

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Children Learn Better When They Select Their Own Data

Permalink

<https://escholarship.org/uc/item/2km363bs>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Sim, Zi L

Tanner, Michelle

Alpert, Nina Y

et al.

Publication Date

2015

Peer reviewed

Children Learn Better When They Select Their Own Data

Zi L. Sim (zi@berkeley.edu)

Michelle Tanner (michellemtanner@gmail.com)

Nina Y. Alpert (nalpert@berkeley.edu)

Fei Xu (fei_xu@berkeley.edu)

Department of Psychology, University of California, Berkeley
Berkeley, CA 94703 USA

Abstract

Human learners ask questions, manipulate objects, and perform interventions on their environment. These behaviors are true of adults, but even more so for young children. Recent studies have demonstrated that adults learn better under conditions of *selection learning*, where they can make decisions about the information they wish to acquire, as compared to *reception learning*, where they merely observe data that happens to be available to them. Yet to date, it remains unclear whether this advantage is available to children, and if so, does it arise because children can gather data in a non-random way? In the current study, we show that 7-year-old children show superior learning under conditions of selection in a category-learning task, and that their information gathering is systematically driven by uncertainty.

Keywords: self-directed learning; active learning; education

Introduction

“You’re speaking too loudly! No, that’s too soft; you have to speak up!” What volume do these adults mean exactly? As a young child, learning how to modulate our speaking volume is an important aspect of learning how to socialize with others. However, instructions that adults give can be quite opaque at times. What is a child to do?

One solution might simply be to observe what others are doing; taking note of the volumes that they are speaking at. Or one might choose to actually collect the necessary data: try a variety of different volumes (hopefully spread out across time!), and observe how mommy responds.

Children probably use a mix of strategies to learn the right volume to speak with, but as can be seen from the above example, there are at least two modes of learning that people engage in to refine their knowledge about the world: *reception learning*, in which learners merely observe data that happens to be available and attempt to find structure within them; and *selection learning*, in which learners are allowed to make decisions about the information they wish to acquire (Bruner, Goodnow, & Austin, 1956; Bruner, 1961).

Much of cognitive research has focused on the former mode of learning. Researchers study category and concept learning in experiments where they tightly control the exemplars that are presented to the participants (e.g. Medin & Schaffer, 1978; Shepard, Hovland, & Jenkins, 1961). Language learning has also traditionally been examined in the laboratory by presenting infants and young children with repetitive sentences, speech streams, or word-object pairings

(e.g. Saffran, Aslin, & Newport, 1996; Waxman & Gelman, 2009; Xu & Tenenbaum, 2007).

Selection learning, in contrast, has found its niche mostly in the domain of causal learning, because certain causal networks can only be distinguished with data gained from intervention, rather than mere observation. In other words, the data generated by intervention simply cannot be acquired through observation. In such cases, researchers have empirically shown that selection learning has distinct advantages over reception learning (Sobel & Kushnir, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). For example, Sobel and Kushnir (2006) demonstrated that when learners observed the data that they generated themselves, they were better at learning the underlying causal structure than learners who observed data that others generated.

Recent cognitive research with adults has gone on to study this advantage in domains outside of causal learning, especially in domains where it is possible to generate the same information from both selection and reception learning (Gureckis & Markant, 2012). Studies by Castro et al. (2008) and Markant and Gureckis (2013) have successfully shown that learners benefit from selection in category learning as well. In Castro et al. (2008), adults were presented with novel 3D shapes that varied continuously only in how spiky their edges were. They were told that these shapes were alien eggs: spiky eggs would most likely hatch into alien *snakes*, while smooth eggs would most likely hatch into alien *birds*. The task for each participant was therefore to find out the precise egg shape (category boundary) for which eggs that were any spikier would hatch into snakes, while eggs that were any smoother would hatch into birds. Critically, participants in a selection condition were allowed to choose which eggs they wanted to learn about, while participants in a reception condition were presented with randomly generated egg shapes. Both groups observed whether each egg hatched into a snake or a bird after each selection or presentation.

The results of this study were striking. Participants who were allowed to actively select samples to learn about had more accurate guesses about the category boundary as compared to participants who could only observe samples that were randomly generated for them. This result was successfully replicated in Markant and Gureckis (2013) using a slightly modified procedure. However, these results do come with some caveats: the selection advantage is only present at low noise levels, i.e. when the spiky and smooth eggs reliably hatched into snakes and birds respectively

(Castro et al., 2008), and in low complexity tasks, e.g. when the classification rule is based on only one dimension (e.g. spikiness only), rather than multiple dimensions (e.g. a combination of spikiness and size).

The same authors also performed a comprehensive analysis aimed at uncovering the psychological processes underlying the found selection advantage, concluding that learners benefit from selection learning because they can gather data in a “non-random, useful way” that maximizes their *own* future learning (Markant & Gureckis, 2013).

Do these results naturally extend to young children? It is indubitable that young children often engage in some forms of selection learning; one only needs to recall their incessant questions, or their mucking around the house and whatnot. Does this effortful form of learning where children have to both generate and learn from the data benefit them, as compared to the less demanding form of learning where they simply observe data that happens to be available to them? If so, are the psychological processes underlying the selection advantage similar between children and adults? When given the opportunity, do children gather data in a “non-random, useful way”? Addressing these questions would provide insights into the developmental origins of selection learning and its underlying mechanisms.

However, these questions remain mostly unaddressed in the literature. At this point, there is still a lack of empirical evidence demonstrating that children actually benefit from selection, relative to reception. What we do know, though, is that young children may be able to gather data in non-random manner (Cook, Goodman, & Schulz, 2011; Kidd, Piantadosi, & Aslin, 2012; Legare, Mills, & Souza, 2013; Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014; Ruggeri & Lombrozo, under review; Schulz & Bonawitz, 2007; Sim & Xu, 2014). For example, Schulz and Bonawitz (2007) showed that preschoolers prefer to explore a toy for which the causal structure remained ambiguous to them, over a completely novel toy. Nelson et al. (2014) also demonstrated that 10-year-old German children had good intuitions about how useful various questions would be in sequential search tasks that resembled games such as “Guess Who?” Children were also able to search adaptively, varying their questions according to the statistical structure of the environment they were presented with (e.g. when the population in a “Guess Who?” game was modified such that asking about gender first would no longer be quite as useful, children were less likely to ask about it at the beginning).

But such evidence does not necessarily imply that children will benefit from selection over reception when it comes to refining their beliefs about the world. Indeed, it would be quite a leap to make the claim that just because children are exploring in a systematic way, they are learning from that form of exploration.

Furthermore, although Castro et al. (2008) provides a formal proof for the advantage of selection learning over reception learning in deterministic (noise = 0) environments, there is currently no evidence that children are optimal in their information gathering either. Without this evidence, it

is difficult to support a theoretical argument that selection is necessarily more efficient than reception for learning.

To begin examining selection learning in children, at minimum, we need to establish three points within the very same task: (1) children can learn successfully under conditions of selection, (2) they can gather data in a systematic manner, and (3) selection learning has distinct advantages over reception learning. We address these points in the current study by examining whether children perform better at a category-learning task when they can select the information they wish to acquire, as compared to when they are merely presented with randomly generated data.

In an experimental design inspired by Castro et al. (2008) and Markant and Gureckis (2013), 7-year-old children were presented with a row of identical worms that were ascending in size, and told that the worms live in either a green house or a blue house. The house that each worm lives in depended on its size, so the goal of the game was to figure out the category boundary as quickly as possible in order to bring them home before a thunderstorm arrives. Each child was randomly assigned to one of two conditions: selection, where they could choose sequentially which worms to learn about, or reception, where they were presented with randomly generated worms one after the other. There were 4 test blocks, and within each block, children learned about 2 worms and then were given a classification task. The design of this task allowed us to examine the children’s learning performance and their information gathering strategy (for example, were children taking advantage of feedback generated by previously selected worms?) when they are given the opportunity to actively make decisions about the information they wish to acquire.

Method

Participants

Sixty-four English-speaking 7-year-olds (23 boys and 42 girls) with a mean age of 88.4 months (range = 74.6 to 104.3 months) were tested. All were recruited from schools and museums in Berkeley, California, and its surrounding communities. An additional 8 children were tested but excluded due to difficulties in following task instructions (e.g. indicating that a worm, which had a little blue reminder house beneath it, lived in the green house; $N = 6$), technical error ($N = 1$), and experimenter error ($N = 1$). Each child was randomly assigned to a Selection condition or a Reception condition.

Materials

The experiment was presented in the form of an interactive PowerPoint presentation. Each presentation sequentially showed 3 sets of animals, with each set consisting of 13 identical animal images that varied only in their size, i.e. their heights and widths.

These animals were arranged from smallest to largest (left to right). The animals lived in either a green house or a blue house, and these houses were represented by colored images

placed on the top left and top right of the screen respectively (Figure 1). When an animal image was clicked on, it would move across the screen towards its designated house, disappearing upon arrival. A “reminder house,” which is a scaled down version of its house, would then appear in the space below where the animal was located previously.




Figure 1: Thirteen worms that can be categorized into the green house or the blue house.

Procedure

Children were tested individually in our laboratory, a quiet room in their elementary school, or in a quiet area at a museum. An experimenter sat next to the child to control the slide show. The procedure for both the Selection and Reception conditions consisted of a demonstration phase, and 4 test blocks (each with 2 sampling trials). Each block consisted of a sampling phase, followed by a classification phase. The experiment lasted about 10 minutes.

Demonstration Phase The demonstration phase consisted of two practice trials. These practice trials were to establish to the child that (1) the displayed animals lived in one of the two houses, (2) the house that each animal lived in was determined by an invisible category boundary that divided the animals into two groups, and (3) that the boundary location was different for each set of animals.

In the first practice trial, the participant was shown a row of 13 spiders that increased in size, together with a green house and a blue house placed at the top corners of the screen. When the experimenter clicked on each house, a flashing box surrounding the spiders that lived in the selected house appeared. The experimenter subsequently pointed at two spiders, one at a time, asking the child “Does this spider live in the green house or the blue house?” The experimenter praised the child if he/she answered accurately (“Good job!”), and corrected the child otherwise (“No, that spider actually lives in the green house!”).

The second practice trial that followed was identical to the first, except that we used a row of frogs instead, and a new category boundary.

Test Block: Sampling Phase Children were presented with a row of 13 worms. 12 category boundaries were possible,

but only the 3rd through the 10th boundary were used in this experiment. This step was taken to ensure that there was at least a small number of worms that lived in each house. For each participant, a boundary location was randomly generated, and this location was used for all test blocks.

To begin the sampling phase, the experimenter informed the child that she would be asked to figure out which house each worm lived in. The experimenter then clicked on the 1st and 13th worm, showing that they lived in the green house and the blue house respectively. As mentioned above, an appropriately colored “reminder house” subsequently appeared below the worm that had just been selected.

An image of a storm then appeared. In the Selection condition, the experimenter told the child that there was only time left to tap on one worm, and asked the child to choose one worm to “figure out which worms live in the green house, and which worms live in the blue house”. The experimenter clicked on the chosen worm, which moved to its given house as determined by the category boundary. The child was then told that the storm had not arrived yet, so there was still time to learn about another worm. After the child made this second selection, the experimenter clicked on the worm to show where it lived. Reminder houses appeared after each worm was selected. The key feature in the Selection condition was thus that the child was allowed to independently generate data about the worms in order to learn about their category structure.

In the Reception condition, a program was ran such that one worm would be randomly selected at appropriate time points. Based on information obtained about children’s choices during pilot testing of the Selection condition, the script was constrained such that 1) a single worm cannot be selected twice within each critical block, and 2) a previously selected worm can be reselected in a later critical block. Within each test block, two worms were randomly selected one after another. Upon being selected, the worm wiggled to attract the child’s attention before moving to the house that it lived in. Again, reminder houses appeared to provide a visual memory aid of where the selected worms lived. The key feature in the Reception condition was therefore that the child could only observe, but not generate, data about the worms to learn about their category structure.

Test Block: Classification Phase After the sampling phase, the experimenter informed the child that the storm was almost here, so they had to take the rest of the worms home. The child was asked to point to all the worms that lived in the green house, as well as all the worms that lived in the blue house. If the child skipped the classification of some worms, the experimenter pointed to each of these skipped worms and asked, “Which house does this worm live in?” The children’s answers allowed us to determine where they believed the boundary was located. After all the worms had been classified, they disappeared and the experimenter told the child, “Phew, all the worms are safe! But we don’t know if they went to their correct houses.”

The test blocks were repeated until the child had classified all the worms correctly, or when the child had engaged in 4 test blocks (i.e. viewed a maximum of 8 worms), whichever occurred first.

Coding

In the Selection condition, we recorded the worms that each child selected during the sampling phase. We then measured the sampling distance, i.e. the distance between each of their selections and the true category boundary. For example, if the child selected a worm that was adjacent to the category boundary (left and right), the sampling distance was 0. The sampling distances allowed us to examine how children were sampling across time. This measure was recorded in the Reception condition as well, although note that these “selected” worms were randomly generated.

For each child, we also obtained a classification accuracy score for all test blocks. Each correctly classified worm was scored as 1 point, so the maximum score in each block was 13. The children’s scores were then converted into a percentage of classification accuracy.

Results

An alpha level of 0.05 was used in all statistical analyses. Preliminary analyses found no effects of gender or location of boundary on children’s accuracy on classification trials. Subsequent analyses were collapsed over these variables.




Figure 2: Sampling distance from the category boundary in the two conditions. Dashed line indicates average sampling distance expected by a random-sampling strategy. Error bars show standard error.

Information Sampling

After learners have acquired some data in a category learning task, they would easily classify items that are far from the true category boundary, but are more uncertain about items that are near the boundary. Following the analyses in Markant & Gureckis (2013), we thus examined children’s sampling distances, i.e. the distance between the

children’s selections and the true category boundary, as a general measure of uncertainty-driven information selection.

As Figure 2 indicates, children in the Selection condition were sampling closer to the true category boundary over time. Using the children’s average sampling distance for each test block, we performed a 2x4 repeated measures analysis of variance (ANOVA) with Condition (Selection vs. Reception) as a between-subjects factor and Test Block (1–4) as a within-subjects factor. There were significant main effects of Condition, $F(1, 62) = 15.2, p < .001, \eta^2 = .197$, and Test Block, $F(3, 60) = 10.77, p < .001, \eta^2 = .350$. There was also a significant interaction between the two factors, $F(3, 61) = 8.58, p < .001, \eta^2 = .30$.

Planned comparisons showed that average sampling distance in the Selection condition was significantly smaller than expected by a random-sampling strategy by the second test block, $t(31) = 2.34, p = .026, d = .413$, while the average sampling distance of the randomly generated data points in the Reception condition never differed from chance, e.g. in the fourth test block, $t(31) = .684, p = .50, d = .121$.




Figure 3: Classification accuracy in the Selection and Reception conditions. Error bars show standard error.

Classification

Using children’s average classification accuracy across the four blocks, we then performed a 2x4 repeated measures analysis of variance (ANOVA) with Condition (Selection vs. Reception) as a between-subjects factor and Test Block (1–4) as a within-subjects factor. There was only a main effect of Test Block, $F(3, 60) = 14.5, p < .001, \eta^2 = .42$. No other main effects or interaction was found.

Planned comparisons revealed that the overall classification accuracy for children in both the Selection ($M = .935, SD = .064$) and Reception conditions ($M = .924, SD = .059$) was significantly different from chance (0.689). For the Selection condition, $t(31) = 21.8, p < 0.01, d = 3.85$. For the Reception condition, $t(31) = 22.7, p < 0.01, d = 4.01$.

Although children’s classification accuracy did increase steadily in both conditions, their classification accuracy diverged over time. By the final block, children in the Selection condition were significantly more likely to classify the worms correctly ($M = .986, SD = .046$) than

children in the Reception condition ($M = .954$, $SD = .064$), $t(62) = 2.24$, $p = .029$, $d = .574$.

Discussion

The present study examined whether 7-year-old children had the capacity to engage in and benefit from selection learning. Using a category learning task, we demonstrate that young children can learn successfully under conditions of selection, that they can gather data in a systematic manner, and that selection learning has distinct advantages over reception learning.

First, our results indicate that children can learn successfully when they are allowed to make decisions about what information they wish to gather. The overall classification accuracy in the Selection condition was very high, suggesting that children are perfectly capable of learning from the data they generate by themselves. Their performance was comparable to that of children in the Reception condition, the latter of which should not be surprising given previous research showing that children are proficient at learning categories using randomly-generated exemplars when the classification rule is based only on a single dimension (i.e. rule-based category structure) similar to that used in our experiment (Huang-Pollock, Maddox, & Karalunas, 2011; Minda, Desroches, & Church, 2008). It should also be noted that the task may have been too easy for children, resulting in near-ceiling performance in both conditions. Ongoing work improves the current design by increasing the number of classification items and removing the “reminder houses.”

Second, 7-year-olds are able to gather data in a systematic way. As our results show, children sampled closer to the true category boundary over time. This result suggests that the children’s information gathering was informed by uncertainty and previous feedback, leading them to sample items that were near the true category boundary. Such a strategy would allow children to avoid generating redundant information, and focus on collecting data that is expected to help them learn effectively and efficiently.

Third, and most importantly, children showed better learning under conditions of selection as compared to reception over time. By the final block, the classification accuracy obtained by children in the Selection condition was reliably higher than that of children in the Reception condition. Given the extremely small amount of information observed by the children over four blocks (as compared to previous adult studies), we found this measure to be more revealing of children’s learning under different modes of information gathering than that of average classification accuracy, which unduly weighs children’s early guesses.

Establishing these findings within a single task suggests that children benefit from selection learning over reception learning partly because they are able to gather data in a systematic, non-random fashion. Researchers have previously examined the systematicity and optimality of children’s exploration strategies, but few have shown that these strategies have consequences on children’s learning.

The current study thus adds an important piece to the puzzle by demonstrating that when given the opportunity, children can gather data in a systematic manner, and this uncertainty-driven data generation is associated with superior performance during category learning.

One notable difference between the selection and the reception conditions is that the learners observed different data points. Thus, to further establish the advantages of self-directed data generation, ongoing work in our lab examines how children perform in a “yoked” condition (Gureckis & Markant, 2012). In such a condition, each child will be presented with the same sequence of worms that was generated by another child in the Selection condition. If the learners in the Selection and the Yoked condition perform differently despite having observed the same data, this result would provide additional evidence that being able to gather data that systematically addresses one’s *own* regions of uncertainty is crucial for selection to result in more effective and more efficient learning (Markant & Gureckis, 2013).

Our discussion above offers a cognitive explanation for the selection advantage. Children performed better under conditions of selection because they generated data that was informative for them. However, the present results cannot speak directly to other psychological processes that may also drive the advantage found for selection learning. A variety of different psychological factors have been posited to account for such an advantage: enhanced memory encoding (Metcalf & Kornell, 2005); deeper processing of the problem structure (Sobel & Kushnir, 2006); attention and motivation (Corno & Mandinach, 1983; Kersh, 1962), etc. Given the design of our experiment in which children in both conditions were provided with visual reminders of the house that each worm lives in, we are inclined to believe that the advantage found for selection over reception learning cannot be attributed to enhanced encoding of the presented information. As for other psychological factors, we do not think that they run contrary to our arguments – after all, those processes could have certainly been recruited when children were deciding which items to learn about.

Another important note is that even though we have demonstrated that children learn better in the Selection condition as compared to the Reception condition, it is highly unlikely that the children’s information gathering was normatively optimal. In this two-category learning task, the optimal strategy is to engage in a binary search, such that the learner should always sample the item that is in the middle of the region of uncertainty (e.g. the space between the worm that one is certain lives in the green house, and the worm that one is certain lives in the blue house). By using such a strategy, the learner’s error in estimating the category boundary should exponentially converge (Castro et al., 2008). In our task, optimal learners need to sequentially sample at least 3 worms, but at most 4 worms, to discover the category boundary. However, most of the children in the Selection condition did not appear to have used such a strategy, as only 7 out of 32 children successfully classified all the worms in Test Block 2 (having sampled 4 worms).

Thus, like adults (Castro et al. 2008), children were not able to take full advantage of being able to select their own data.

Self-directed learning has been a hugely influential and long-standing debate in education. While educators have consistently encouraged their young students to engage in hypothesis testing and self-directed exploration in order to boost learning, there has been a relative dearth of empirical evidence supporting such a belief. Our results provide strong evidence that in a simple two-category learning task, children do perform better under conditions of selection, and this phenomenon stems from them being able to gather information in a systematic, non-random way. That being said, we believe that self-directed learning might not necessarily be beneficial at all developmental levels, or in all situations (Castro et al., 2008; Markant & Gureckis, 2013). More research is thus necessary to plug these gaps before work in this field can properly guide educators.

Acknowledgments

We thank Kuldeep Mahal and the Berkeley Early Learning Lab (BELL) for their help in testing and recruitment, as well as the parents and children for their participation.

References

- Bruner, J. (1961). The act of discovery. *Harvard Educational Review*, 31, 21–32.
- Bruner, J., Goodnow, J., & Austin, G. (1956). *A Study of Thinking*. New York, NY: John Wiley and Sons.
- Castro, R., Kalish, C., Nowak, R., Qian, R., Rogers, T., & Zhu, J. (2008). Human active learning. *Advances in Neural Information Processing Systems (NIPS)*, 21, 241–249.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3), 341–9.
- Corno, L., & Mandinach, E. B. (1983). The role of cognitive engagement in classroom learning and motivation. *Educational Psychologist*, 18(2), 88–108.
- Gureckis, T. M., & Markant, D. B. (2012). Self-Directed Learning: A Cognitive and Computational Perspective. *Perspectives on Psychological Science*, 7(5), 464–481.
- Huang-Pollock, C. L., Maddox, W. T., & Karalunas, S. L. (2011). Development of implicit and explicit category learning. *Journal of Experimental Child Psychology*, 109(3), 321–335.
- Kersh, B. Y. (1962). The motivating effect of learning by directed discovery. *Journal of Educational Psychology*, 53(2), 65–71.
- 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.
- Landau, B., Smith, L., & Jones, S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3(3), 299–321.
- 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.
- Markant, D. B., & Gureckis, T. M. (2013). Is It Better to Select or to Receive? Learning via Active and Passive Hypothesis Testing. *Journal of Experimental Psychology. General*, 142(2).
- Medin, D., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207–238.
- Metcalfe, J., & Kornell, N. (2005). A Region of Proximal Learning model of study time allocation. *Journal of Memory and Language*, 52, 463–477.
- Minda, J. P., Desroches, A. S., & Church, B. a. (2008). Learning rule-described and non-rule-described categories: a comparison of children and adults. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 34(6), 1518–1533.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, 130(1), 74–80.
- Ruggeri, A., & Lombrozo, T. (under review). Learning By Asking: How Children Ask Questions To Achieve Efficient Search.
- 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.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75(13), 1–42.
- Sim, Z., & Xu, F. (2014). Acquiring Inductive Constraints from Self-Generated Evidence. In P. Bello, M. Guarini, M. McShane, & S. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1431–1436). Austin, Texas: Cognitive Science Society.
- Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, 34(2), 411–9.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453–489.
- Waxman, S. R. S., & Gelman, S. a S. (2009). Early word-learning entails reference, not merely associations. *Trends in Cognitive Sciences*, 13(6), 258–263.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–72.