The Role of Linguistic Labels in Inductive Generalization

Wei (Sophia) Deng and **Vladimir M. Sloutsky**

Department of Psychology and Center for Cognitive Science, The Ohio State University

Abstract

What is the role of linguistic labels in inductive generalization? According to one approach, labels denote categories and differ from object features, whereas according to another approach, labels start out as features and may become category markers in the course of development. This issue was addressed in four experiments with 4- to 5-year-olds and adults. In Experiments 1-3, we replicated Yamauchi & Markman's findings (1998, 2000) with adults and extended the paradigm to young children. In Experiment 4, we compared effects of labels to those of highly salient visual features. Overall, results of experiments provide strong support for the idea that early in development, labels function the same way as other features, but they may become category markers in the course of development. A related finding is that whereas categorization and induction may be different processes in adults, they seem to be equivalent in young children. These results are discussed with respect to theories of development of inductive generalization.

Induction, or generalizing knowledge from known to novel, is a critical component of learning and cognition: induction enables us to apply learned knowledge to new situations. Some examples of inductive generalization include (a) inferring a property of a novel item given that a known item has this property or (b) inferring a category of a novel item given category membership of a known item. The former is referred to as *projective induction* and the latter as *categorization*. The term *induction* is often used to refer to both projective induction and categorization (Sloutsky & Fisher, 2004a).

Induction may have humble beginnings: it has been well established that induction appears early in development (Gelman & E. Markman, 1986; Mandler & McDonough, 1996; Sloutsky & Fisher, 2004a). There is also much evidence demonstrating that even early in development linguistic labels may affect inductive generalization (Gelman & E. Markman, 1986; Sloutsky & Fisher, 2004a; Sloutsky, Lo, & Fisher, 2001; Welder & Graham, 2001). However, the mechanism underlying the role of labels in early induction is hotly debated. Do labels start out as category markers (i.e., symbols denoting the category) or do they start as features and potentially become category markers in the course of development. In what follows, we consider both possibilities in greater detail.

Putative Mechanisms Underlying Effects of Labels on Generalization

Some researchers have argued that from early in development children expect linguistic labels (primarily in the form of count nouns) to mark categories (Waxman & Markow, 1995)

and facilitate inductive generalization (e.g., Gelman, 2003; Welder & Graham, 2001). According to this view, a common label suggests common category (i.e., if two items are called "dog", they are likely to belong to the same kind), whereas common category suggests that the items may share multiple properties. Therefore, when performing induction, people may first use a category label to identify the category the entity belongs to and then generalize properties of that entity to other members of the target category. For example, in a series of experiments, Gelman and E. Markman (1986) presented young children with triads consisting of a target and two test items. One test item shared the label with the target, but looked dissimilar from it, whereas the other test item looked similar to the target, but had a different label. Children were informed that one test item had a particular hidden property (e.g., "hollow bones") and the other test item had a different hidden property (e.g., "solid bones"), and asked to decide which hidden property the target had. The results indicated that children were more likely to base their inference on the common label than on perceptual similarity (but see Sloutsky & Fisher, 2004a, Experiment 4, for diverging evidence and counterarguments). This and similar findings have been interpreted as evidence that children's induction is based on category membership, which is denoted by a particular label.

There is also evidence that count nouns are most likely to guide induction than other word forms. For example, Gelman and Heyman (1999) reported that young children were more willing to generalize properties of a person from one context to another when the person was referred to by a count noun (i.e., "carrot-eater") than when referred to by a descriptive sentence (e.g., "likes to eat carrots").

These findings, however, do not lend unequivocal support to the idea that words are category markers. For example, some researchers suggested that the contribution of linguistic labels is driven by attentional rather than conceptual factors (Napolitano & Sloutsky, 2004; Sloutsky & Napolitano 2003). There is also evidence that labels contribute to the overall similarity of compared entities (Sloutsky & Lo, 1999; Sloutsky & Fisher, 2004a) and thus to both categorization and induction. In one experiment using items that had been previously used by Gelman and E. Markman (1986), Sloutsky and Fisher (2004a) demonstrated that similarity computed over labels and appearances can accurately predict young children's responses, whereas a model that assumes reliance only on labels fails to predict children's performance. Proponents of this view have also argued that early in development, labels may function like other features (e.g., shape, color, size, etc.), although they may become category markers as a result of development (Deng & Sloutsky, 2012; Sloutsky & Fisher, 2004a; Sloutsky & Lo, 1999; Sloutsky & Fisher, 2004a; Sloutsky, Lo, & Fisher, 2001; Sloutsky, 2010).

In short, according to one approach labels start out as category markers: even early in development they denote categories, and as such, they differ from other features. In contrast, labels may become category markers as a result of development, whereas early in development labels do not qualitatively differ from other features.

Experimental Distinction between Labels-as-Features and Labels-as-Category-Markers

In an attempt to distinguish between labels being features and category markers, Yamauchi and A. Markman (1998, 2000) developed an innovative paradigm potentially capable of settling the issue. The paradigm is based on the following idea. Imagine two categories A (labeled "A") and B (labeled "B"), each having four binary dimensions (e.g., Size: large vs. small, Color: black vs. white, Shape: square vs. circle, and Texture: smooth vs. rough). The prototype of Category A has all values denoted by "1" (i.e., "A", 1, 1, 1, 1) and the

prototype of Category B has all values denoted by "0" (i.e., "B", 0, 0, 0, 0). There are two inter-related generalization tasks – categorization (referred to as "classification" by the authors) and projective induction (referred to as "inference"). The goal of classification is to predict category membership (and hence the label) on the basis of presented features. For example, participants are presented with all the values for an item (e.g.,?, 0, 1, 1, 1) and have to predict category label "A" or "B". In contrast, the goal of inference is to predict a feature on the basis of category label and other presented features. For example, given an item (e.g., "A", 1,?, 1, 0), participants have to predict the value of the missing feature. A critical manipulation that could illuminate the role of labels is the "low-match" condition. For low-match inference, participants were presented with an item "A",?, 0, 1, 0, 0 (which had more features in common with the prototype of Category B, but label "A") and asked to predict the missing feature. For low-match classification, participants were presented with an item "?", 1, 0, 1, 0, 0 (which again had more features in common with the prototype of Category B) and asked to predict the missing label.

These researchers argued that if the label is just a feature then performance on the low-match classification and inference tasks should be symmetrical. However, if labels are more than features and are treated as category markers, then predicting a label when features are provided (i.e., a classification task) should elicit different performance from a task of predicting a feature when the label is provided (i.e., an inference task). Specifically, category-consistent responding should be more likely in low-match inference tasks (where participants could rely on the category label) than in low-match classification tasks (where participants had to infer the category label).

Upon finding predicted asymmetries between the two conditions, these researchers concluded that category labels differed from other features (see also Rehder, et al, 2009 for supporting eye tracking evidence). These findings have been replicated in a series of follow-up studies (see A. Markman & Ross, 2003, for a review) and have been successfully modeled (see Love, et al, 2004).

What is at Stake: Why is the Difference between Labels-as-Features and Labels-as-Category-Markers Important?

Why is understanding of the role of label early in development important? We believe that there are at least two reasons. First, this understanding is necessary for identifying the mechanism of early generalization and its change in the course of development and this knowledge, in turn, may elucidate more general principles of cognitive development. In particular, if labels function as features, they contribute to generalization in a bottom-up manner (by contributing to the featural overlap among the compared items), whereas if they are category markers they may guide the process in a top-down manner (by triggering a search for overlapping features). Each of these possibilities has long-ranging consequences for our understanding of cognitive development. If from early in development language exerts top-down influences on category learning, then even early in development, the lower-level processes (such as discrimination and generalization) are subject to top-down control. Therefore, the ability to exert top-down control, as well as cognitive and neural mechanisms that sub-serve this ability, has to exhibit early onset. Alternatively, if words acquire the ability to guide cognition in the course of development, then top-down control does not have to exhibit early onset and could be itself a product of development.

And second, the role that labels play in generalization may elucidate relationships between categorization and induction. Note that some researchers have argued that that the two tasks are functionally equivalent for adults (e.g., Anderson, 1991) and children (e.g., Sloutsky & Fisher, 2004a), whereas others have argued that the tasks are functionally different (see

Markman & Ross, 2003, for a review). If the tasks are equivalent, then representation formed in the course of classification and inference training should be equivalent as well. In contrast, if the tasks are functionally different and thus classification and inference training should result in different representations. For example, Markman and Ross (2003) presented an extensive argument regarding potential differences in representations between classification and inference training and presented evidence supporting this distinction in adults (see also Hoffman & Rehder, 2010, for eye tracking evidence; Love, et al., 2004, for a computational model). However, if early in development, labels function as features, then, classification and inference tasks should be equivalent, which, in turn suggests that extensive differences observed between classification and induction in adults are a product of development. We return to this issue in the *General Discussion* section.

Present Research

Yamauchi and Markman's paradigm has been successfully applied for examining the role of labels in adults' generalization and could be applied for examining possible developmental changes in the effect of labels on generalization. Does the asymmetry between low-match classification and low-match inference characterizing adults' performance also characterize children's performance? Finding such an asymmetry would suggest that labels play a similar role across development, indicating that even for young children labels are more than features. However, as argued above, it is possible that labels function differently across development: whereas labels may denote function as category markers in adults, they may function as perceptual features in young children. If this is the case, then unlike adults, children may exhibit symmetrical performance in low-match classification and low-match inference.

The reported experiments were designed to address these issues. In Experiments 1-3, we replicated Yamauchi & Markman' findings with adults and extended the paradigm to young children. In Experiment 4, we compared effects of labels to those of highly salient visual features.

EXPERIMENT 1

The goal of Experiment 1 was to (1) replicate Yamauchi and A. Markman's (2000) paradigm with adults and (2) examine the role of labels in early generalization by extending the paradigm to young children. Similar to Yamauchi and Markman (2000), participants learned two categories of creatures and then were given classification and inference trials, half of which were high-match and half were low-match. There were small procedural differences between the current procedure and the one used by Yamauchi and Markman (2000). Most importantly, in contrast to Yamauchi and Markman (2000), where labels were presented as a single written word, labels in Experiment 1 were presented auditorily in a carrier phrase (e.g. "This is a Flurp").

Based on Yamauchi and A. Markman's (2000) results, we expected that adult participants would make category-consistent responses in low-match inference, but not in low-match classification. This finding would be consistent with the idea that adults treat labels as category markers. Finding such an asymmetry in young children would support the idea that even for children labels are more than features.

EXPERIMENT 1A

Method

Participants—Participants were 12 adults (3 women) and 12 preschool children ($M = 56.0$ months, range 51.9-59.4 months; 7 girls). In this and all other experiments reported here, children were recruited from childcare centers, located in middle-class suburbs of Columbus, Ohio and tested in a quiet room in their preschool by a female experimenter. All adults were undergraduate students from the Ohio State University participating for course credit.

Materials—In all reported experiments reported, materials were colorful drawings of artificial creatures measuring 17.0 cm by 23.5 cm (see Figure 1). The items had five features varying in color and shape and formed two categories determined by feature values. Artificial labels ("Flurp" or "Jalet" printed above each creature) were used to refer to the categories.

As shown in Tables 1-2, the two categories have a family-resemblance structure, which is derived from two prototypes (F0 and J0) by modifying the values of one of five features (see Figure 1). For example, stimulus F1, has four features consistent with the prototype F0 and one feature (i.e., antenna) consistent with the prototype J0. The degree of similarity between a test stimulus and the prototype is defined by the number of matching features of the test stimulus to the prototype of the corresponding category (see Tables 1-2).

There were two levels of similarity: high-match and low-match. In the high-match condition, each test stimulus had four features in common with the prototype of the corresponding category and one feature in common with the prototype of the contrasting category. In the low-match condition, each test stimulus had two features in common with the prototype of the corresponding category and three features in common with the prototype of the contrasting category.

Design and Procedure—All experiments reported here had a two (Test Condition: Classification vs. Inference) by two (Feature Match: High vs. Low) within-subjects design, and the procedure consisted of two phases, training and testing. The experiments were administered on a 17-inch computer monitor and controlled by E-prime 2.0 software. Classification and Inference test trials were presented in blocks and the order of the blocks was counterbalanced. The order of test trials within each block was randomized for each participant.

All experiments started with the training phase. At the beginning of training both adult and children were instructed that there were two groups of creatures "Flurps" and "Jalets." They were then presented with creatures (one at a time, with 4000 ms per item), each accompanied by a category label presented auditorily in a carrier phrase (e.g., "This is a Flurp"). The carrier phrase was pre-recorded and presented by a computer. The phrase had the same onset as the beginning of the trial, with the total duration of approximately 1800 ms. There were 36 training trails (i.e., 18 Flurps and 18 Jalets) and this part of the experiment lasted for approximately 3-4 minutes. No participant response was required during this phase.

The testing phase, consisting of 92 trials (12 with feedback and 80 without feedback), was administered immediately after the training phase (see Figure 2 for examples of testing trials). Half of these trials were Classification and half were Inference, with each trial presented in a self-paced manner. The Classification and Inference testing conditions differed in what participants had to predict. On Classification trials, participants predicted

the label of an item, given information about all five features (they were instructed that they would be presented with creatures and they would need to decide whether the creature was a Flurp or a Jalet). On Inference trials, participants predicted a missing (i.e., covered) feature, given the other four features and the label (they were instructed that they would be presented with creatures with a covered body part and they will have to decide which body part is under the cover). For both children and adults, this part of the experiment lasted for approximately 14-15 minutes.

The procedures were identical for both adult and child participants except for the way the instructions and test questions were presented and the data were recorded. Adults read instructions and the test questions on the computer screen and responded by pressing an appropriate key on the keyboard. For children, all instructions and questions were presented by a female experimenter and she recorded children's verbal responses by pressing the keyboard.

To familiarize participants with the testing task, yes/no feedback was given on 12 test trials -- the first six test trials of Classification and Induction testing condition (each of these were high-match trials). In this and other experiments reported here, children were above 85.2% accuracy on these trials and adults were above 68.5%, all above chance, $ps < .05$. No feedback was given on the remaining 80 testing trials (40 in each testing condition, half high-match and half low-match) and only these trials were used in the reported analyses. The proportion of responses consistent with the category from which the exemplar was derived (called "category-accordance responses" by Yamauchi & A. Markman, 2000) was the dependent variable.

In addition, a memory check was administered after the main experiment to examine whether participants remembered two categories after completing all the tasks. There were five memory check trials, with participants being presented with stimuli randomly generated from the training structure (see Table 1). On each trial, participants were asked to recall the corresponding label of each stimulus. Children and adults exhibited memory accuracy of 91.7% and 78.0% respectively. One adult answered fewer than three out five memory check questions correctly and these data were excluded from the analysis.

Results and Discussion

The main results are presented in Figure 3. As can be seen in the figure, adults exhibited equivalent performance in the high-match condition (i.e., no differences between Classification and Inference), whereas there was a marked difference in the low-match condition. In particular, adults were more likely to produce category-consistent responding in the low-match Inference condition than in the low-match Classification condition. In contrast, young children produced high levels of category-consistent responding only in the high-match conditions, whereas this was not the case in the low-match conditions.

Note that all experiments that involved different age groups (Experiments 1-3) revealed a significant 3-way (i.e., Age × Testing Type × Feature Match) interaction (all $Fs > 4.1$, $ps < .06$, $\eta_p^2 > .176$). To interpret the interaction, we conducted separate 2 (Testing Type: Classification vs. Induction) by 2 (Feature Match: High vs. Low) within-subjects ANOVAs for each age level.

For adults, there was a significant testing type by feature match interaction, $F(1, 10) = 24.63$, $MSE = 0.32$, $p = .001$, $\eta_p^2 = .711$. A paired-samples t-test indicated that in the high-match condition there were no differences between Inference and Classification $t(10) = 0.30$, $p = .772$, whereas in the low-match condition participants were more likely to make category-

consistent responses in the Inference than in the Classification condition, $t(10) = 3.89$, $p = .003$, $d = 1.71$.

For children, there was a main effect of feature match, $F(1, 11) = 43.56$, $MSE = 1.33$, $p = .001$, $\eta_p^2 = .798$, with participants being more likely to provide category-consistent responses in the high-match than in the low match condition. There was also a main effect of testing type (there were more category consistent responses in the Classification than in the Inference condition), $F(1, 11) = 14.77$, $MSE = 0.16$, $p = .003$, $\eta_p^2 = 0.573$, which was different from adults. This main effect may reflect the advantage of Training-Testing correspondence (recall that participants in this experiment were trained by classification) and we will come back to this issue in Experiments 1B and 3.

There were several important differences between children and adults. First, in contrast to adults, for children there was no significant interaction between testing type and feature match, $F(1, 11) = 0.02$, $MSE = 0.00$, $p = .90$. Second, unlike adults who were above chance in relying on category information in low-match inference, one-sample $t(10) = 2.66$, $p = .024$, $d = 0.80$, young children performed significantly below chance in relying on the label to predict missing features; they relied instead on the overall similarity, one-sample $t(11) = 4.47$, $p = .001$, $d = 1.29$. And finally, in contrast to adults, children's performance in low-match Inference did not exceed low-match Classification. In fact, the opposite was true – children were somewhat more likely to generate category-consistent responses in low-match Classification than in low-match Inference, paired-sample $t(11) = 2.32$, $p = .041$, $d = 0.85$.

These results extended those of Yamauchi & Markman (2000), suggesting that labels may play a different role for adults and children. In particular, similar to Yamauchi & Markman (2000), adults treated labels differently from other features: in low-match inference they relied primarily on labels. In contrast, children relied on the overall similarity: the proportion of category-consistent responses in low-match Inference was below chance and did not exceed that in low-match Classification, which provided little evidence that for young children labels are category markers.

In sum, the asymmetry between low-match Inference and low-match Classification in adults suggests that for adults labels are processed differently from other features. In contrast, children's tendency to rely on the overall similarity in both low-match Inference and Classification indicates that young children did not treat category labels differently from other features.

Note that Experiment 1A used only Classification training, whereas participants were tested in both classification and inference. To ensure that the observed effects are not specific to a particular training condition used in Experiment 1A, we conducted Experiment 1B, in which participants were given Inference training.

EXPERIMENT 1B

Method

Participants—Participants were 11 adults (6 women) and 12 preschool children ($M = 54.2$ months, range 50.9-56.9 months; 4 girls).

Materials, Design, and Procedure—The materials, design and procedure were similar to those in Experiments 1A, with several differences. First, in contrast to Experiment 1A, participants were given Inference training. Before training, both adult and child participants were instructed that there were two groups of creatures, with members of each group having something special inside its body. One group of creatures was said to have a flurp inside its

body, whereas the other was said to have a jalet inside its body. On each training trial, one creature was presented and participants were told: "This one has a flurp (or jalet) inside its body."

There were also some differences in testing. In the Classification task, participants were asked to predict the feature label of an item, given information about all five features (e.g., Is it from the group with a flurp or with a jalet inside the body?). In the Inference task, they were asked to predict the value of one of five features, given the other four features and the label.

Similar to Experiment 1A, a memory check was administered after the main experiment with all child and adult participants exhibiting memory accuracy of 83.3% and 72.7 % respectively. Two adults answered fewer than three out five memory check questions correctly, and these data were excluded from the analysis.

Results and Discussion

The main results are shown in Figure 4. Patterns of responding were very similar to those in Experiment 1A. In adults, there were no differences between Classification and Inference in the high-match condition, whereas there was a marked difference in the low-match condition. In particular adults were more likely to produce category-consistent responding in low-match Inference than in low-match Classification. In contrast, young children were likely to produce high levels of category-consistent responding only in the high-match condition, but not in the low-match condition.

The data were submitted to two separate 2 (Testing Type: Classification vs. Inference) by 2 (Feature Match: High vs. Low) within-subjects ANOVAs. For adults, there was a interaction, $F(1, 8) = 9.92$, $MSE = 0.26$, $p = .014$, $\eta_p^2 = .553$. Specifically, in the low-match condition, participants were more likely to provide category-consistent responses in Inference than in Classification, paired-samples $t(8) = 2.79$, $p = .023$, $d = 1.62$, which was not the case for the high-match condition, paired-samples $t(8) = 0.53$, $p = .613$.

For children, there was a main effect of feature match, $F(1, 11) = 86.84$, $MSE = 2.50$, $p = .001$, $\eta_p^2 = .888$, with participants being more likely to provide category-consistent responses in the high-match than in the low match condition. There was also a marginally significant interaction, $F(1, 11) = 4.22$, $MSE = 0.04$, $p = .064$, $\eta_p^2 = .277$, with participants showing equivalent performance in the high-match conditions, paired-samples $t(11) = 0.11$, $p = .913$, but higher performance in the low-match inference than low-match classification, paired-samples $t(11) = 2.79$, $p = .018$, $d = 0.91$. However, unlike adults who were above chance in relying on category information in low-match inference, one sample $t(8) = 2.57$, $p = .032$, $d = 0.84$, young children were not significantly different from chance, $p = .185$.

Furthermore, a comparison with Experiment 1A (where low-match Classification performance was somewhat higher than low-match Inference performance) suggests that these differences may be indeed training-specific and we address this issue directly in Experiment 3. However, critically, in contrast to adults in Experiments 1A and 1B, low-match inference in young children did not exceed chance performance.

These results extend those of Experiment 1A. Adults again showed asymmetric performance in Classification and Inference tasks, with performance in low-match induction being above chance. Therefore, when a label indicated one prototype and the majority of perceptual features indicated another prototype, adults tended to rely on the label. In contrast, young children exhibited little evidence of relying on the label. In addition, unlike adults, children

exhibited evidence of training-specific effects, and we will further examine these effects in Experiments 2-3.

The goal of Experiment 2 was to examine the generality of effects observed in Experiment 1. In Experiment 1 labels were novel count nouns and only adults, but not children, exhibited evidence of consistent reliance on labels in the low-match induction. Would these effects hold for different labels? Would children more readily rely on labels in low-match induction if verbal information is familiar? To answer this question, we conducted Experiment 2A, in which participants were presented with familiar count nouns. To avoid providing children with information they know to be false (e.g., naming the present robot-like stimuli as "bear" or "rabbit"), we used more general count nouns "friendly pet" and "wild animal".

In Experiment 2B, we further examined the generality of effects by replacing count nouns with descriptors of the habitat (e.g., "lives in the forest" vs. "lives in in the sea"). Experiment 2B introduces an important additional control: if adults continue relying on verbal information, this would indicate that for adults verbal information does not have to be presented in the form of count noun to be treated as a category marker.

EXPERIMENTS 2A and 2B

Except for the verbal information given to participants, experiments 2A and 2B were isomorphic to Experiments 1A and 1B (i.e., in Experiment 2A participants were presented with Classification training, whereas in Experiment 2B they were presented with Induction training). We therefore describe these experiments (including relevant procedural information) in Table 3.

The main results of Experiments 2A and 2B are shown in Figures 5-6. The overall pattern is strikingly similar to that in Experiment 1. In adults, there were no differences between Classification and Inference in the high-match condition, whereas there was a marked difference in the low-match condition. In particular, adults were more likely to produce category-consistent responding in low-match Inference than in low-match Classification. In contrast, young children were likely to produce high levels of category-consistent responding only in the high-match condition, but not in the low-match condition.

The testing data of Experiment 2A were analyzed with two separate 2 (Testing Type: Classification vs. Inference) by 2 (Feature Match: High vs. Low) within-subjects ANOVAs. For adults, there was a significant interaction between testing type and feature match, $F(1, 12) = 36.23$, $MSE = 0.48$, $p = .001$, $\eta_p^2 = .751$. Similar to previous results, adults were more likely to provide category-consistent responses in low-match Inference than in low-match Classification, paired-samples $t(12) = 4.78$, $p = .001$, $d = 2.15$. At the same time, there was no significant difference between the high-match conditions, paired-samples $t(12) = 0.76$, $p = .461$.

For children, there was a significant main effect of feature match, $F(1, 8) = 35.66$, $MSE = 2.15$, $p = .001$, $\eta_p^2 = .817$, with children being more likely to provide category-consistent responses in the high-match than in the low-match condition. At the same time, neither the main effect of testing nor the interaction approached significance, $ps > .28$. In addition, unlike adults who were above chance in relying on category information in low-match inference, one sample $t(12) = 4.31$, $p = .001$, $d = 1.19$, young children were marginally below chance, $t(8) = 1.88$, $p = .097$, $d = 0.63$.

Testing data of Experiment 2B were also analyzed with two separate 2 (Testing Type: Classification vs. Inference) by 2 (Feature Match: High vs. Low) between-subjects

ANOVAs. For adults, there was a significant testing type by feature match interaction, $F(1, 10) = 37.56$, $MSE = 0.38$, $p = .001$, $\eta_p^2 = .790$. Similar to other experiments, adult participants were more likely to give category-consistent responses in low-match Inference than in low-match Classification, paired-samples $t(10) = 4.08$, $p = .002$, $d = 1.79$, which was not the case for the high-match condition.

For children, similar to the previous experiments there was a significant main effect of feature match, $F(1, 8) = 24.63$, $MSE = 1.50$, $p = .001$, $\eta_p^2 = .755$. Specifically, children were more likely to provide category-consistent responses in the high-match than in the low-match condition. In addition, similar to Experiment 1B, there was a significant testing type by feature match interaction, $F(1, 8) = 15.75$, $MSE = 0.07$, $p = .004$, $\eta_p^2 = .663$. Similar to adults, in the low-match condition, they were more likely to provide category-consistent responses in the Inference than in the Classification condition, paired-samples $t(8) = 2.51$, $p = .036$, $d = 1.13$, which was not the case in the high-match condition, paired-samples $t(8) = 0.92$, $p = .38$. However, unlike adults who were above chance in relying on category information in low-match inference, one sample $t(10) = 2.73$, $p = .021$, $d = 0.83$, young children were not different from chance, $p = .56$.

Overall, results of Experiment 2 replicated and further extended results of Experiment 1. Most critically, differences between children and adults persisted across different ways of presenting verbal information: similar to Experiment 1, adults consistently relied on verbal information, whereas children did not. Furthermore, neither familiar count nouns (e.g., "friendly pet") used in Experiment 2A increased children's reliance on labels, nor descriptors (e.g., "lives in the forest") used in Experiment 2B attenuated adults' reliance on labels. Therefore, whereas young children exhibited a broad tendency to rely on the overall similarity, regardless of the familiarity or the form of the label, adults tended to rely on labels, also regardless of the familiarity or the form of the label.

Finally, a closer examination of the difference between children's performance in Classification and Induction tasks in Experiments 1A and 2A versus 1B and 2B suggests that child participants may be affected by different training procedures. In Experiments 1A and 2A, children were trained by classification and during testing they exhibited somewhat better performance in low-match Classification than in low-match Inference. In contrast, in Experiments 1B and 2B children were given inference training and during testing they exhibited somewhat better performance in low-match Inference than in low-match Classification. To equate effects of training, we conducted Experiment 3, in which all participants received both Classification and Inference training.

EXPERIMENT 3

Method

Participants—Participants were 26 adults (8 women) and 20 preschool children ($M = 55.6$ months, range 48.3-70.0 months; 13 girls). One child participant was interrupted during the experiment and these data were excluded from the analysis.

Materials, Design, and Procedure—The experiment had two between-subjects training conditions: (1) classification-label (CL) and inference-descriptor (ID) and (2) classification-descriptor (CD) and inference-label (IL). The orders of CL vs. ID and CD vs. IL were counterbalanced. Across the conditions, participants were presented with the same visual stimuli used in previous experiments. Participants were trained with 24 classification trials and 24 inference trials. The corresponding testing trials (46 classification trials and 46 induction trials for each training condition) were administered after all training trials. Similar to previous experiments, yes/no feedback was given on 12 test trials -- the first six

test trials of Classification and Inference testing (each of these were high-match trials). No feedback was given on the remaining 80 testing trials (40 in each testing condition, half high-match and half low-match) and only these trials were used in the reported analyses.

The CL training trials were identical to those of Experiment 1A and the ID trials were identical to those of Experiment 2B. For example, participants in the CL-ID condition were trained by classification with labels (e.g., "This is a Flurp.") and inference with descriptions (e.g., "This one lives in the forest."). They were then presented with test trials of both Classification (to predict the label given all other features) and Inference (to predict a missing feature given other four features and the living place).

The procedure of CD and IL training condition was similar, but the CD trials were identical to those of Experiment 2A and the IL trials were identical to those of Experiment 1B. Specifically, participants in the CD-IL condition were trained by classification with descriptors (e.g., "This one lives in the forest") and inference with feature labels (e.g., "This one has a jalet inside its body"). During testing, participants were presented with test trials of both Classification (to predict the descriptor given all other features) and Inference (to predict a missing feature given other four features and the feature label).

Similar to previous experiments, a memory check was administered after the main experiment and child and adult participants exhibited memory accuracy of 88.4% and 72.8% respectively. Five adults and one child answered fewer than six out of ten memory check questions correctly and these data were excluded from the analysis.

Results and Discussion

Since for both children and adults, the effect of training condition (i.e., CL-ID vs. CD-IL) was not significant and did not interact with Testing Type or Feature Match, all $ps > .16$, the data were collapsed across two conditions (see Figure 7).

For adults, there was a significant testing type by feature match interaction, $F(1, 20) = 39.09$, $MSE = 0.56$, $p = .001$, $\eta_p^2 = .662$. Specifically, in the low-match condition, participants were more likely to provide category-consistent responses in the Inference than in the Classification condition, paired-samples $t(20) = 4.50$, $p = .001$, $d = 1.26$, which was not the case in the high-match condition, paired-samples $t(20) = 1.79$, $p = .089$.

For children, there was a main effect of feature match, $F(1, 17) = 603.60$, $MSE = 2.94$, $p = .001$, $\eta_p^2 = .973$, with participants being more likely to provide category-consistent responses in the high-match than in the low match condition, but in contrast to adults, there was no significant interaction between testing type and feature match, $p > .56$. Furthermore, unlike adults who were above chance in relying on category information in low-match inference, one-sample $t(20) = 2.80$, $p = .011$, $d = 0.61$, young children performed significantly below chance in relying on the label in low-match inference; they relied instead on the overall similarity, one-sample $t(17) = 2.54$, $p = .021$, $d = 0.60$.

Critically, when participants received a combination of Classification and Inference training, they exhibited the same patterns as when they received only one type of training (Experiments 1-2). Specifically, adults relied on labels, whereas children relied on the overall similarity.

To further examine differences between children and adults in their reliance on labels, we analyzed individual patterns of responding of children and adults across all experiments that used count nouns (i.e., Experiments 1, 2A, and 3). Participants who made at least 13 out of 20 testing trials category-consistent responses (above chance, binomial $p = .07$) in the high-

match inference were selected for the analysis of the response pattern in the low-match induction. Those providing category-based (or label-bases) responses on 13 out of 20 testing trials were classified as category-based responders, whereas those providing at least 13 out 20 responses based on the overall similarity were classified as feature-based (or similarity-based) responders. The rest were classified as mixed responders. The proportions of label-based, similarity-based and mixed responders are presented in Table 4. Critically, while the majority of adults (i.e. over 80%) were consistent label-based responders, only 6% of children were. Instead children were equally split between similarity-based and mixed responders, and the pattern found in children differed significantly from that found in adults, $\chi^2(1, 69) = 41.3$, $p < 0.0001$.

Overall across Experiments 1-3 when the label (or descriptor) was pitted against appearance similarity, adults tended to rely on labels when making inductive generalization, which was not the case for young children. Is it possible that children's failure to rely on labels stemmed from fatigue resulting from multiple trials? Although this possibility seemed unlikely because the overall experiment lasted for less than 20 minutes, we deemed it necessary to compare children's performance for the first and the second halves of each experiment. If failure to rely on labels stemmed from fatigue, then reliance on labels should be significantly higher in the first part of the experiment than in the second half. Our analyses of Experiments 1A, 1B, 2A, 2B, 3A and 3B indicated that this was not the case: in none of the experiments reliance on labels decreased significantly from the first half of the experiment to the second half, all Bonferroni adjusted ps > .7.

Taken together, results of Experiments 1-3 suggest that whereas for adults labels could be category markers, for young children, they are no more than features. Although these results are informative, it could be argued that young children do understand that the labels are category markers, but the results reflect inability of young children to rely on a single feature when this feature is pitted against multiple features. The goal of Experiment 4 is to address this possibility.

EXPERIMENT 4

In Experiments 1-3, adult and child participants consistently showed different patterns of responding. The consistent reliance on labels in low-match induction suggests that adults treated labels differently from other features, perhaps as category markers, which was not the case for young children.

In contrast to adults, children's induction was similarity-based: across all the experiments they relied on featural overlap rather than on the label (or descriptor) when performing induction. However, while Experiments 1-3 present evidence that labels are not category markers for young children, they do not eliminate one important alternative. It is possible that children do understand that the labels are category markers, but they miss the ability to rely on a single feature, especially when this single feature is pitted against multiple features, such as in the low-match inference. While still advancing our understanding of the role of labels in early induction, this latter explanation does not eliminate the possibility that labels are category markers.

To address this issue in Experiment 4, we made one of the non-linguistic features more salient than any other feature or the label (see *Method* section for explanation of how this was ascertained). To achieve this goal, the creatures' head was made to move. One type of head motion was consistent with Category 1 and another with Category 2. The rest of the procedure was similar to that in Experiments 1-3.

Method

Participants—Participants were 12 preschool children (M = 53.9 months, range 49.6-59.3 months; 5 girls).

Materials, Design, and Procedure—The visual stimuli were identical to previous experiments except for the following differences. First, to set up a proper competition between the category information (which did not vary across the exemplars), and a feature, the value of one feature (the head) was also fixed within each category (see Tables 5-6).

And second, to make the fixed feature highly salient, the head was animated using Macromedia Flash MX software. The head of one category was pink and moved up and down; whereas, for the other category, the head was blue and moved sideways. When asked after the experiment what they noticed about the items all children mentioned the moving head. Two children also mentioned the category label. Therefore, it was concluded that the moving head was more salient than any other feature or the label.

As a result, the learning structure was changed for the part of the head and is shown in Table 5. Similar to previous experiments, Experiment 4 consisted of two phases, training and testing. Two levels of feature match between the test item and the prototype of the corresponding category were used, high and low (see Table 6). As shown in the table, in the low-match condition there was only one feature (i.e., the moving head) in common with the respective prototype; whereas, in the high-match condition there were four such features. The critical condition was low-match inference where only the moving head was in common with the prototype of the corresponding category; whereas, three features and the category information (i.e., descriptors of the habitat: "lives in the forest" vs. "lives in in the sea") were common with the prototype of the contrasting category. Therefore, if participants rely on multiple features they should infer the feature from the contrasting category, thus exhibiting a high level of category-based responding. In contrast, if they rely on the highly salient moving head, they should exhibit a low level of category-based responding. In all other conditions, there was no conflict between the category information and the moving head, and thus reliance on the moving head would result in a high level of category-based responding (see Table 6).

The overall procedure was similar to Experiment 2B (i.e., the participants received inference training) with the following difference: both the training and testing phases were shortened, with the procedure including 24 training trials and 44 testing trials (similar to previous experiments, the first 12 testing trials were high-match trails accompanied by feedback and these were not included analyses). We shortened the procedure to eliminate the possibility that failure to rely on labels in Experiments 1-3 did stemmed from fatigue resulting from multiple testing trials. Recall that comparison of the first half of and the second half of testing in Experiments 1-3 undermined this possibility, shortening the procedure allowed us to address this issue directly. Similar to previous experiments, a memory check was administered after the main experiment with all participants exhibiting high memory accuracy (93%), with no participant answering correctly fewer than three out of five memory check questions correctly.

Results and Discussion

The main results of Experiment 4 are shown in Figure 8. The data were analyzed with a 2 (Testing Type: Classification vs. Inference) by 2 (Feature Match: High vs. Low) within-subjects ANOVA. Most importantly, there was a significant testing type by feature match interaction, $F(1, 11) = 129.76$, $MSE = 1.51$, $p = .001$, $\eta_p^2 = .922$. In the high-match condition, there was no difference between Classification and Inference, paired-samples

$t(11) = 0.84$, $p = .417$, whereas in the low-match condition, participants were more likely to make category-consistent responses in Classification than in Inference condition, paired-samples $t(11) = 19.90$, $p = .001$, $d = 8.36$. Most importantly, when the category descriptor, appearance similarity, and the moving head all indicated the same category (i.e., high-match Inference), children were above chance in providing category-consistent responses, one sample $t(11) = 11.86$, $p = .001$, $d = 3.47$. In contrast, when the descriptor denoting a category was pitted against the salient feature (i.e., in low-match inference), children performed significantly above chance in relying on the moving head to infer missing features, one-sample $t(11) = 13.68$, $p = .001$, $d = 3.90$. Therefore, while in Experiments 1-3, there was no evidence that children rely on label or category information in the low-match induction, they had no difficulty relying on a single highly salient feature in the current experiment.

Unlike in other experiments reported here, in Experiment 4 children relied on a single feature (i.e., the moving head) rather than on multiple features. Therefore, children's failure to rely on labels in Experiments 1-3 is unlikely to stem from their inability to rely on a single feature when it is pitted against multiple features. Although no difference was found between labels and descriptors in Experiments 1B and 2B, it could be argued that results of Experiment 4 would be different had we used count nouns. We therefore replicated Experiment 4 with 13 additional children who were given category labels presented as count nouns (e.g., "friendly pet" or "wild creature") instead of descriptors. The results of this experiment were equivalent to those of Experiment 4: children were below chance in the low-match inference (31% of category-consistent responses, $t(12) = 8.40$, $p = .001$, $d = 2.33$) and they were above chance in the other three conditions (category consistent responses ranged from 77% to 84%, all ps = .001, ds > 2.6). Therefore, even when the highly salient moving head was pitted against a count noun, young children relied on the moving head to predict missing features. Furthermore, as was shown in a recent study, young children relied on the moving head when labels were presented as either novel count nouns, such as flurp vs. jalet or as familiar count nouns, such as carrot-eater vs. meat-eater (Deng & Sloutsky, 2012).

Overall, across all the reported experiments children failed to rely either on the label or the category descriptor (Experiments 1-3), whereas they relied on a salient perceptual feature (Experiment 4). In contrast, adults tended to rely on label (or category information) and not on the overall similarity. These results point to important developmental differences in the role of labels in generalization: Whereas adults are likely to treat labels as category markers, there is little evidence that for young children linguistic labels are more than features.

GENERAL DISCUSSION

The reported research presented six experiments designed to examine the role of labels in early generalization and changes in this role in the course of development. To achieve this goal, we built on the paradigm pioneered by Yamauchi & Markman (1998, 2000). Several major findings stem from the reported experiments.

First, in all experiments adults relied on category labels when the label was pitted against appearance similarity, which was not the case for young children. In contrast, under no condition did young children exhibit sole reliance on labels in their induction. These effects could not have stemmed from poor category learning or poor memory for labels: across all experiments children were exceedingly accurate on memory checks, exhibiting memory accuracy of 88%. At the same time, when a highly salient visual feature was introduced (Experiment 4), young children did perform induction by relying on this feature. Taken together these results offer little evidence that labels are category markers for young

children, whereas labels may become category markers in the course of development. As we discuss below, these results have important implication for understanding of the development and mechanism of generalization and for theories of categorization.

Labels and the Mechanism of Generalization

Although researchers agree that from early in development people are capable of performing inductive generalization, the underlying mechanism is a matter of debate. Some have argued that induction is category-based in that when performing induction, people access the category of items in question (see Gelman, 2003 Murphy, 2002, for reviews). Others have presented an alternative argument that, at least early in development, induction is driven by similarity of compared entities rather than by a common category membership (see Murphy, 2002; Sloutsky, 2010, for reviews). However, under typical circumstances, it is difficult to distinguish between these possibilities. There have been at least two proposals as to how such a distinction could be made.

First, there is an argument that category-based and similarity based induction may result in different memory traces for studied items, with similarity-based induction resulting in more detailed verbatim memories and category-based induction resulting in less detailed gist-type memory (Sloutsky & Fisher, 2004b; Fisher & Sloutsky, 2005, but see Wilburn & Feeney, 2008 and Sloutsky, 2008). This argument has resulted in a set of studies demonstrating that (a) young children (who presumably perform similarity-based induction) retain more accurate memories of the studied items than adults (who presumably perform category-based induction) and (b) training young children to perform category-based induction attenuates their memory accuracy to the level of adults (Fisher & Sloutsky, 2005; Sloutsky & Fisher, 2004b).

The second idea is to experimentally dissociate category membership and similarity. If induction is category-based, it should follow category information, whereas if it is similarity-based, it should follow similarity information. In one such study (Sloutsky, et al, 2007a), 4-5-year-olds learned two rule-based categories, with similarity not being predictive of category membership. Upon learning the rule-based categories, participants were presented with a set of induction trials in which they could rely on either category information or similarity information. Despite the fact that children successfully acquired the categories and retained this knowledge throughout the experiment, their induction was similarity-based (see also Gelman & Waxman, 2007 and Sloutsky, Kloos, & Fisher, 2007b, for further discussion; Griffiths, Hayes, & Newell, 2012, for cases of non-category-based induction in adults).

The third idea is to examine the role of labels in induction: finding evidence that category labels are different from other features and they guide inductive inference would support the idea of category-based induction. The current work finds such evidence for adults, but not for young children. And if early in development labels are indeed features rather than category markers, then the ability of young children to perform category-based induction is highly questionable. Therefore, current research, in conjunction with earlier reported findings (Fisher & Sloutsky, 2005; Sloutsky & Fisher, 2004a; Sloutsky, et al., 2001; Sloutsky, et al 2007a), presents further evidence that early induction is similarity-based, but it may become category-based in the course of development.

Language and Cognition: Are Labels Features or Category Markers

The question of how labels affect generalization is critically important for understanding the mechanism of generalization, but it has broader implications for understanding of the role of language in cognition and cognitive development. At the computational level of analysis

(Marr, 1982), the labels-as category-markers approach assumes that words are not merely a part of stimulus input, but rather fulfill the role of supervisory signals directing and guiding learning. Thus, if two discriminable items share the same count noun (e.g., both are called "a dax"), the name serves as a top-down signal that the items are equivalent in some way (cf. Gliga, Volein, & Csibra, 2010). In contrast, if words are like any other perceptual feature, they are part of perceptual input contributing to the overall category structure.

Each of these possibilities presumes a different mechanism and dedicated neural architecture, and a different developmental trajectory. Distinguishing among them and understanding the mechanisms underlying the effect of words on category learning is critically important for understanding cognitive development. If from early in development words are supervisory signals, then top-down effects have to play a significant role in early cognitive development. Perhaps the most important implication is that at both the cognitive and the neural levels, the lower-level processes (such as discrimination and generalization) are subject to top-down control. Alternatively, if words become supervisory signals in the course of development, then top-down control does not have to exhibit early onset and could be itself the product of development. Therefore, understanding the role of labels in generalization has implications for most fundamental aspects of cognitive development as well as for understanding of the interaction between language and cognition.

Although additional research is needed, present research indicates that even at 4-5 years of age, labels function more like features than category markers. When and how do labels become category markers? It seems that there are two possible ways of approaching these questions, pessimistic and optimistic. According to a pessimistic view, the presented results severely undermine the claim that labels do have the special status for young children, while not providing conclusive evidence for the special status of the label even in adults. Indeed, results of Experiment 4 suggest that even overwhelming reliance on the label may not be indicative of the fact that the label is a category marker: although children in Experiment 4 overwhelmingly relied on the moving head, we cannot envision a claim that the moving head is a category marker. However, if one assumes a more optimistic view -- that labels eventually become category markers for adults – then a theoretical and empirical challenge is to establish the developmental mechanism of this process.

Labels-as-Features vs. Labels-as-Category-Markers: Implications for the Relationship between Categorization and Induction

Recall that much evidence suggests that inference and classification learning are not equivalent for adults, who form different representations in the course of classification and inference training. In particular, under most conditions, categorization training results in the discovery of the features that distinguish among the contrasting categories, whereas inference training results in the discovery of features that are most common in the given category and of inter-feature relations (see Markman & Ross, 2003, for a review; Chin-Parker & Ross, 2004; Sakamoto & Love, 2010; Yamauchi & Markman, 1998; see also Love, et al, 2004, for computational modeling).

These differences in representation also manifest themselves in differences in learning rates: for most family-resemblance categories inference training is faster than classification (see Markman & Ross, 2003, for a review), whereas the opposite is true for non-linearly-separable categories (Love, et al., 2004). At the same time, little is known about the development of these differences. Current research reveals no systematic difference between early categorization and induction (sometimes categorization exceeded induction, sometimes the opposite was the case, and sometimes they were statistically equivalent). These findings, in conjunction with evidence that early in development labels function as features, suggest that early categorization and induction could be functionally equivalent. Although we did

not examine how children represent categories in the course of classification and inference training (this question is for future research), the current findings allow us to predict that children may form equivalent representations in the course of classification and inference training. If the task does include an exceedingly salient feature (as Experiment 4), the current results suggest that even in inference training children will learn this highly salient and diagnostic feature rather than the interrelationships among the features (as was the case in previous research with adults). Therefore, the profound differences between classification and inference learning found in adults may not be a fixed property of the tasks; instead these differences may emerge in the course of development.

Markman and Ross (2003) argued that the differences between categorization and induction pose a challenge to many existing theories of categorization. It seems that the idea that the distinction may emerge in the course of development adds to this challenge.

From Features to Markers? The Changing Role of Category Label in Generalization

If we accept that labels do become category markers later in development, it is reasonable to ask: what changes in the course of development? One answer can be provided at the computational level: for example a model of category learning SUSTAIN (Love, et al., 2004) introduces a parameter of "category focus" (λ) that governs how much attention is placed on the category label. Depending on the value of the parameter, the label could be similar to other features or it could be a category marker. This parameter change offers a mechanistic way of understanding development, but it is also important to understand what triggers this change.

One possible idea that was discussed elsewhere (e.g., Sloutsky, 2010) is that the contribution of labels to categorization and category learning hinges on (a) the ability to process cross-modal information and (b) the ability to attend selectively. Although neither of these abilities might be sufficient, both seem to be necessary, and both may be relatively immature early in development.

First, there is a growing body of evidence that auditory input may affect attention allocated to corresponding visual input (Napolitano & Sloutsky, 2004; Robinson & Sloutsky, 2004; Sloutsky & Napolitano, 2003; Sloutsky & Robinson, 2008). In particular, linguistic labels may strongly interfere with visual processing in pre-linguistic infants, but these interference effects may weaken when children start acquiring language (Sloutsky & Robinson, 2008, see also Robinson & Sloutsky, 2007a; 2007b). Given that category learning depends critically on visual processing, labels may hinder learning of new categories in both infants and young children. Therefore, the ability to efficiently process and integrate auditory and visual input appears to be a critical (yet by no means sufficient) step in labels becoming category markers.

And second, in order for a label to be used as a category marker, participants should be able to selectively attend to relevant information and ignore irrelevant information. However, research published over the last 30 years suggests that young children miss this ability (see Dempster & Corkill, 1999; Hanania & Smith, 2010; Lane & Pearson, 1982, for comprehensive reviews). These difficulties have been linked to the fact that the regions sub-serving selectivity (most importantly, the prefrontal cortex) undergo protracted development (Bunge & Zelazo, 2006; Davidson, Amso, Anderson, & Diamond, 2006) and exhibit critical immaturities throughout infancy and preschool years. In short, the ability to integrate cross-modal information and to attend selectively seem to be necessary steps for labels to become category markers.

To summarize, results reported here present evidence that labels may function differently across development: whereas labels are likely to function as features early in development, they may become category markers later in development. Although the ability to integrate cross-modal information and to attend selectively could be necessary steps in the changing role of labels, precise mechanisms underlying this transition remain unknown. Therefore, much research is needed to understand why, how, and when labels become category markers.

Conclusion

Current research presents extensive evidence that (a) early in development labels are features rather than category markers, but they may become category markers in the course of development and (b) categorization and induction are likely to be equivalent in children, but not in adults. The remaining challenge is to understand why and how these transitions take place.

Acknowledgments

References

1. Anderson JR. The adaptive nature of human categorization. Psychological Review. 1991; 98:409–429.

2. Bunge SA, Zelazo PD. Brain-based account of the development of rule use in childhood. Current Directions in Psychological Science. 2006; 15:118–121.

3. Chin-Parker S, Ross BH. Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. Journal of Experimental Psychology: Learning, Memory and Cognition. 2004; 30:216–226.

4. Davidson MC, Amso D, Anderson LC, Diamond A. Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. Neuropsychologia. 2006; 44:2037–2078. [PubMed: 16580701]

5. Dempster FN, Corkill AJ. Interference and inhibition in cognition and behavior: Unifying themes for educational psychology. Educational Psychology Review. 1999; 11:1–88.

6. Deng W, Sloutsky VM. Carrot-eaters and moving heads: Salient features provide greater support for inductive inference than category labels. Psychological Science. 2012; 23:178–186. [PubMed: 22228644]

7. Fisher AV, Sloutsky VM. When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. Child Development. 2005; 76:583–597. [PubMed: 15892780]

8. Gelman, SA. The essential child: Origins of essentialism in everyday thought. Oxford University Press; New York: 2003.

9. Gelman SA, Heyman GD. Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. Psychological Science. 1999; 10:489–493.

10. Gelman SA, Markman E. Categories and induction in young children. Cognition. 1986; 23:183–209. [PubMed: 3791915]

11. Gelman SA, Waxman SR. Looking beyond looks: Comments on Sloutsky, Kloos, and Fisher (2007). Psychological Science. 2007; 18:554–555. [PubMed: 17576270]

12. Gliga T, Volein A, Csibra G. Verbal labels modulate perceptual object processing in one-year-old infants. Journal of Cognitive Neuroscience. 2010; 22:2781–2789. [PubMed: 20044900]

13. Griffiths O, Hayes BK, Newell BR. Feature-based versus category-based induction with uncertain categories. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2012; 38:576–595.

14. Hanania R, Smith LB. Selective attention and attention switching: towards a unified developmental approach. Developmental Science. 2010; 13:622–635. [PubMed: 20590726]

15. Hoffman AB, Rehder B. The costs of supervised classification: The effect of learning task on conceptual flexibility. Journal of Experimental Psychology: General. 2010; 139:319–340. [PubMed: 20438254]

16. Lane DM, Pearson DA. The development of selective attention. Merrill-Palmer Quarterly. 1982; 28:317–337.

17. Love BC, Medin DL, Gureckis TM. SUSTAIN: A network model of category learning. Psychological Review. 2004; 111:309–332. [PubMed: 15065912]

18. Mandler J, McDonough L. Drinking and driving don't mix: Inductive generalizastion in infancy. Cognition. 1996; 59:307–335. [PubMed: 8706380]

19. Markman AB, Ross BH. Category use and category learning. Psychological Bulletin. 2003; 129:592–613. [PubMed: 12848222]

20. Marr, D.; Vision, WH. Freeman. San Francisco, CA:

21. Murphy, GL. The Big Book of Concepts. MIT Press; Cambridge, MA: 2002.

22. Napolitano AC, Sloutsky VM. Is a picture worth a thousand words? The flexible nature of modality dominance in young children. Child Development. 2004; 75:1850–1870. [PubMed: 15566384]

23. Rehder B, Colner RM, Hoffman AB. Feature inference learning and eyetracking. Journal of Memory & Language. 2009; 60:394–419.

24. Robinson CW, Sloutsky VM. Auditory dominance and its change in the course of development. Child Development. 2004; 75:1387–1401. [PubMed: 15369521]

25. Robinson CW, Sloutsky VM. Linguistic labels and categorization in infancy: Do labels facilitate or hinder? Infancy. 2007a; 11:233–253.

26. Robinson CW, Sloutsky VM. Visual processing speed: Effects of auditory input on visual processing. Developmental Science. 2007b; 10:734–740. [PubMed: 17973789]

27. Robinson CW, Sloutsky VM. Effects of auditory input in individuation tasks. Developmental Science. 2008; 11:869–881. [PubMed: 19046156]

28. Sakamoto Y, Love BC. Learning and Retention through Predictive Inference and Classification. Journal of Experimental Psychology: Applied. 2010; 16:361–377. [PubMed: 21198253]

29. Sloutsky VM. From perceptual categories to concepts: What develops? Cognitive Science. 2010; 34:1244–1286. [PubMed: 21116483]

30. Sloutsky VM, Fisher AV. Induction and categorization in young children: A similarity-based model. Journal of Experimental Psychology: General. 2004a; 133:166–188. [PubMed: 15149249]

31. Sloutsky VM, Fisher AV. When learning and development decrease memory: Evidence against category-based induction. Psychological Science. 2004b; 15:553–558. [PubMed: 15271001]

32. Sloutsky VM, Fisher AV. Attentional learning and flexible induction: How mundane mechanisms give rise to smart behaviors. Child Development. 2008; 79:639–651. [PubMed: 18489418]

33. Sloutsky VM, Kloos H, Fisher AV. When looks are everything Appearance similarity versus kind information in early induction. Psychological Science. 2007a; 18:179–185. [PubMed: 17425540]

34. Sloutsky VM, Kloos H, Fisher AV. What's beyond looks? Reply to Gelman and Waxman. Psychological Science. 2007b; 18:556–557.

35. Sloutsky VM, Lo Y-F. How much does a shared name make things similar? Part 1: Linguistic labels and the development of similarity judgment. Developmental Psychology. 1999; 35:1478–1492. [PubMed: 10563736]

36. Sloutsky VM, Lo Y-F, Fisher AV. How much does a shared name make things similar? Linguistic Labels and the development of inductive inference. Child Development. 2001; 72:1695–1709. [PubMed: 11768140]

37. Sloutsky VM, Napolitano A. Is a picture worth a thousand words? Preference for auditory modality in young children. Child development. 2003; 74:822–833. [PubMed: 12795392]

38. Waxman SR, Markow DB. Words as invitations to form categories: Evidence from 12-13-month-old infants. Cognitive Psychology. 1995; 29:257–302. [PubMed: 8556847]

39. Welder AN, Graham SA. The influences of shape similarity and shared labels on infants' inductive inferences about nonobvious object properties. Child Development. 2001; 72:1653–1673. [PubMed: 11768138]

40. Wilburn C, Feeney A. Do development and learning really decrease memory? On similarity and category-based induction in adults and children. Cognition. 2008; 106:1451–1464. [PubMed: 17604015]

41. Yamauchi T, Markman AB. Category learning by inference and classification. Journal of Memory and Language. 1998; 39:124–148.

42. Yamauchi T, Markman AB. Inference using categories. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2000; 26:776–795.

Highlights

- This research focuses on the mechanism underlying effects of count nouns on inductive generalization across development.

- Some argue that even early in development linguistic labels are category markers, whereas other argue that label start out are features, but may become category markers in the course of development.

- Results of 6 experiments with 4-to-5-year-olds and adults support the idea that labels function as features for young children, but not for adults.

- Results of experiments also indicate that whereas categorization and induction are the same processes for young children, they are likely to be different processes for adults.
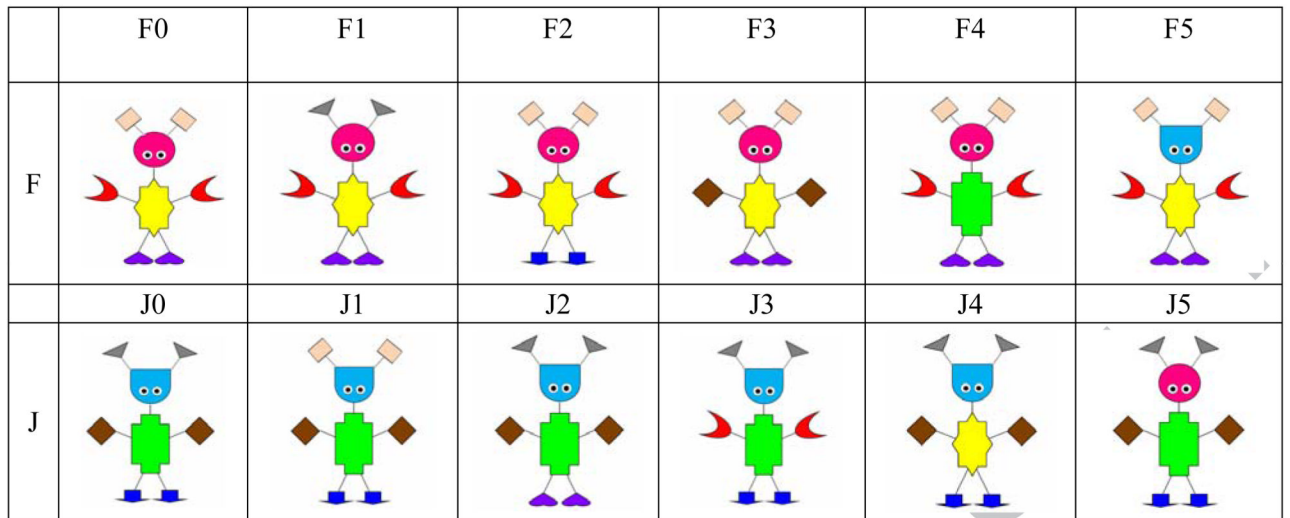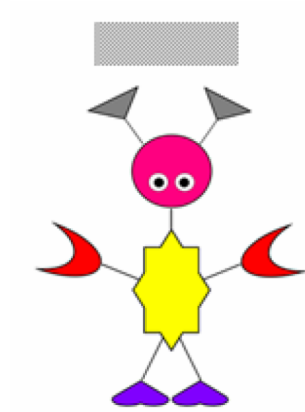
Figure 1.
Stimuli examples from two categories used in Experiments 1-3. F = Flurp; J = Jalet. F0 and J0 are prototypes of each category and F1/J1-F5/J5 are individual exemplars.
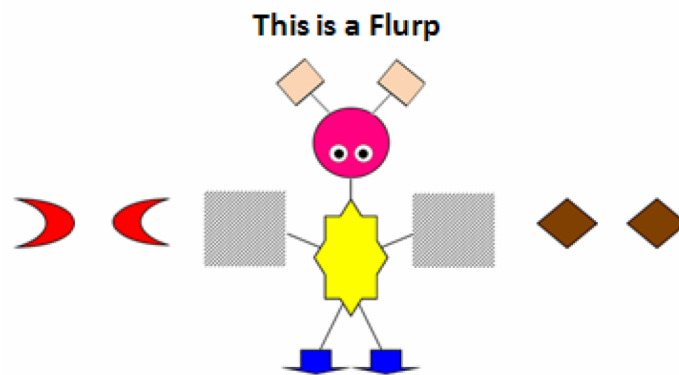
A. Classification

B. Induction

Figure 2.
Examples of Classification and Induction test trials in Experiments 1-4. A. On classification trials, participants were presented with stimuli and asked: whether the item was Flurp or Jalet? B. On induction trials, participants were presented with stimuli and asked: which body part was under the cover?

A. Adults
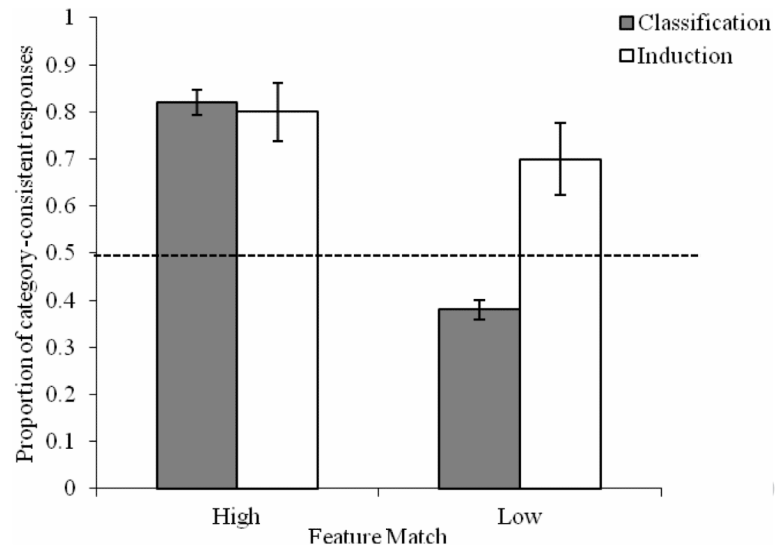
B. Children

Figure 3.
Proportion of category-consistent responses by feature match and testing condition in
Experiment 1A. *Note.* Error bars represent standard error of the mean.

A. Adults

B. Children

Figure 4.
Proportion of category-consistent responses by feature match and testing condition in Experiment 1B. *Note.* Error bars represent standard error of the mean.

A. Adults

B. Children

Figure 5.
Proportion of category-consistent responses by feature match and testing condition in Experiment 2A. *Note.* Error bars represent standard error of the mean.
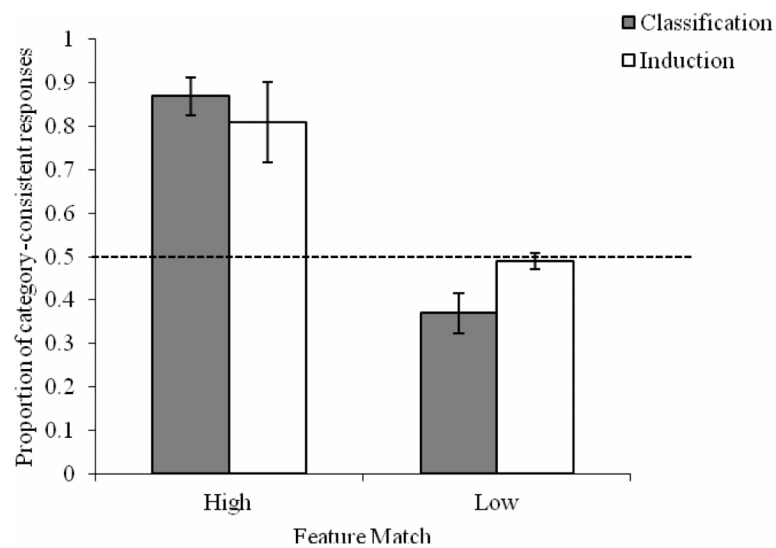
A. Adults

B. Children

Figure 6.
Proportion of category-consistent responses by feature match and testing condition in Experiment 2B. *Note.* Error bars represent standard error of the mean.

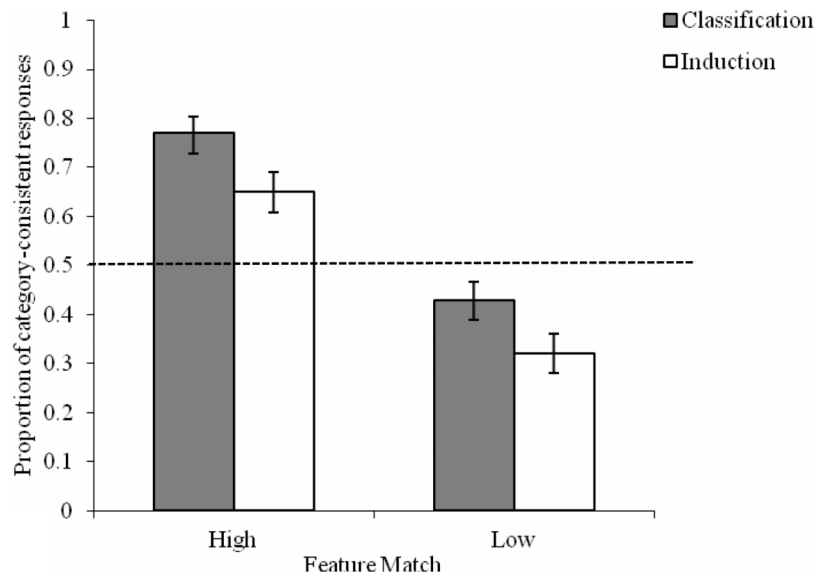A. Adults
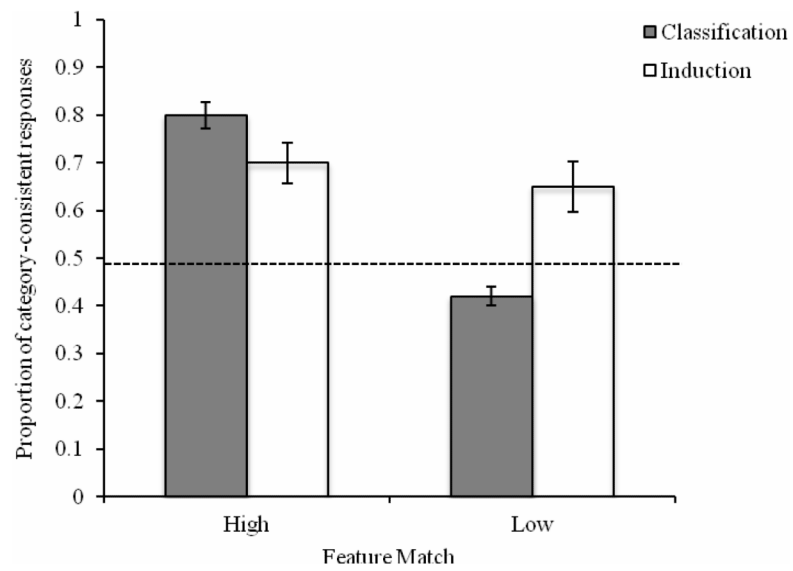
B. Children
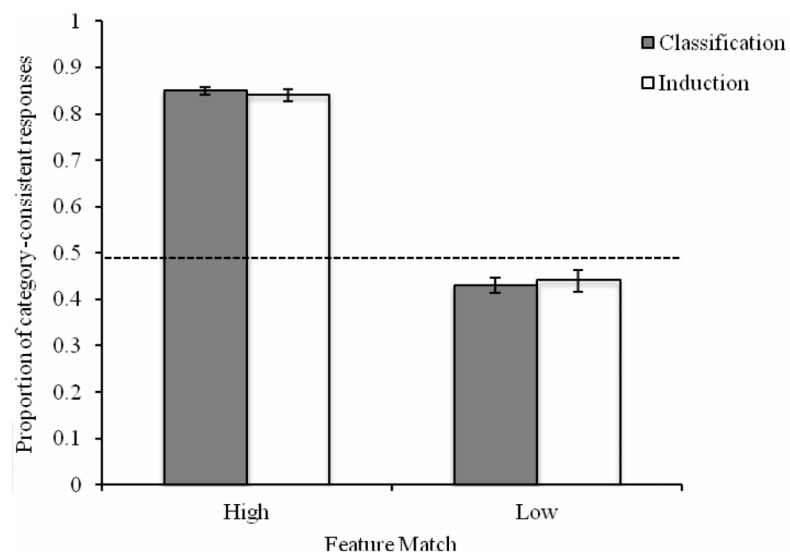
Figure 7.
Proportion of category-consistent responses by feature match and testing condition in Experiment 3. *Note.* Error bars represent standard error of the mean.

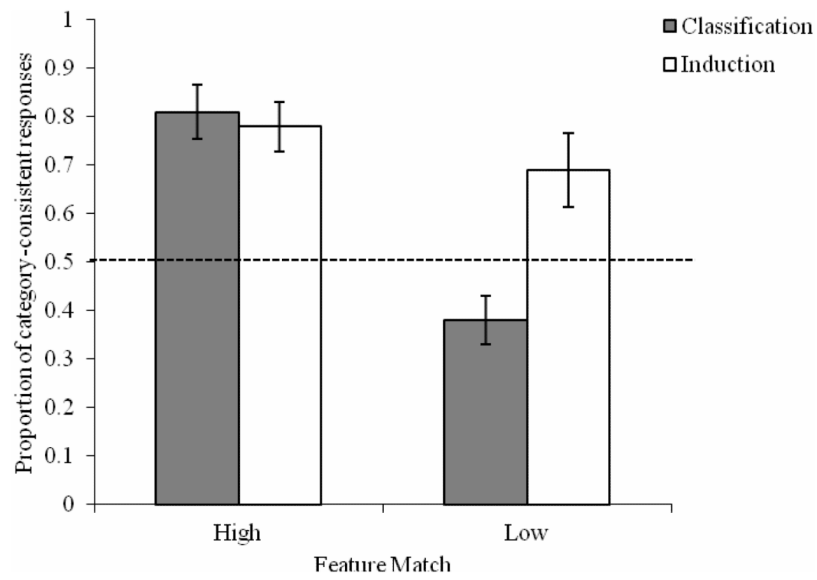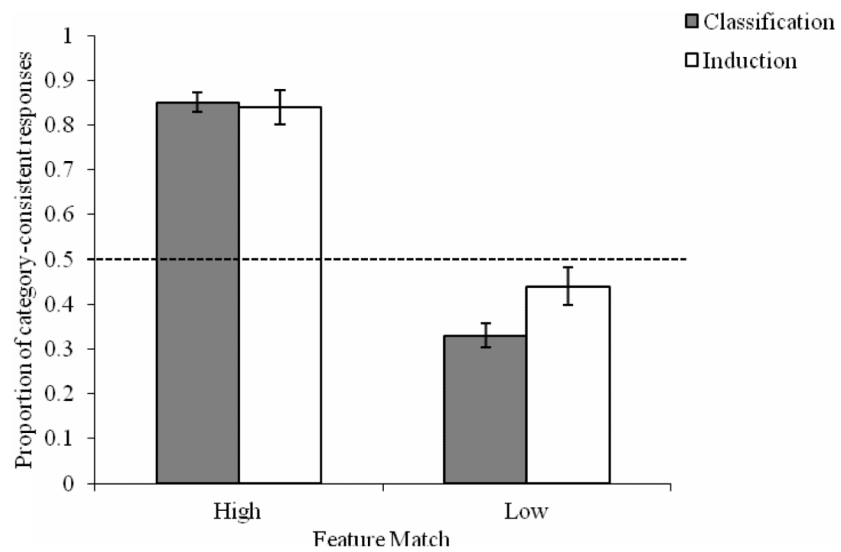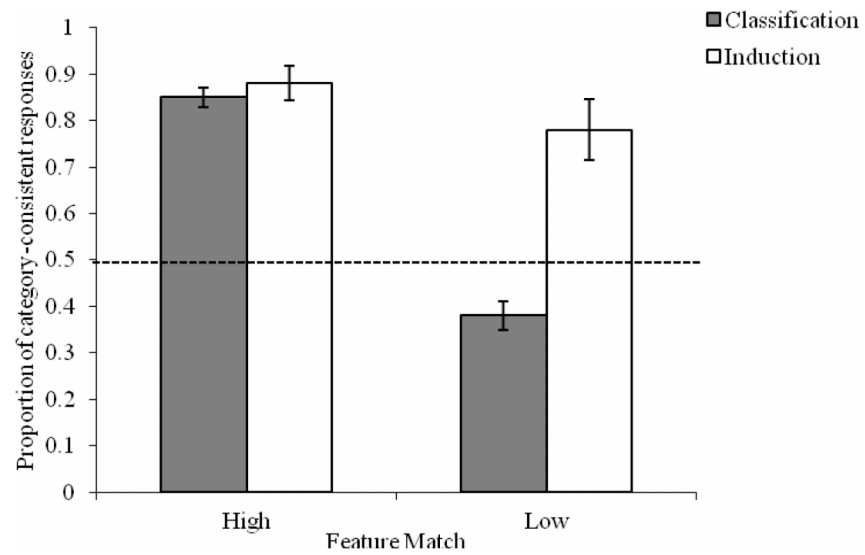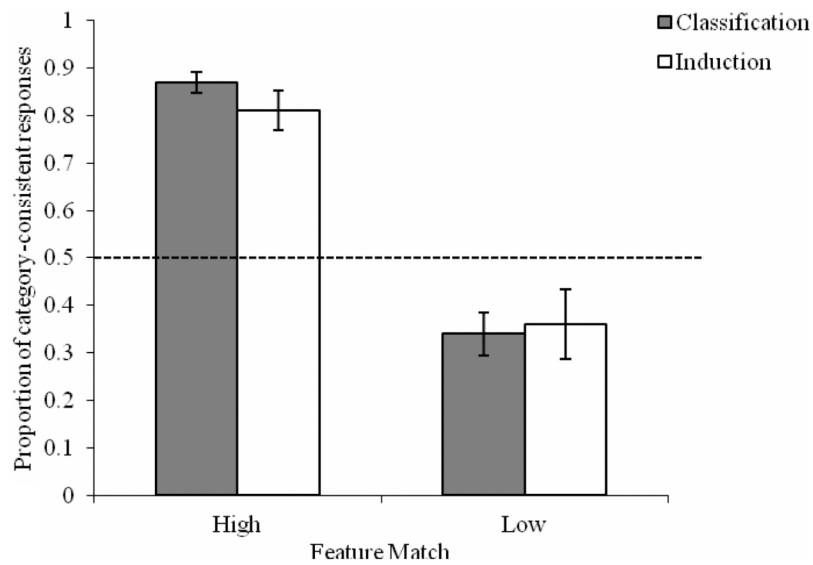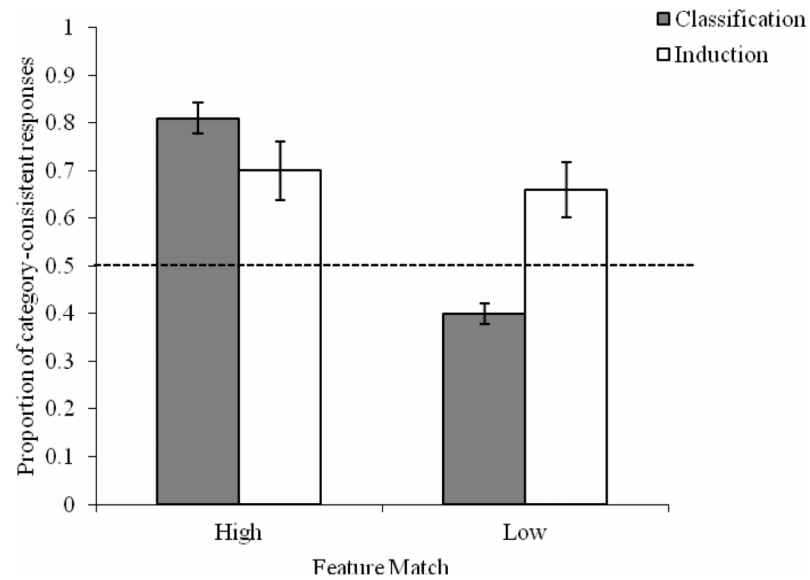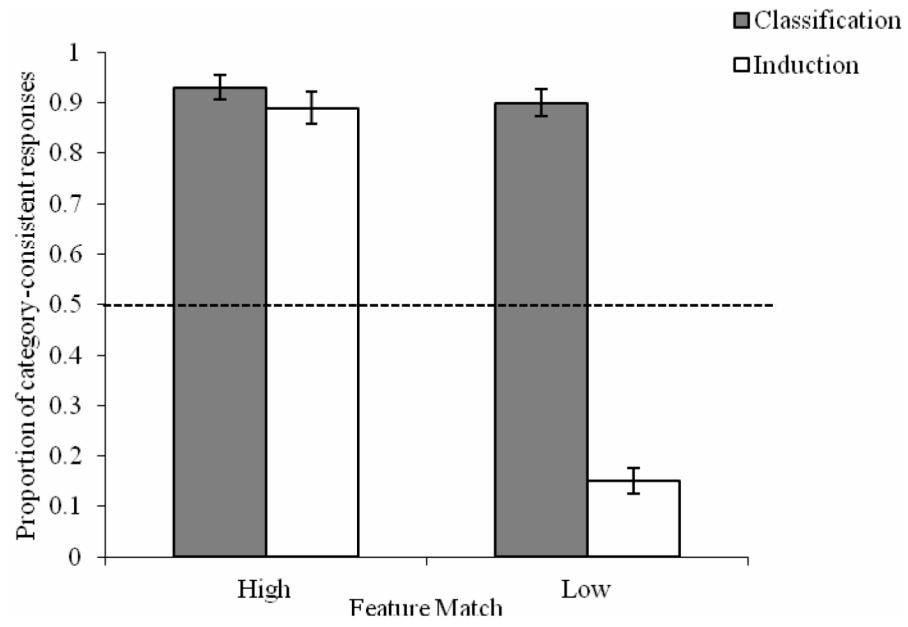**Figure 8.**
Proportion of category-consistent responses by feature match and testing condition in
Experiment 4. *Note*. Error bars represent standard error of the mean.

Table 1

Category structure used in learning in Experiment 1-4.

| | Flurp | | | | | | | Jalet | | | | | | |
Stimuli	Head	Body	Hands	Feet	Antenna	Label	Stimuli	Head	Body	Hands	Feet	Antenna	Label
F1	1	1	1	1	0	1	J1	0	0	0	0	1	0
F2	1	1	1	0	1	1	J2	0	0	0	1	0	0
F3	1	1	0	1	1	1	J3	0	0	1	0	0	0
F4	1	0	1	1	1	1	J4	0	1	0	0	0	0
F5	0	1	1	1	1	1	J5	1	0	0	0	0	0
F0	1	1	1	1	1	1	J0	0	0	0	0	0	0

Note. The value 1 = any of five dimensions identical to "Flurp" (see Figure 1). The value 0 = any of five dimensions identical to "Jalet" (see Figure 1). F = Flurp; J = Jalet. F0 and J0 are prototypes of each category.

Table 2

A. Structure of testing stimuli in Classification used in Experiments 1-3.

| | Flurp | | | | | | | Jalet | | | | | | |
Stimuli	Head	Body	Hand	Feet	Antenna	Target Label	Match	Stimuli	Head	Body	Hand	Feet	Antenna	Target Label
F11	1	1	1	1	0	?		J11	0	0	0	0	1	?
F12	1	1	1	0	1	?		J12	0	0	0	1	0	?
F13	1	1	0	1	1	?	High	J13	0	0	1	0	0	?
F14	1	0	1	1	1	?		J14	0	1	0	0	0	?
F15	0	1	1	1	1	?		J15	1	0	0	0	0	?
F21	1	0	1	0	0	?		J21	0	1	0	1	1	?
F22	0	1	0	1	0	?		J22	1	0	1	0	1	?
F23	0	0	1	0	1	?	Low	J23	1	1	0	1	0	?
F24	1	0	0	1	0	?		J24	0	1	1	0	1	?
F25	0	1	0	0	1	?		J25	1	0	1	1	0	?

B. Structure of testing stimuli in Induction in Experiments 1-3.

| | Flurp | | | | | | | Jalet | | | | | | |
Stimuli	Head	Body	Hand	Feet	Antenna	Target Label	Match	Stimuli	Head	Body	Hand	Feet	Antenna	Target Label
F11	1	?	1	1	0	1		J11	0	?	0	0	1	0
F12	1	1	?	0	1	1		J12	0	0	?	1	0	0
F13	?	1	0	1	1	1	High	J13	?	0	1	0	0	0
F14	1	0	1	1	?	1		J14	0	1	0	0	?	0
F15	0	1	1	?	1	1		J15	1	0	0	?	0	0
F21	?	0	1	0	0	1		J21	?	1	0	1	1	0
F22	0	?	0	1	0	1		J22	1	?	1	0	1	0
F23	0	0	?	0	1	1	Low	J23	1	1	?	1	0	0
F24	1	0	0	?	0	1		J24	0	1	1	?	1	0
F25	0	1	0	0	?	1		J25	1	0	1	1	?	0

Note. High and low are two levels of feature match. F = Flurp; J = Jalet. Category-consistent responses were the ones consistent with the values indicated in the target features and target labels.

Table 3

An overview of Experiments 2A and 2B.

Experiment	Participants	Training Details	Testing
2A	13 adults (2 women) 9 children (M = 56.6 months, range 52.8-60.1 months; 5 girls)	Training Procedure: Classification Verbal information: *Friendly pet* vs. *Wild creature* Memory Check: Children (83.3%) Adults (86.2%)	Classification and Induction
2B	12 adults (6 women) 9 children (M = 53.6 months, range 51.1-58.6 months; 3 girls)	Training Procedure: Induction Verbal information: *This one lives in the forest* vs. *This one lives in the sea* Memory Check: Children (84.4%), adults (76.7%)	Classification and Induction

Table 4

Numbers (and percentages in parentheses) of label-based, similarity-based and mixed responders in Experiments 1, 2A, and 3.

	Children	Adults
Label-based responders	2 (6.0%)	29 (83%)
Similarity-based responders	16 (47%)	1 (3%)
Mixed responders	16 (47%)	5 (14%)

Table 5

Category structure used in learning in Experiments 4.

	Category A							Category B					
Stimuli	Body	Hands	Feet	Antenna	Label	Head	Stimuli	Body	Hands	Feet	Antenna	Label	Head
A1	1	1	1	0	1	1	B1	0	0	0	1	0	0
A2	1	1	0	1	1	1	B2	0	0	1	0	0	0
A3	1	0	1	1	1	1	B3	0	1	0	0	0	0
A4	0	1	1	1	1	1	B4	1	0	0	0	0	0
A0	1	1	1	1	1	1	B0	0	0	0	0	0	0

Note. The value 1 = any of six dimensions identical to Category A (see Figure 1). The value 0 = any of six dimensions identical to Category B (see Figure 1). A = Category A; B = Category B. A0 and B0 are prototypes of each category and A1/B1 –A4/B4 are individual exemplars.

Table 6

A. Structure of testing stimuli in Classification used in Experiments 4.

	Category A						Match		Category B					
Stimuli	Body	Hand	Feet	Antenna	Label	Head		Stimuli	Body	Hand	Feet	Antenna	Label	Head
A11	1	1	1	0	?	1		B11	0	0	0	1	?	0
A12	1	1	0	1	?	1	High	B12	0	0	1	0	?	0
A13	1	0	1	1	?	1		B13	0	1	0	0	?	0
A14	0	1	1	1	?	1		B14	1	0	0	0	?	0
A21	0	1	0	0	?	1		B21	1	0	1	1	?	0
A22	1	0	0	0	?	1	Low	B22	0	1	1	1	?	0
A23	0	0	0	1	?	1		B23	1	1	1	0	?	0
A24	0	0	1	0	?	1		B24	1	1	0	1	?	0

B. Structure of testing stimuli in Induction used in Experiments 4.

	Category A						Match		Category B					
Stimuli	Body	Hand	Feet	Antenna	Label	Head		Stimuli	Body	Hand	Feet	Antenna	Label	Head
A11	?	?	1	0	1	1		B11	?	0	0	1	0	0
A12	1	?	0	1	1	1	High	B12	0	?	1	0	0	0
A13	1	0	?	1	1	1		B13	0	1	?	0	0	0
A14	0	1	1	?	1	1		B14	1	0	0	?	0	0
A21	0	?	0	0	0	1		B21	1	?	1	1	1	0
A22	?	0	0	0	0	1	Low	B22	?	1	1	1	1	0
A23	0	0	0	?	0	1		B23	1	1	1	?	1	0
A24	0	0	?	0	0	1		B24	1	1	?	1	1	0

Note. High and low are two levels of feature match. A = Category A; B = Category B.