

# Acquiring Inductive Constraints from Self-Generated Evidence

Zi L. Sim (zi@berkeley.edu)  
Fei Xu (fei\_xu@berkeley.edu)

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

## Abstract

Several recent theoretical proposals suggest that young children are rational, constructivist learners (e.g. Gopnik & Wellman, 2012; Xu & Kushnir, 2012; 2013). One of the claims made under constructive learning is that children are active learners – they selectively attend and explore their environment in order to maximize information gain (e.g., Kidd, Piantadosi, & Aslin, 2012; Schulz & Bonawitz, 2007). Most studies to date, however, have focused on how efficiently children learn when they are given evidence by an experimenter (‘teacher’), under conditions of training: children receive a restricted set of evidence, and they are subsequently tested on their learning. Yet children are not mere observers; they actively engage their environment to supplement their learning. In our experiment, 3-year-old children successfully acquired higher-order generalizations using self-generated evidence during free play, suggesting an early capacity to engage in self-directed learning.

**Keywords:** self-directed learning; free play; generalization; causal learning

## Introduction

What is the nature of early learning? In recent years, several theoretical accounts of cognitive development have emerged, each describing the young child as a rational, constructivist learner (e.g. Gopnik & Wellman, 2012; Xu & Kushnir, 2012; 2013). These proposals make two key claims: first, learners form meaningful generalizations based on limited evidence obtained from their environment, and second, the child is an active learner.

First, much of early learning may be characterized as inductive learning, i.e. making principled and meaningful generalizations based on limited amounts of data. Research has repeatedly demonstrated that young children engage in such learning proficiently: they generalize non-obvious properties to novel objects after just a short demonstration (e.g., Welder & Graham, 2006), and they learn the physical rules of occlusion with just a single trial (e.g., Wang & Baillargeon, 2005). Other domains of knowledge in which young children also show such sophisticated inductive inferences include language (Chomsky, 1980), causality (Gopnik & Sobel, 2000), and biological kinds (Gelman & Wellman, 1991).

In addition, young learners often make generalizations at multiple levels of abstraction, which is even more important for building large conceptual structures. Not only do they make first-order generalizations (e.g. dogs like to eat bones; rabbits like to eat vegetables), but they also make sophisticated second, third, or fourth-order generalization

(e.g. each kind of animal has a favored food; each kind of animal has its own unique traits).

This view of early learning approaches the issue of the origins of inductive constraints and biases from the perspective that early input provides the basis for developing such constraints, and subsequent learning is guided by these learned constraints. Computational cognitive scientists have developed formal models, in particular Bayesian models that capture the idea of learning to learn (e.g. Kemp, Perfors & Tenenbaum, 2007) across a variety of domains, from causal learning to categorization to word learning, and from whole grammars to intuitive theories (Griffiths & Tenenbaum, 2009; Tenenbaum, Griffiths & Niyogi, 2011). Recent empirical work has also provided evidence for such a capacity early on in development: Looking-time experiments with 9-month-old infants indicate that they can form second-order generalizations such as “boxes contain objects that are uniform in color” (Dewar & Xu, 2010); and Macario, Shipley & Billman (1990) showed that 4-year-old children could rapidly construct higher-order generalizations about how objects were being categorized, successfully classifying new exemplars into novel categories by shape or color. These two lines of research, together, provide strong evidence for the first key claim advanced by rational constructivist accounts of learning.

The second key claim put forth by these theoretical proposals of cognitive development is that children can influence their own learning outcomes. Two sub-claims underlie this argument: one is that children are smart and sophisticated *processors* of data, and the other is that children may be smart *generators* of data as well: they can independently generate the data that is necessary for learning.

The first sub-claim has been well-researched over the last two decades. A myriad of experiments on this topic have produced ample evidence demonstrating that children have powerful domain-general learning mechanisms that allows them to keep track of complex statistics in their input (e.g. Aslin, Saffran & Newport, 1998; Gopnik et al., 2004; Kirkham, Johnson & Slemmer, 2002; Saffran, Aslin & Newport, 1996; Xu & Garcia, 2008; among many others). This sophisticated input *processing* also enables children to identify what to learn and when to make further inferences. For example, after infants had used the statistics in a speech stream to carry out word segmentation, they attached these newly segmented words to objects (Graf Estes, Evans, Alibali & Saffran, 2007).

In contrast, the second sub-claim has been much less studied to date. In this claim, it is argued that children can sometimes *generate* the relevant and necessary evidence themselves, even in the absence of explicit instructions or demonstrations. In some ways, this process is much more similar to real world learning, where young children engage in free play, and their attention is drawn to whichever aspects of the environment that appeal to them.

In the present study, we examine this theoretical description of the young child as a rational constructivist learner by exploring the two key claims in tandem – we investigate whether children are able to generate the relevant evidence themselves, in order to discover second-order generalizations. The formation of such generalizations is the beginning of building larger pieces of the conceptual structure, such as intuitive theories.

As mentioned, there is strong evidence that children can form higher-order generalizations based on the input they receive. However, these studies have exclusively examined children under conditions of training: infants and children receive a restricted set of evidence provided by the experimenter, and they are subsequently tested on their learning. Furthermore, the evidence presented during such training is usually helpful, created to lead children towards making the correct generalizations (e.g. training objects had shapes that were perfectly correlated with their names in Smith, Jones, Landau, Gershkoff-Stowe & Samuelson, 2002).

The results from these training studies thus beg the question of whether children can drive their own learning, such that they can successfully acquire these inductive constraints without externally-generated helpful evidence. To date, few studies have investigated this question, and there remains a gap in our knowledge about children's capacity to engage in such self-directed learning. Many related questions remain unanswered. Can children independently generate evidence? Do they gather sufficient evidence to support their own learning? Can they properly incorporate the information obtained into their own knowledge?

Developmental researchers have made some inroads into addressing these gaps by examining children's exploration and question-asking behaviors (e.g. Frazier, Gelman, & Wellman, 2009; Schulz & Bonawitz, 2007). However, these studies have focused on the information-seeking behaviors themselves, and not learning outcomes. In studies that have investigated learning, results have generally been less than promising – young children often fail to ask enough questions to gather sufficient information, and they frequently fail to remember the relevant information. For example, when 4- to 6-year-olds were told to ask questions in order to determine a card (out of an array of 24 cards) that was previously chosen by an experimenter, Legare, Mills & Souza (2013) found that the children were able to ask good questions some of the time, but the overall accuracy at test was quite low.

Yet there is reason to believe that the research on question-asking behaviors may underestimate the young child's capacity for self-directed learning. Due to the verbal nature of these tasks, the child's capacity to gather required information may be masked by difficulties in explicitly formulating and generating the necessary questions. Instead, a nonverbal task that taps into the child's ability to independently generate evidence may be more revealing of his/her capacity for active learning.

One way that young children may actively engage their environment in order to supplement their learning is through free play. Young mammals spend an extensive amount of time playing, and it has been often argued that free play is an important aspect of cognitive development, allowing the child to assimilate knowledge (Piaget, 1962), solve problems (Bruner, 1972) and create new knowledge (Vygotsky, 1978). One reason that such development can occur through free play is because playful behavior is less constrained by functional pressures, allowing the child to discover novel behavior combinations (Bruner, 1972). We thus hypothesize that during free play, children may generate the evidence necessary for their own learning, and in the context of this experiment – for the formation of appropriate higher-order generalizations.

Some previous empirical work examining free play has indicated that 4-year-old children are adept in utilizing the experience gained from either a play opportunity or a training situation to accurately complete a problem-solving task (Smith & Dutton, 1979; Simon & Smith, 1983). In one study, experimenters presented children with some sticks and blocks. Half of the children were given a play opportunity, in which they were allowed to freely explore the sticks and blocks by themselves, while the other half of the children were given a training experience, in which they followed an experimenter's instructions to learn how to join the sticks together with a block. After the play/training, the experimenter presented the child with a lure-retrieval task, in which it was necessary to join two sticks together with a block to retrieve a marble. Interestingly, children performed equally well in both conditions! Given these results, we further hypothesize that children would be equally successful in forming these generalizations based on *experimenter-generated* and *self-generated* evidence.

To test our hypotheses, we designed a causal learning experiment. Three-year-old children were provided with either a training experience in which an experimenter demonstrated the activations of three different categories of machines (Training condition), or a play opportunity in which children could freely interact with the different machines (Free Play condition). They were then asked to make first-order generalizations, where they had to choose from a new set of blocks to activate a previously seen machine, and second-order generalizations, where they had to choose from a new set of blocks to activate a novel machine. Success in the Free Play condition would suggest that young children can generate evidence to support their learning, indicating an early capacity for active learning.

## Method

### Participants

Fifty-six English-speaking 3-year-olds (22 boys and 34 girls) with a mean age of 35.9 months (range = 30.3 to 42.3 months) were tested. All were recruited from Berkeley, California, and its surrounding communities. An additional 6 children were tested but excluded due to refusal to make a choice at test ( $N = 2$ ), parental interference ( $N = 3$ ), no attempt to make any activations ( $N = 1$ ) and experimenter error ( $N = 1$ ). Each child was randomly assigned to a Training condition or a Free Play condition.

### Materials

Four categories of toy machines were used in this experiment, with two identical machines in each category. The categories differed in shape and color, i.e. machines in Category 1 were blue and rectangular; machines in Category 2 were red and triangular; machines in Category 3 were green and circular; and machines in Category 4 were orange and L-shaped (each approximately 30 cm x 10 cm x 5 cm). Each set of machines also produced a unique effect when activated, e.g. made a sound, lit up with flashing lights, or played a song.

A variety of small blocks (approximate 4 cm x 2 cm x 1 cm) with different shapes and colors were used to activate these machines. Some of these blocks matched the toy machines in shape but not color (shape-match blocks), some matched the machines in color but not shape (color-match blocks), and others did not match the machines in shape or color (distractor blocks). There was an additional cross-shaped yellow machine that was used only in the Free Play condition, and it was activated by a cross-shaped yellow block.

In the Training condition, 3 white trays with separators were used to easily present the activator blocks during the training phase and the test phase. In the Free Play condition, three plastic bins were used to present the toy machines with their corresponding activator blocks.

### Procedure

Children were tested individually in the laboratory. The parents were also present in the testing room, but sat about 80 cm behind the children throughout the experiment, in order to not influence their actions and choices. Children were introduced to the machines and blocks under the pretext of the experimenter showing them her toys.

The Training condition consisted of two phases: a training phase and a test phase; while the Free Play condition consisted of three phases: a familiarization phase, a free play phase and a test phase. Within each condition, half of the children were presented with machines that were activated using a shape rule: a shape-match block had to be used to activate the machine's effect, while the other half of the children were presented with machines that were activated using a color rule: a color-match block had to be used to activate the machine's effect (See Figure 1).

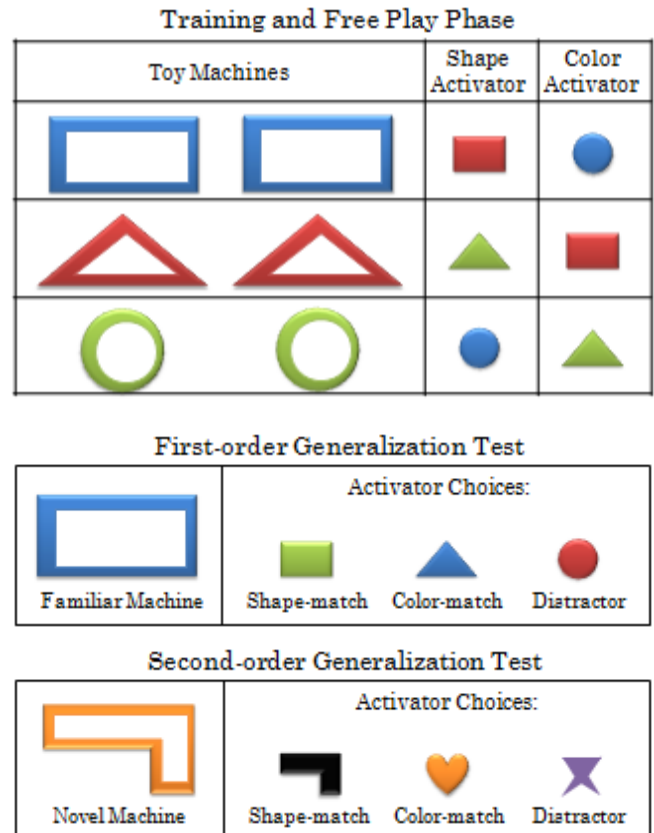


Figure 1: Schematic diagram of materials and procedure.

**Training Condition** In the Training condition, each child was seated opposite an experimenter across a long table. To begin the training phase, the experimenter presented a white tray containing three blocks differing in shape and color. The child was free to play with these blocks for about 20 seconds. After this exploration, the blocks were returned onto the tray and pulled close to the experimenter, but remained visible to the child.

The experimenter then presented the first toy machine (e.g. blue rectangular machine), and activated the machine with one of the three blocks by placing it on top of the machine (e.g. red rectangular block, if the machines were being activated by a shape rule; blue triangular block, if the machines were being activated by a color rule). Upon the machine's activation, the experimenter drew attention to the event by saying, "Look! The block made the machine go; it made it go!" The experimenter next showed the child another machine that was identical to the first one, and activated it using the same block. This first set of two machines was then cleared from the table. The experimenter repeated this procedure with two other sets of training machines, activating them with their respective shape-match or color-match blocks.

A total of six machines were presented during the training phase, and each child saw each machine being activated only once. The three categories of machines chosen as the training set were randomized, leaving the fourth category of

machines for the test phase (i.e. each category could be used as a training machine or a test machine). The order of presentation for the categories of training machines was also counterbalanced. The duration of the training phase was about 4 minutes.

A test phase immediately followed the training phase. The test phase consisted of a first-order generalization test and a second-order generalization test. In the first-order test, each child was presented with a *familiar machine*, which is a machine that was previously presented in the training phase. Then, the child was provided with 3 novel choice blocks in a white tray: a shape-match block, which is similar to the target machine in shape but not color, a color-match block, which is similar to the target machine in color but not shape, and a distracter block, which differed from the target in both color and shape. The experimenter requested the child to hand her a block that made the target machine go, “Can you give me the block that makes this machine go?”

In the second-order test, each child was presented with a *novel machine*, which is a machine that was not previously presented in the training phase. The child was again asked to activate the machine by choosing among 3 novel choice blocks: a shape-match block, a color-match block, and a distracter block. The duration of the test phase was about 1 minute.

**Free Play Condition** In the Free Play condition, each child sat opposite an experimenter on the floor. To begin the familiarization phase, the experimenter presented the child with a cross-shaped yellow machine, together with its activator block. This block matched the machine both in shape and color. The familiarization phase served to introduce the child to the sound-making function of these novel machines. This phase was not necessary earlier, since the machines’ function would be introduced in the course of training. The experimenter then activated the machine, drawing attention to the event by saying, “Look! The block made the machine go; it made it go!” The child was then given the activator block, and was allowed to activate the machine freely. The experimenter ensured that each child saw at least two activations of this familiarization machine.

A free play phase followed the familiarization phase, and this phase began by the experimenter exclaiming, “Oh no! I just remembered that I have some work to do. While I’m doing my work, you can play with some of my toys!” The experimenter then laid out three plastic bins, each consisting of two identical machines and their corresponding activator block. The experimenter subsequently moved to a table and pretended to work, telling the child, “You can go ahead and play!” Each child was given 5 minutes to play freely with the machines and blocks. After 5 minutes, the experimenter announced that she was done with her work and that it was time to put the toys away.

The test phase that immediately followed the free play was identical to that of the Training condition.

## Coding

The children’s responses in the test trials were scored for accuracy. For the children who were exposed to the shape rule during the training or free play phases, choosing a shape-match block was scored as 1 point. Correspondingly, for children exposed to the color rule, choosing a color-match block was scored as 1 point. The maximum score for each child was 2 points, as there were 2 test trials in total. The children’s scores were then converted into percentage of accuracy. A second coder recoded all of the children’s responses, and the level of agreement between the coders was 100%.

## Results

An alpha level of 0.05 was used in all statistical analyses. Preliminary analyses found no effects of gender, median age-split (whether the children were younger or older than the median age of the group), trial order (first trial vs. second trial), presentation order of the training machines (e.g. whether machines from Category 1 were presented first, second or third during the training phase), and rule type (shape rule vs. color rule) on children’s accuracy on the test trials. Subsequent analyses were collapsed over these variables.

Due to the free play nature of the Free Play condition, several parameters varied across children in this condition: the number of activations for each category of machines ( $M = 5.07$ ,  $SD = 5.12$ ; recall that children in the Training condition each saw 2 activations per category of machines), the number of times that “negative evidence” was generated, as defined by the number of times the child placed an activator block on a machine from a different bin ( $M = 3.7$ ,  $SD = 5.11$ ), and the length of the free play phase ( $M = 226s$ ,  $SD = 126.2s$ ). 75% of the children in the Free Play condition activated every category of machines that was presented during the free play phase.

As Figure 2 shows, children performed accurately during the test trials, selecting the correct activator block in both the Training and the Free Play conditions. Using the children’s percentage accuracy, a 2x2 repeated measures analysis of variance (ANOVA) was performed with Condition (Training vs. Free Play) as a between-subjects factor and Type of Generalization (first-order generalization vs. second-order generalization) as a within-subjects factor. There was neither a main effect of Condition,  $F(1, 54) = .011$ ,  $p = .98$ , nor Type of Generalization,  $F(1, 54) = .54$ ,  $p = .47$ ,  $\eta^2 = .01$ , and there was also no interaction between the two factors,  $F(1, 54) < 1$ .

Critically, we were interested in the effects of training and free play on children’s accuracy on the test trials. We used a conservative chance criterion of .5, even though children were offered three activator choices at each test trial. Planned comparisons indicated that children in the Training condition were significantly more likely to choose the correct activator block ( $M = .72$ ,  $SD = .38$ ) as compared to chance,  $t(31) = 3.26$ ,  $p = .003$ ,  $d = .57$ . Children in the Free Play condition were also significantly more likely to choose

the correct activator block ( $M = .73$ ,  $SD = .36$ ) as compared to chance,  $t(23) = 3.11$ ,  $p = .005$ ,  $d = .62$ . An additional comparison also revealed that children in the two conditions were equally likely to choose the correct activator block in both the first-order generalization test,  $t(54) = 1.65$ ,  $p = .870$ ,  $d = .45$ , and the second-order generalization test,  $t(54) = 0$ ,  $p = 1.00$ ,  $d = 0$ .




Figure 2: Percentage accuracy in the two conditions. Dashed line indicates chance performance. Error bars represent SE.

## Discussion

The present study examined whether three-year-old children can form higher-order generalizations in the causal domain, based on both *experimenter-generated* and *self-generated* evidence. The results demonstrate that the children can, and they are *equally effective* in learning from both types of evidence. In both Training and Free Play conditions, three-year-olds rapidly made first-order and second-order generalizations about how the machines and the activator blocks interacted with one another, and they extended these generalizations appropriately to novel toy machines.

These results make two important new contributions to the literature. First, children's success in the Training condition constitutes the first demonstration that children can simultaneously acquire generalizations at multiple levels in the causal domain. Developmental research in this area have largely focused on word learning and categorization tasks, leaving open the question of whether "learning to learn" is a perspective that is limited to only a few specific domains. Given that causal knowledge constitutes the foundation of intuitive theories, and that these theories are present in multiple domains (Carey, 1985; 2009; Gopnik & Meltzoff, 1996), there is reason to believe that the ability to "learn to learn" is a domain-general one.

Second, and even more striking, is children's robust success in the Free Play condition. Children consistently chose the correct blocks to activate the machines presented in both the first-order and second-order generalization tests.

This success suggests that even in the absence of explicit instructions, children can, in the course of free play, generate the relevant evidence themselves, in order to form appropriate generalizations at multiple levels. This learning condition is much closer to what children encounter in the real world, where preschoolers are often allowed to play freely, and engage with whatever aspects of the environment they find interesting and appealing. Our data provide strong evidence that (1) preschool children are motivated to understand what rules govern the behavior of the objects around them, and (2) they may have some rudimentary capacity to systematically generate the relevant evidence to support such learning.

Ongoing work in our laboratory takes a closer look at the types of evidence generated by children during the free play phase, and how these different types of evidence is related to their subsequent accuracy at the generalization tests. Preliminary analyses show that the number of activations and the number of negative evidence generated predicted accuracy at test.

Several previous studies have shown that young learners' attention is allocated in systematic ways that reflects active learning. For example, Kidd et al. (2012) showed that 8-month-old infants preferentially looked at series of stimuli that provide the most information gain; Gerken, Balcomb & Minton (2011) found that 17-month-old infants devote more attention to aspects of their environment that are learnable, rather than unlearnable; and Schulz and Bonawitz (2007) demonstrated that preschoolers selectively played with a box that had produced ambiguous evidence for its causal structure. These studies, however, have not shown whether children acquired specific pieces of knowledge through their own free play and visual exploration. Here we provide the first clear demonstration that preschoolers are capable of doing so. Interestingly, using the same task, preliminary results from our lab suggest that toddlers may not be able to generate the necessary evidence to support their learning. Future work will focus on charting the developmental trajectory for children's ability to engage in active learning.

Put together, the present study provides strong evidence for emerging theoretical proposals that children are rational constructivist learners (e.g. Gopnik & Wellman, 2012; Xu & Kushnir, 2012; 2013). Our results are consistent with the two key claims put forth by these proposals: first, children form principled and meaningful generalizations based on the inputs from the environment, and second, they are self-directed learners, actively engaging different aspects of their environment to supplement their own learning. Future research should also investigate the optimality of the active learning that children partake in, as well as its limits, to shed light on how early learning actually occurs in the real world.

## Acknowledgments

We thank Sarah Blacher, Sophie Beiers, and the Berkeley Early Learning Lab (BELL) for their help in testing and recruitment, as well as the parents and children for their

participation. We also thank Caren Walker and Alison Gopnik for insightful discussions.

## References

- Aslin, R. N., & Newport, E. L. (2012). Statistical Learning: From Acquiring Specific Items to Forming General Rules. *Current Directions in Psychological Science*, 21(3), 170–176.
- Bruner, J. S. (1972). Nature and uses of immaturity. *American Psychologist*, 27(8), 687.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.
- Carey, S. (2009). *The origin of concepts*. Oxford University Press, USA.
- Chomsky, N. (1980). Rules and representations. *Behavioral and Brain Sciences*, 3(1), 1-15.
- Dewar, K. M., & Xu, F. (2010). Induction, overhypothesis, and the origin of abstract knowledge. Evidence from 9-month-old infants. *Psychological science*, 21(12), 1871–7.
- Frazier, B. N., Gelman, S. A., & Wellman, H. M. (2009). Preschoolers' search for explanatory information within adult-child conversation. *Child development*, 80(6), 1592-1611.
- Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings non-obvious. *Cognition*, 38, 213–244.
- Gerken, L., Balcomb, F. K., & Minton, J. L. (2011). Infants avoid “labouring in vain” by attending more to learnable than unlearnable linguistic patterns. *Developmental science*, 14(5), 972–9.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological Review*, 111(1), 3–32.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Sobel, D. M. (2000). Detectingblickets: how young children use information about novel causal powers in categorization and induction. *Child development*, 71(5), 1205–22.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: causal models, Bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6), 1085–108.
- Graf Estes, K., Evans, J. L., Alibali, M. W., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological science*, 18(3), 254–60.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological review*, 116(4), 661–716.
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental science*, 10(3), 307–21.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, 7(5), e36399.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35-B42.
- Legare, C. H., Mills, C., & Souza, A. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, 114, 63–76.
- Macario, J. F., Shipley, E. F., & Billman, D. O. (1990). Induction from a single instance: formation of a novel category. *Journal of experimental child psychology*, 50(2), 179–99.
- Piaget, J. (1962). *Play, dreams and imitation* (Vol. 24). New York: Norton.
- Saffran, J. R. J., Aslin, R. N. R., & Newport, E. L. (1996). Statistical Learning by 8-Month-Old Infants. *Science*, 274(5294), 1926 – 1928.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, 43(4), 1045–50.
- Simon, T., & Smith, P. (1983). The study of play and problem solving in preschool children: Have experimenter effects been responsible for previous results? *British Journal of Developmental Psychology*, (1), 289–297.
- 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), 13–9.
- Smith, P. K., & Dutton, S. (1979). Play and Training in Direct and Innovative Problem Solving. *Child development*, 50, 830–836.
- Tenenbaum, J. B., Griffiths, T. L., & Niyogi, S. (2007). Intuitive theories as grammars for causal inference. In A. Gopnik, & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation*. Oxford: Oxford University Press.
- Vygotsky, L. S. (1978). The role of play in development. *Mind in society*, 92-104.
- Wang, S., & Baillargeon, R. (2005). Inducing infants to detect a physical violation in a single trial. *Psychological science*, 16(7), 542–9.
- Welder, A. N., & Graham, S. A. (2006). Infants' categorization of novel objects with more or less obvious features. *Cognitive psychology*, 52(1), 57–91.
- Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 105(13), 5012–5.
- Xu, F., & Kushnir, T. (2013). Infants Are Rational Constructivist Learners. *Current Directions in Psychological Science*, 22(1), 28–32.
- Xu, F. & Kushnir, T. (Eds.) (2012) Rational Constructivism in Cognitive Development. *Advances in Child Development and Behavior*, Vol. 43. Waltham, MA: Academic Press.