# Acquiring Word Learning Biases

**Zi L. Sim (zilin@berkeley.edu)**
**Sylvia Yuan (shyuan@berkeley.edu)**
**Fei Xu (fei_xu@berkeley.edu)**

Department of Psychology, University of California, Berkeley
Berkeley, CA 94720 USA

## Abstract

Previous research has shown that infants can acquire a "shape bias" in word learning when presented with labels that are perfectly correlated with object shape. However, little research examines whether children can acquire a non-shape word learning bias. Research on inducing this bias can help inform the origins of the shape bias. In our experiment, 3 year-old children successfully acquired a new bias even with relatively few objects and a short training session, illustrating the relevance of the overhypothesis formation account in explaining the acquisition of early inductive constraints.

**Keywords:** shape bias; overhypothesis; word learning; inductive constraints

## Introduction

How does a child, learning his/her first language, map words to their referents? According to Quine (1960), a young child learning his first language, is similar to a traveler in a new place without any knowledge of the local language. When a native speaker says "Gavagai!", there is an infinite number of referents for the word, such as whole objects, parts of objects, or properties of objects. How does the traveler, or the child, then map the word "gavagai" to the correct referent?

Early word learning is thus an inductive problem, where the young child attempts to make inferences about words and their referents. Children appear to be remarkably proficient at solving such problems of induction. They circumvent the mapping problem, and quickly map the words they hear to the correct object category (Bloom, 2000; Markman, 1989; Carey & Bartlett, 1978).

The apparent ease at which young children learn suggests that their learning is supported by strong inductive constraints (Markman, 1989). In the domain of word learning, a possible inductive constraint is the shape bias (Landau, Smith & Jones, 1988), the tendency for children to extend new object names on the basis of shape similarity. This shape bias has been demonstrated in various studies (e.g. Landau et al., 1988; Soja, Carey & Spelke, 1991; Samuelson, 2002) and has also been consistently linked with an increased rate of noun acquisition in children (e.g. Smith, Jones, Landau, Gershkoff-Stowe & Samuelson, 2002), suggesting that this particular bias may aid children in mapping new words to their referents quickly.

How does the shape bias emerge? While one would be hard-pressed to find a truly nativist account of the shape bias, some might still argue that the shape bias is an innate constraint that a child is born with, especially because it emerges early in development. The nativist account posits that children are biased to pay attention to shape of objects, and to generalize labels according to shape. However, such a nativist account neither addresses how the shape bias came to be an innate constraint, nor the processes and mechanisms behind the bias (Xu, Dewar & Perfors, 2009). Furthermore, children do not approach word learning with an expectation that extension of labels is solely guided by shape similarity. For example, children consider the taxonomic categories (Cimpian & Markman, 2005) and the solidity (Soja et al., 1991) of objects. Hence, a nativist account of the shape bias is one that is difficult to defend.

While a nativist account of the shape bias has been widely rejected, other theories fall on a continuum between the constraint being innate, and the constraint being learnt. The "shape-as-cue" account, for example, (Bloom, 2000; Markson, Diesendruck & Bloom, 2008) falls at the center of such a continuum.

The proponents of the "shape-as-cue" account argue that children organize objects by shape because shape serves as a reliable cue to object kinds. In this account, a shape bias emerges from two different understandings: first, children understand that count nouns refer to object categories (Dewar & Xu, 2007) and second, they perceive shape to be a reliable and available cue to object category membership.

Consequently, on this account, the perceptual dimension of shape is not special – when other cues such as texture or function are more reliable cues for object category membership, children use them in their extension of novel names (Booth & Waxman, 2002; Ware & Booth, 2010), and a shape bias will not emerge.

However, the shape-as-cue account leaves the origins of the shape bias unaddressed. For example, how do children come to believe that shape is a reliable cue of object category membership? We investigate this question in the current study, by asking whether children can acquire a different bias when given the relevant input.

Another account, the Attentional Learning Account (ALA) proposed by Smith and her colleagues (Smith et al., 2002), addresses the origins of the shape bias more fully, by positing that the bias is a learnt constraint. According to ALA, the development of word learning biases is supported by connectionist networks, which allow children to learn from the many "correlations among linguistic devices, object properties, and perceptual category organization" (Colunga & Smith, 2005) in their learning environment.

For example, solid rigid objects tend to be organized into categories according to their shapes. Children notice such regularities and begin to form associations between object names and object shapes in individual categories. This association forms the first-order generalization, which refers to the structure of individual categories: "cups are cup-shaped; balls are ball-shaped; etc.". Then children make associations across these learned categories, forming the second-order generalization: "X's are X-shaped," and that categories are organized by the similarities in shape. These learned associations later direct children to attend to shape when learning new names (Smith et al., 2002).

In summary, the ALA argues that the shape bias is a learnt constraint, and derives entirely from the evidence that children encounter in their learning environment and the associations that they make between names of object categories and the properties of objects.

While ALA makes a strong case for how the shape bias emerges from regularities in children's environment, there are alternative explanations for the bias' origins. A recent framework construes the shape bias as an overhypothesis, describing its emergence via computational principles of hierarchical Bayesian model (Goodman, 1955; Kemp, Perfors & Tenenbaum, 2007; Xu et al., 2009). According to Goodman (1955), overhypothesis refers to an over-arching generalization that one makes across different categories. As such, an overhypothesis is analogous to the second-order generalization identified in the ALA by Smith et al. (2002).

The point of divergence between the ALA and the overhypothesis formation account occurs in their respective explanations for how children acquire the higher order generalizations that aid their noun extensions. On the one hand, the ALA appears to be a bottom-up associative process, as suggested by the 4-step model describing the ALA in Smith et al. (2002). In relation to higher order generalizations, the authors argue that "the child makes a higher-order generalization across learnt categories about the common structure of named object categories," implying that the first order generalizations form the basis for higher order generalizations. Thus, children first begin by collecting evidence for lower levels of generalizations, and then move into making higher order generalizations.

On the other hand, the overhypothesis formation account describes learners as collecting statistical evidence and making generalizations at multiple levels simultaneously. In fact, it may be possible to make a second order generalization before the first order generalization, or to be more confident about the higher-order generalization (Kemp et al., 2007; Xu et al., 2009).

Accordingly, ALA strikes one as a rather laborious process that requires a lot of time for learning and large amounts of data, while the overhypothesis formation account suggests that learning can occur at rapid speeds with relatively sparse data.

Therefore, the overhypothesis formation account provides a new framework to explain how the shape bias may be quickly acquired from inputs in the child's environment. For example, when children encounter evidence that cups have cup-shapes, their hypothesis that "cups are all cup-shaped" has an increased probability of being true, while other hypotheses, such as "cups are all red in color," decrease in probability correspondingly. At the same time, their overhypothesis that "X's are all X-shaped" has an increased probability of being true as well. Through the collection of data describing labeling instances and adjusting their hypotheses accordingly, children can acquire a shape bias that manifests itself as a constraint, an expectation that objects in the same category tend to have the same shape (Kemp et al., 2007). More importantly, the overhypothesis account also predicts that learners can rapidly learn to apply these generalizations to new instances, extending brand-new labels even when provided with only one or a few positive exemplars. The overhypothesis formation account is thus a powerful mechanism that can account for the acquisition of this inductive constraint.

Despite their differences, both accounts have a common premise – that the shape bias is learnt. Two lines of empirical research can demonstrate that the shape bias is indeed acquired from the statistical regularities in the environment: one is to conduct training studies on the shape bias itself; the other is to see if young children can be trained to acquire a different bias given the relevant input.

In the first line of research, training studies are used to induce a shape bias in children younger than the age at which they show a reliable tendency to extend novel names by shape. In Smith et al. (2002), experimenters labeled objects that were organized perfectly according to shape, once a week for 7 weeks to 17-month-old infants. Objects of the same shape had the same label, and this training facilitated the emergence of a shape bias. Clearly, the shape bias is learnt from input in the environment, justifying the Attentional Learning Account.

However, note that these results are also consistent with the overhypothesis formation account. In fact, the sparseness of the data presented, coupled with the speedy success of the infants, suggests that the overhypothesis formation account may be more in line with the results than the ALA (Xu et al., 2009).

In the present study, we pursue the second line of research, attempting to show that the shape bias is learnt by inducing a non-shape word learning bias in children through training. If the shape bias is indeed acquired from data in the environment, presenting children with objects organized and named according to similarities in a different dimension, such as color or pattern, should cause them to pick up a color or pattern bias respectively.

One study that had earlier explored this line of research is Samuelson (2002), which provided 15 – 20 month-old infants with intensive naming experience using solid exemplars either organized by shape similarity (Shape-biased condition) or material similarity (Material-biased condition). Infants in the Shape-biased condition extended labels by shape reliably in the generalization task. However,

infants in the Material-biased condition did not reliably extend labels by material.

As the author speculates, the children did not acquire a material bias potentially because of their previously learnt noun labels, which prevented them from developing a material bias. However, we contend that the non-emergence of the material bias may have resulted from material being a more subtle dimension than shape. Hence, the infants might have been overwhelmed by the salience of the shape dimension, and failed to notice that the exemplars were labeled and organized by their material. Accordingly, using a different dimension with a higher perceptual salience may facilitate the emergence of a non-shape word learning bias.

Recent research has illustrated that infants have the capacity to form overhypotheses over different dimensions, such as shape and color. In Dewar & Xu (2010), an experimenter sampled from 3 boxes and produced 4 objects from each box that were identical in shape. 9 month-old infants looked significantly longer when the objects taken from a fourth box did not match in shape, indicating that they had formed an overhypothesis from the familiarization trials that "objects from the same box are uniform in shape". Furthermore, this learning was not limited to the shape dimension; infants could form an overhypothesis about color uniformity given the appropriate input.

These results suggest that the mechanism for overhypothesis formation is present very early, and that infants have the capacity to acquire a non-shape bias, at least in non-linguistic tasks. Thus, we consider it likely that children can rapidly develop a different bias when presented with appropriate regularities in a word learning context.

We designed an experiment asking if children can acquire a non-shape word learning bias when provided with the appropriate regularities. The procedure was modeled after Smith et al. (2002), except that the training phase was limited to a single 10-minute session. Three year-old children were provided with a naming experience consisting of exemplars categorized according to either shape or pattern. They were then asked to make first-order generalizations, where they had to extend a label that they had previously heard on to a novel object, and second-order generalizations, where they had to extend a novel label to a new object. Also, we chose to use pattern as the dimension of interest, as it is another perceptual dimension that is readily available. With such a short training session, a rapid induction of a non-shape word learning bias would suggest that overhypothesis formation can better account for the developmental origins of the shape bias.

## Method

### Participants

Forty-eight English-speaking 3-year-olds (23 boys and 25 girls) with a mean age of 38.7 months (range = 29.9 to 46.9 months) were tested. All were recruited from Berkeley, California, and its surrounding communities. An additional 9 children were tested but excluded due to failing to make a correct choice during the practice trials (N = 3), side bias[1] (N = 3), parental interference (N = 2) or refusal to point (N = 1).

### Materials and Procedure

The children were tested individually in the laboratory in front of a laptop, with their parents seated behind them in the testing room. Each child was randomly assigned to a Shape condition, a Pattern condition, or a Control condition. They were introduced to novel nonsense words under the cover story of two stuffed animals (Doggy and Mr. Crocodile) sharing their language with us, an animal for each critical block. The cover story served to discourage the children from possibly using the biases they may have previously acquired. Novel objects with various shapes and patterns were presented on the laptop. The words labeling the novel objects were used in sentences with count nouns syntax, such as "Look, this is a *blick*!".

**Shape and Pattern Conditions** The procedure for the experimental conditions (Shape or Pattern) consisted of a practice phase and 2 critical blocks of trials. Each critical block consisted of a training phase, followed by a test phase. There was a 1 minute interval between the 2 critical blocks, to re-engage the child in the task.

In the practice phase, each child received 2 practice trials. In each practice trial, the child was presented with a target item – a novel-looking object – that was labeled with a novel name (e.g. blick, geel, toopa). Then, the child was presented with 2 choice items: an identity match, which was identical to the target item, and a distracter, which differed from the target item in both shape and pattern.

In the training phase of each critical block, each child saw 16 novel objects, presented either in 2 categories with 8 members each (2-category set), or 4 categories with 4 members each (4-category set). In the shape condition, each category had objects that were identical in shape but differed in pattern, and in the pattern condition, each category had objects that were identical in pattern but differed in shape. (See Figure 1.) Objects in the same category were labeled with the same nonsense word. Each object category was presented 3 times, and was associated with its label 5 times.

In the test phase, each child was given a first-order generalization test, and a second-order generalization test. In the first-order test, each child was presented with a *familiar target object*, an object previously shown in the training phase. This object was paired with its name, "Doggy said this was a X." At test, the child was shown 2 new objects: a *shape-match* object, which was identical to the target in shape but not pattern, and a *pattern-match* object, which was identical to the target in pattern but not shape. (See Figure 2.) The child was then requested to choose between the two objects by being asked, "Can you point to another X?"

---

[1] 5 or more points to the same side out of 6 trials. See procedure.

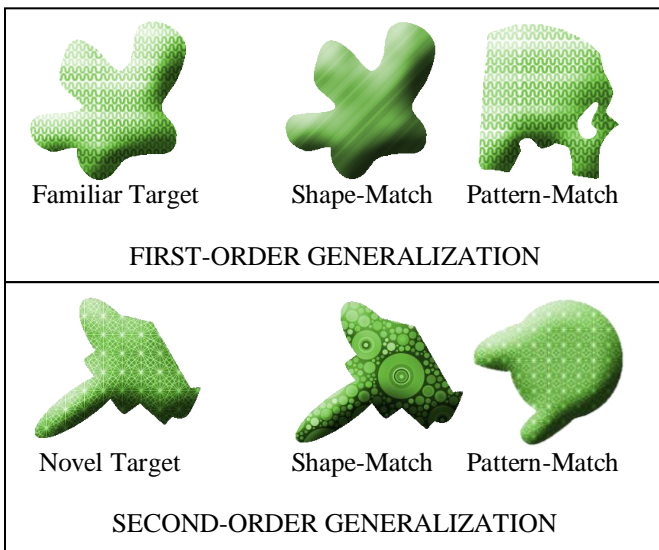Figure 1: One set of stimuli used in training phase.



Figure 2: One set of stimuli used for test trials.

In the second order test, each child was presented first with a *novel target object*, an object that had never been used in the training phase. This novel object was then paired with a novel label, "Doggy told me this is a X." At test, the child was asked to extend the novel label by choosing between 2 new objects: a *shape-match* object and a *pattern-match* object. (See Figure 2.)

**Control Condition** In the control condition, the training phases were removed from the critical blocks, and the children entered the test phase immediately after the practice phase. The control condition provided a baseline measure of the children's extension of novel words according to shape or pattern.

**Coding**

In the test trials of the 3 conditions, choosing a *pattern-match* object was scored as 1 point, and the maximum score for each child was 4 as there were 4 test trials in total. The children's scores were then converted into percentage of

pattern-extension (extending labels to novel objects according to pattern).

## Results

An alpha level of 0.05 was used in all statistical analyses. Preliminary analyses of the percentage of pattern-extension found no effects of gender, age-split (whether the children were younger or older than the median age of the group) or category structure (2- vs. 4-category set used in training phase). Subsequent analyses were thus collapsed over these variables.



Figure 3: Percentage of pattern-extension in pattern, shape and control conditions. Error bars represent SE.

As Figure 3 indicates, the children's percentage of pattern-extension varied according to the condition they were assigned to. Using the children's percentage of pattern-extension in the test trials, a 3x2 repeated measures analysis of variance (ANOVA) was performed with Condition (Shape vs. Pattern vs. Control) as a between-subjects factor and Order of Generalization ($1^{st}$ order generalization vs. $2^{nd}$ order generalization) as a within-subjects factor. The analysis revealed a significant main effect of Condition, $F (2, 45) = 6.91$, $p = .002$; the children are significantly more likely to make extensions by pattern in the Pattern condition than in the Shape or Control conditions. The analysis also revealed no main effect of Order of Generalization, $F (1, 45) < 1$; children were equally successful in the $1^{st}$-order generalization task and the $2^{nd}$-order generalization task. There was no significant interaction between Condition and Order of Generalization, $F (2, 45) < 1$.

Critically, we were interested in the effect of training in both the Shape and Pattern conditions. Children in the Shape condition extended labels to new instances by pattern 16% (*SE = 6*) of the time, Patten condition 59% (*SE = 10*) of the time and Control condition 27% ( *SE = 10*) of the time. Children in the Pattern condition were significantly more likely to make extensions according to pattern, as compared to the children in the Shape condition, $t(30) = 3.58$, $p = .001$ and the children in the Control condition, $t(30) = 2.42$, $p = .02$. The children in the Shape and Control

conditions did not differ in the percentage of pattern-extension, t(30) < 1. Note that the comparisons were not made against .5, as children were not equally likely to extend labels by shape and pattern without training. In fact, at chance levels, children's percentage of pattern extension is .27, as demonstrated in the Control condition.

## Discussion

The present study examined whether 3 year-olds can acquire a new bias in a word learning context. The results demonstrate that the children can. After receiving rather sparse evidence in a short training phase (about 10 minutes), the children quickly made first-order and second-order generalizations about how the objects were being labeled and categorized. Using these generalizations, they extended labels to novel objects by shape or pattern, depending on the training condition that they were assigned to. Children in the Pattern condition were significantly more likely to make extensions of novel words according to pattern, as compared to children in the Shape condition or the Control condition. These results provide evidence that children are capable of acquiring non-shape word learning biases from appropriate regularities in the environment.

Furthermore, children performed similarly in the first and second order generalization tasks after a short training session, suggesting that they were collecting statistical evidence at multiple levels simultaneously.

The effect of training in the Pattern condition appears to be strong, as indicated in Figure 3. In the Control condition, we measured children's percentage of pattern-extension at baseline, without any exposure to training that consistently organized and labeled objects by shape or pattern. The children in the Control condition extended labels according to pattern 27% of the time (i.e., by shape 73% of the time), suggesting that our child participants had a pre-existing shape bias, possibly acquired from the naming experiences that they had in their learning environment. Yet, the children in the Pattern condition were significantly more likely to extend labels by pattern, demonstrating that the pattern training had facilitated the emergence of a "pattern bias" in a word learning context. Clearly, children do not blindly attend to the shape of objects, and do not broadly apply the same shape bias in extending labels for all objects.

Regrettably, our results do not rule out another plausible explanation: instead of learning that the objects were labeled according to pattern, the children in the Pattern condition simply learnt that the objects were not labeled by shape. Even if this alternative explanation was really the case, the results still suggest that the Pattern training causes the children's shape bias to be temporarily overridden, and also interestingly induces a different sort of bias, albeit not the one we expected. Future studies can rule out this alternative explanation by using a third object in test trials.

The speedy induction of a non-shape word learning bias with rather sparse evidence from the study's training phase suggests the presence of a mechanism in children that allows them to acquire word learning biases very rapidly with few exemplars. The results support the overhypothesis formation account, suggesting that children are rational learners. By using the evidence that they observe, children update their hypotheses at multiple levels simultaneously, acquiring inductive constraints that in turn support their learning. In our case, these constraints manifest as word learning biases, aiding children in mapping labels accurately to their referents.

Although the hierarchical Bayesian model that underlie the overhypothesis formation account do not specify the exact process of acquiring hypotheses, we can speculate that comparison might be one way to account for how hypotheses arise. When children compare objects, structural alignment invites extraction of common structures, which translates into the generation of new hypotheses (Christie & Gentner, 2010). At present, Gentner and her colleagues have focused on how comparison is applied to first-order generalizations, limiting the application of analogical learning models to the discussion of shape bias. The contributions of these models will certainly be magnified with future studies that examine the role of comparison on higher order generalizations.

In principle, connectionist models that underline the Attentional Learning Account (ALA) (Smith et al., 2002; Smith & Samuelson, 2006; Colunga & Smith, 2005) can also explain the results from the present study. Despite the rapid speed of learning displayed by the children in the Pattern condition, these results can be accounted by connectionist models with the appropriate parameters. However, we believe that this account is less likely, as connectionist models tend to be slow and laborious, requiring extensive amounts of evidence and long periods of time for learning before making appropriate generalizations.

Future research into the points of divergence between these connectionist models, which describe the ALA, and the hierarchical Bayesian model, which describes the overhypothesis formation account, will shed light on the nature of inductive constraints that support children's early learning.

One such point of divergence relates to the category structure of stimuli presented during training sessions. Connectionist models posit that first-order generalizations are made before second-order generalizations, and thus predicts that the strength of first-order generalizations increases as number of exemplars per category increases. Conversely, the hierarchical Bayesian model posits that the levels of generalization are made simultaneously, thus the strength of the learner's generalizations depends more greatly on the number of categories presented (Kemp et al., 2007).

Consequently, the hierarchical Bayesian model predicts that the children's accuracy in the generalization tasks will be higher if they were presented with the 4-category set of stimuli than if they were presented with the 2-category set of stimuli. However, no main effect was found for Category Structure in the present study. A likely explanation for this null result is that the difference between the 4-category and

2-category set is too small to influence the children's accuracy in the generalization task. Hence, future studies can make this difference more extreme, e.g. 2 categories with 8 members each, vs. 8 categories with 2 members each.

Our results provide strong support for the existence of an inductive learning mechanism in children that requires little data and training. Furthermore, this mechanism operates at multiple levels of generalization and thus explains how inductive biases can emerge rapidly early in development.

This mechanism is also most likely domain-general and can be applied to both word learning and non-linguistic tasks. Recent studies have demonstrated that children display different biases for objects in different domains; for example, they consider both shape and texture when extending labels to novel animate objects (Booth, Waxman & Huang, 2005). Overhypothesis formation can potentially account for all these results, as it provides a framework in which the child can use both conceptual and perceptual knowledge as data to update their hypotheses about how objects are labeled and categorized in their environment.

Finally, our results have implications for how other cognitive biases appear early in development. As our results indicate that children can quickly acquire a new word learning bias, they counter intuitions that constraints emerging early in development must be innate. Inductive constraints can be acquired rapidly through a child's interactions with his environment during his early years.

In summary, the present study attempts to address the developmental origins of the shape bias. This study adds to other training studies that have induced word learning biases in children (Smith et al., 2002; Samuelson, 2002) and extends these studies by documenting the acquisition of a non-shape word learning bias. To our knowledge, the present study is the first to use the pattern dimension to induce a non-shape bias, and the results are encouraging. These results strongly suggest that the shape bias is a learnt constraint, acquired from children's observations of how objects are consistently named and categorized in their learning environment. In the process of collecting evidence and making generalizations about an individual category, children simultaneously make higher order generalizations, facilitating the rapid emergence of an overhypothesis, which manifests naturally in development as a shape bias. However, this shape bias is not a static inductive constraint – it can be overridden when other properties of objects are more reliable cues to object category membership.

## Acknowledgments

## References

Bloom, P. (2000). *How children learn the meaning of words*. Cambridge, MA: MIT Press.

Booth, A. E., & Waxman, S. R. (2002). Word learning is 'smart': Evidence that conceptual information affects preschoolers' extension of novel words. *Cognition, 84*(1), B11 - B22.

Booth, A. E., Waxman, S. R., & Huang, Y. T. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology, 41*(3), 491-505.

Carey, S. & Bartlett, E. (1978). Acquiring a single new word. *Proceedings of the Stanford Child Language Conference,* 15, 17-29.

Christie, S. & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development, 11(3),* 356 – 378.

Cimpian, A., & Markman, E. M. (2005). The absence of a shape bias in children's word learning. *Developmental Psychology, 41*(6), 1003-1019.

Colunga, E., & Smith, L. B. (2005). From the lexicon to expectations about kinds: A role for associative learning. *Psychological Review, 112*(2), 347-382.

Dewar, K., & Xu, F. (2007). Do 9-month-old infants expect distinct words to refer to kinds? *Developmental Psychology, 43*(5), 1227 - 1239.

Dewar, K., & Xu, F. (2010). Induction, Overhypothesis, and the Origin of Abstract Knowledge: Evidence from 9-month-old Infants. *Psychological Science, 21*, 1871-1877.

Goodman, N. (1955). *Fact, Fiction and Forecast.* Cambridge, MA: Harvard University Press.

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical bayesian models. *Developmental Science, 10*(3), 307-321.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3*(3), 299-321.

Markman, E. (1989). *Categorization and Naming in Children.* Cambridge, MA: The MIT Press.

Markson, L., Diesendruck, G., & Bloom, P. (2008). The shape of thought. *Developmental Science, 11*(2), 204-208.

Quine, W. (1960). *Words and Objects.* Cambridge, MA: The MIT Press.

Samuelson, L. K. (2002). Statistical regularities in vocabulary guide language acquisition in connectionist models and 15-20-month-olds. *Developmental Psychology, 38*(6), 1016-1037.

Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological Science, 13*(1), 13-19.

Ware, E. A., & Booth, A. E. (2010). Form follows function: Learning about function helps children learn about shape. *Cognitive Development, 25*(2), 124-137.

Xu, F., Dewar, K., & Perfors, A. (2009). Induction, overhypotheses, and the shape bias: Some arguments and evidence for rational constructivism. In B. M. Hood & L. Santos (Eds.), *The origins of object knowledge.* New York, NY: Oxford University Press.