developmental Psychology*, *28*, 273 – 289.
Smith, L. B., Jones, S. S., & Landau, B. (1996). Naming in young children: A dumb attentional mechanism? *Cognition*, *60*, 143 – 171.
Yoshida, H., & Smith, L. B. (2005). Linguistic cues enhance the learning of perceptual cues. *Psychological Science*, *16*, 90 – 95.
Young, M. E., & Wasserman, E. A. (2004). Theories of learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 161 – 182). Thousand Oaks, CA: Sage.