

Sampling to learn words: Adults and children sample words that reduce referential ambiguity

Martin Zettersten (zettersten@wisc.edu) & Jenny Saffran (jsaffran@wisc.edu)

Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706 USA

Abstract

How do learners gather new information during word learning? We present evidence that adult learners will choose to receive additional training on object-label associations that reduce ambiguity about reference during cross-situational word learning. This ambiguity-reduction strategy is related to improved test performance. We find mixed evidence that children (4–8 years of age) show a similar preference to seek information about words experienced in ambiguous word learning situations. In an initial experiment, children did not preferentially select object-label associations that remained ambiguous during cross-situational word learning. However, this may be explained by some children having relatively high certainty about object-label associations for which they did not see evidence disconfirming their initial hypothesis. In a second experiment that increased the relative ambiguity of two sets of novel object-label associations, we found evidence that children preferentially make selections that reduce ambiguity about novel word meanings.

Keywords: cross-situational word learning; mutual exclusivity; active learning; self-directed learning; sampling

Introduction

What makes us seek out new information during learning? One proposal is that information-seeking behavior is driven by uncertainty reduction (e.g., Kidd & Hayden, 2015). A variety of studies have demonstrated that – at least in some contexts – children are motivated to gather information to reduce the uncertainty after ambiguous or surprising events (Schulz & Bonawitz, 2007; Stahl & Feigenson, 2015).

To what extent does ambiguity-reduction play a role in word learning? A classic problem is how learners disambiguate the meaning of words in potentially ambiguous situations (Quine, 1960). One solution is that children can disambiguate word meanings by tracking co-occurrences of object-label pairs across multiple ambiguous situations (Yu & Smith, 2007). This proposal would be particularly powerful if learners are naturally drawn to isolating object-label associations that have remained ambiguous over the course of past learning (Hidaka, Torii, & Kachergis, 2017). A previous study of cross-situational word learning has shown that being able to actively select sets of object-label pairs to learn about increases participants' accuracy compared to a passive condition in which random sets of objects are presented (Kachergis, Yu, & Shiffrin, 2013). However, we still know little about what sampling strategies adult and child learners display when given the opportunity to control their learning input.

In the current work, we investigated whether adult and child learners seek information that aids in reducing

ambiguity about the meaning of novel words. We manipulated the ambiguity of novel word mappings by varying the degree to which object-label pairs co-occurred with one another during cross-situational word learning (Experiments 1A, 1B, 2A) or whether children could use mutual exclusivity to disambiguate the referents of novel words (Experiment 2B). The central question was whether adults and children would choose to learn more about those items that most strongly reduce referential ambiguity.

Experiments 1A & 1B

We tested whether adult learners would seek information that aided in disambiguating reference. Participants completed a cross-situational learning task in which their goal was to learn a set of object-label associations by determining the referent of each label across training. Participants were then given the opportunity to select which object-label association they would hear on the next learning trial. The central question was whether adult learners would make selections that reduce referential ambiguity about the novel object-label associations. We collected data in an online experiment (Experiment 1A) and in an in-lab experiment (Experiment 1B) that we discuss together due to their similarity in design and results.

Method

Participants. For Experiment 1A, we recruited 31 participants through Amazon Mechanical Turk. Three participants were excluded for not passing an initial auditory attention check (2) or for restarting the experiment (1). All participants were assigned to the Fully Ambiguous Condition ($n = 28$) and paid \$0.75 for completing the study.

For Experiment 1B, 62 University of Wisconsin-Madison's undergraduates participated for course credit and were randomly assigned to the Fully Ambiguous Condition ($n = 28$) or the Partially Ambiguous Condition ($n = 34$).

Stimuli. The object stimuli were 8 images of novel 'alien' creatures used in previous word learning studies (Partridge, McGovern, Yung, & Kidd, 2015). 8 novel word stimuli (*beppo*, *finna*, *guffi*, *kita*, *noopy*, *manu*, *sibu*, *tesser*) were recorded by a female native speaker of English and normalized in duration and average loudness. The association between each label and its target referent and the roles of the stimuli within a condition were randomized across participants. The stimuli were presented using a web-based experiment created using jsPsych (de Leeuw, 2014).

Design & Procedure. The experiment was split into a *Training Phase*, a *Sampling Phase*, and a *Test Phase*.

Training Phase. Participants completed 24 cross-situational learning trials (2 blocks of 12 trials), presented in random order. The goal was to learn the association between eight novel labels and their referents. On each training trial, participants were presented with two referents and two labels. The labels appeared sequentially in random order, both visually and auditorily. Consequently, the association between a particular label and its referent remained ambiguous on any single trial, but could be disambiguated by aggregating information across trials. Each object and its label occurred 6 times across the 24 training trials.




Figure 1. Overview over one block of the Training Phase for the Fully Ambiguous Condition and the Partially Ambiguous Condition

We manipulated whether the object-label associations became disambiguated across trials during training, and therefore, how uncertain participants were at the onset of the *Sampling Phase* about the specific object-label pairs. Across Experiments 1A and 1B, participants were assigned to one of two conditions: the Fully Ambiguous condition or the Partially Ambiguous condition. In the Fully Ambiguous condition, half of the object-label pairs remained ambiguous: two sets of two items were yoked together such that they were never disambiguated across training (ambiguous items; Figure 1, top left). The remaining items in the Fully Ambiguous condition were disambiguated across trials, occurring with three different object-label pairs (disambiguated items; Figure 1, right panel). In the Partially Ambiguous condition, two sets of two objects were grouped such that two specific objects co-occurred on 4 out of their 6 occurrences, but each occurred with one other object from the ambiguous object set on the remaining 2 trials (partially ambiguous items; Figure 1, bottom left). The other four objects were disambiguated as in the Fully Ambiguous condition. Note that across both conditions, participants saw each individual object and label equally frequently.

Sampling Phase. Participants next completed four sampling trials. On each trial, all 8 objects appeared in randomized locations. Participants were instructed to select which of the 8 items they wanted to hear in the next cross-situational learning trial. After participants' selection, a second object was chosen at random from the remaining objects. The two objects and their labels then appeared together in a cross-situational word learning trial with the same structure as in the training phase.

Test Phase. Participants' knowledge of the object-label associations was probed in an 8-AFC recognition test. On each test trial, all 8 objects appeared in randomized locations on the screen, along with one of the 8 labels. Participants were then asked to select the object that went with the label. No feedback was provided after a choice. Participants were tested on each label in random order, for a total of 8 recognition test trials.

Predictions. We predicted that participants would be more likely to choose to learn more about the ambiguous items than about the disambiguated items in the sampling phase. For the Partially Ambiguous condition, we expected participants to have a weaker preference for ambiguous items over the disambiguated items, since adults accurately tracking the co-occurrence evidence could successfully learn all word-referent pairs. We did not predict large differences in test accuracy between items. One possible outcome was that test accuracy would be higher for items that were disambiguated during training. However, another possibility was that ambiguous items could be learned at comparable levels to disambiguated items if participants preferentially sampled ambiguous items.

Results

Sampling choices. We report the results combining the data from Experiments 1A and 1B for convenience – however, qualitatively similar results are obtained when considering the data from Experiment 1A or Experiment 1B separately. We used the lme4 package version 1.1-18-1 in R (version 3.5.1) to fit a logistic mixed-effects model testing participants' likelihood of making an ambiguous selection against a chance level of 0.5 (Bates & Maechler, 2009; R Development Core Team, 2018), including by-subject and by-item random intercepts and a fixed effect for condition. In the Fully Ambiguous condition, participants were more likely to choose ambiguous items than disambiguated items, $b = .59, z = 3.61, p < .001$. Participants chose an object from the ambiguous set on 63.4% of trials (95% CI = [55.7%, 71.0%]) (Figure 2A). Participants in the Partially Ambiguous condition selected the partially ambiguous items on 47.8% of trials (95% CI = [39.1%, 56.5%]), thus showing no sampling preference between the two item types ($p = .64$). Participants in the Fully Ambiguous condition were more likely than participants in the Partially Ambiguous condition to select the more ambiguous object-label

associations, $b = .68$, $z = 2.64$, $p = .008$. Non-parametric analyses yielded equivalent results.

Test performance. Overall, participants showed learning of the label-object pairs, accurately selecting the correct referent in both the Fully Ambiguous condition ($M = 69.2\%$, $95\% \text{ CI} = [60.8\%, 77.5\%]$, chance = 12.5%) and in the Partially Ambiguous condition ($M = 77.6\%$, $95\% \text{ CI} = [67.1\%, 88.0\%]$) (Figure 2B). Notably, within the Fully Ambiguous condition, test accuracy was lower for the ambiguous items ($M = 61.6\%$) than for the disambiguated items ($M = 76.8\%$; logistic mixed-effects model with by-subject and by-item random intercepts and a by-subject random slope for item type, $z = 3.25$, $p = .001$).

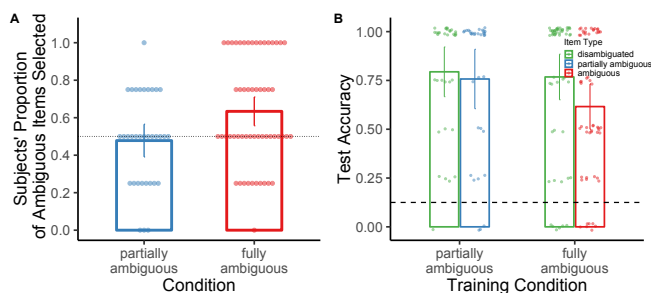


Figure 2. (A) Proportion of more ambiguous items selected in each condition and (B) test accuracy by condition and item type. Error bars in represent within-subject 95% CIs.




Figure 3. Relationship between choosing more ambiguous items and test accuracy for each condition. Error bands represent ± 1 SE.

Relationship between sampling and test performance. In the Fully Ambiguous condition, participants who chose more objects from the ambiguous set during the sampling phase accurately identified more words at test, $r(54) = .48$, $95\% \text{ CI} = [0.25, 0.66]$, $p < .001$ (Figure 3). There was no significant relationship between participants' tendency to select the partially ambiguous items and their accuracy at test ($r(32) = -.11$, $p = .50$).

Discussion

In a cross-situational learning task, adult learners chose to learn more about those object-label pairs that remained ambiguous throughout training. Adults showed this tendency when the object-label pairings were truly ambiguous based on the training evidence (Fully Ambiguous condition), but not when the object-label pairs became disambiguated at any point during training (Partially Ambiguous condition). While participants showed poorer overall learning of the (more difficult) ambiguous object-label pairs, their success at test correlated strongly with the degree to which they chose more ambiguous items during the sampling phase. This experiment provides 'proof-of-concept' evidence that adult learners will seek to reduce ambiguity about object-label associations when given the opportunity to control which items they will learn about.

Experiment 2A

Next, we asked whether children would demonstrate a similar tendency to seek new words that reduce ambiguity during cross-situational learning. As in Experiment 1A, children (4-8 years of age) first completed a cross-situational word learning task. Across training, one set of novel object-label associations could be inferred based on the object-label associations they co-occurred with, while another set of words remained ambiguous. Then, participants were given the opportunity to sample object-label associations presented in isolation, i.e. in unambiguous learning trials. The central question was whether children would prefer to select object-label associations with ambiguous evidence during training, suggesting that children sample words that reduce referential ambiguity.

Method

Participants. We recruited 38 participants ($M = 5.9$ years, range = $4.1 - 8.1$ years, 19 female) at a local children's museum. Two additional participants were excluded due to inattention during experiment.

Stimuli. The object stimuli were 8 images of novel 'alien' creatures used in previous word learning studies (Partridge et al., 2015) and 2 cartoon images of familiar animals (penguin, dog). 8 novel word stimuli (*biffer*, *deela*, *guffi*, *sibu*, *tibble*, *leemu*, *zeevo*, *pahvy*) and two familiar word stimuli (*penguin*, *dog*) were recorded by a female native speaker of English and normalized in duration and average loudness. The association between each novel label and its novel target referent, as well as the particular roles of the novel word-referent stimuli, were randomized across participants. The stimuli were presented using in a web-based experiment created in jsPsych (de Leeuw, 2014).

Design & Procedure. Children were tested in a quiet room in the children's museum on a 10.1" Samsung Galaxy Note tablet. An experimenter guided children through the experiment by giving instructions at the beginning of each

new phase. The experiment was presented as a game in which a cartoon bear named Teddy would first teach children the names of new alien friends, and then ask children to help her find her friends. The experimenter began with the following introduction:

In this game, Teddy went up to space and met a bunch of new alien friends. Teddy is going to tell you the names of aliens, and your job is to try to remember which name goes with which alien. Later, you're going to help Teddy find them.

The experiment then proceeded to a Practice Phase, followed by the main experiment consisting of three phases: the Training Phase, the Sampling Phase, and the Test Phase.

Practice Phase. Participants first completed a practice phase in which they encountered the two familiar word object stimuli and two novel object-label associations. We introduced this short practice phase to give children experience with the overall structure of the main experiment under less demanding circumstances, using a smaller set of items and mixing familiar and novel items. First, children were exposed to 4 training practice trials similar in structure to the training trials in the main experiment. On each trial, two referents appeared on the screen on either side of the Teddy character and children heard two labels, one for each object, in random order. On the first trial, children always saw the two familiar items (i.e., the penguin and the dog), followed by a second trial in which children saw two novel object-label associations (i.e., an ambiguous labeling event). On the final two training practice trials, children saw each of the familiar items occur with one of the two novel items (permitting the disambiguation of the novel object-label associations). Next, children saw two sampling practice trials, in which children had the opportunity to select which of the four items they wanted to learn about next, followed by four practice test trials, in which participants' knowledge of the items was tested in a 4-AFC recognition test. The procedure for each of these practice trial types mirrored the procedure for the Sampling Phase and the Test Phase described in more detail below.

Training Phase. Participants completed 9 cross-situational learning trials (3 blocks of 3 trials each). On each training trial, participants saw two referents appear on the screen on either side of the Teddy character and heard the labels of the two objects presented sequentially in random order. Next, the objects switched locations in a brief animation, and participants heard the same two labels presented in the same order. We introduced this trial repetition with flipped locations in order to reduce children's tendency to interpret the labeling event as moving from left to right on the screen, i.e. assuming that the first label went with the object on the left and the second label went with the object on the right.

As in the Fully Ambiguous condition of Experiment 1A, we manipulated whether the object-label associations could be disambiguated across trials during training (Figure 4). Every object-label pair occurred on three cross-situational training trials. Four of the objects occurred with three

different object label pairs (disambiguated items). The remaining two object-label associations always occurred with one another (ambiguous items), such that children never saw evidence allowing them to link the two words unambiguously with their respective referent.




Figure 4. Overview over the design of the Training and Sampling Phase in Experiment 2A

Sampling Phase. After completing the training phase, participants completed four sampling trials. On each sampling trial, all 6 referents appeared in randomized locations on the screen. Participants were instructed to select which of the 6 items they wanted to learn about next (Figure 4). When participants tapped one of the 6 referents, a brief animation moved the item to the center of the screen while the remaining items disappeared, and the referent was subsequently labeled in isolation.

Test Phase. Participants' knowledge of the object-label associations was probed in a 6-AFC recognition test. On each test trial, all 6 referents appeared in randomized locations on the screen surrounding the Teddy character. When participants tapped Teddy in the center of the screen, they heard one of the 6 labels. Participants were instructed to help Teddy by selecting the friend she was looking for. No feedback was provided after a choice. Participants were tested on each label in random order, for a total of 6 recognition test trials.

Predictions. As in Experiment 1A, our main prediction was that children would preferentially select object-label associations that remained ambiguous during the cross-situational word learning trials of the training phase.

Results

Sampling choices. Contrary to our prediction, children did not preferentially select ambiguous object-label associations during the Sampling Phase, $b = -0.01$, $z = -.11$, $p = .91$.

Participants chose an object from the ambiguous set on 32.9% of trials (95% CI = [27.1%, 38.7%]).

Test performance. Overall, participants showed significant learning of the label-object pairs, choosing the correct object to go with a label at above-chance levels (chance = 0.167), $M = 38.6\%$, 95% CI = [30.7%, 46.5%], $t(37) = 5.65$, $p < .001$. However, surprisingly, children performed more accurately on the ambiguous items ($M = 48.6\%$, 95% CI = [36.9%, 60.4%]) than on the disambiguated items ($M = 33.6\%$, 95% CI = [24.9%, 42.2%]), $b = .68$, $z = 2.23$, $p = .028$. When tested on ambiguous items, children had a strong preference to select one of the two ambiguous objects (61.8% of trials, 95% CI = [50.7%, 72.9%]) over the four disambiguated objects (chance = 0.33). When tested on disambiguated items, children tended not to choose the two ambiguous objects, selecting them on only 18.4% of trials (95% CI = [12.8%, 24.1%]).

Discussion

Unlike adult learners, children did not show a preference for selecting object-label associations for which they had experienced ambiguous evidence during training. Interestingly, children performed better at test for ambiguous object-label associations than for object-label associations that were disambiguated across training trials. There are likely two reasons why children showed higher accuracy on the ambiguous items. First, since the two ambiguous items always co-occurred with one another, the training could help learners constrain the set of possible competitors for a given ambiguous label to two objects (compared to four possible objects for the disambiguated items). Indeed, children appeared to constrain their choices to the two objects that co-occurred on ambiguous trials when tested on their respective labels and rarely chose these objects when tested on the labels that occurred with the disambiguated objects.

Second, anecdotally, we observed that many children explicitly pointed to specific objects during training while listening to each label and even repeated the respective label for each object. This behavior may indicate that some children were making an explicit hypothesis about each word mapping (Trueswell, Medina, Hafri, & Gleitman, 2013). If a child formed a specific hypothesis about the mapping between the two labels and objects on the first ambiguous trial, they would subsequently hear evidence that would appear to confirm their hypothesis: the two labels and the two objects would occur together again on the subsequent two training trials. “Hypothesis-testers” would never experience evidence disconfirming their initial hypotheses and thus have a 50% chance of responding correctly at test for these items (note that our child participants’ test accuracy was 48.6% on average). Crucially, one consequence of learners approaching the task in this manner is that the two object-label associations deemed “ambiguous” according to the experimental design may have actually appeared *less* ambiguous to children

performing the task than the putatively disambiguated items. Thus, in our next step, we adapted the task to create a learning situation in which one set of object-label associations would be more clearly ambiguous from the standpoint of the child learner.

Experiment 2B

In Experiment 2B, we sought to increase the likelihood that children would perceive some novel object-label associations as more ambiguous than others. We used mutual exclusivity to increase the ease with which children could infer word-referent pairs for one set of novel objects (Markman & Wachtel, 1988) while maintaining the ambiguity of a second set of novel word-referent pairs as in the previous experiments. By giving children the opportunity to infer the referents for novel objects occurring in mutual exclusivity trials, we aimed to make it easier for children to recognize the referential ambiguity of novel object-label associations that always co-occurred.

Method

Participants. We recruited 53 participants ($M = 5.7$ years, range = 4.1 – 7.9 years, 32 female) at a local children’s museum. One additional participant was excluded due to experimenter error.

Stimuli. The novel object and word stimuli were six images and recordings composed of a subset of the items used in Experiment 2A. In addition, 4 cartoon images of familiar animals (cow, dog, monkey, pig) along with audio recordings of their respective labels were used. All word stimuli were recorded by the same female native speaker of English and normalized in duration and average loudness.

Design & Procedure. The procedure and testing conditions were identical to Experiment 2A. The experiment followed the same structure as Experiment 2A, beginning with a *Practice Phase* and then proceeding through three phases: *Training Phase*, *Sampling Phase*, and *Test Phase*.

Training Phase. Participants completed 9 cross-situational learning trials (3 blocks of 3 trials each) with 6 object-label pairs, two familiar object-label pairs (e.g., pig and dog) and four novel object-label pairs chosen randomly from the set of novel stimuli. As in Experiment 2A, on each trial, participants saw two referents appear on the screen and heard two labels presented in random order. Two novel object-label associations always occurred with one another (ambiguous items), mirroring the ambiguity manipulation from Experiments 1A/B and 2A. The two remaining novel object-label associations were each yoked to one of the two familiar object-label pairs (i.e., one alien always occurred with the dog image, while the other always occurred with the pig image; mutual exclusivity items). We reasoned that children would successfully disambiguate reference for mutual exclusivity items (i.e., when seeing an image of a dog and a novel “alien”, on hearing the words *leemu* and *dog*, children would successfully infer that *leemu* referred to

the novel alien). This would make it more likely that the ambiguous items would be perceived by child learners as having high referential uncertainty. As in previous experiments, all novel objects and their labels occurred equally frequently across the training phase.

Sampling Phase. Participants next completed two sampling trials. On each trial, the four novel objects appeared on the screen and children were instructed to choose which object they wanted learn more about. The procedure was otherwise identical to Experiment 2A.

Test Phase. Participants' knowledge of the six words from the training phase (4 novel, 2 familiar words) was tested in a 6-AFC recognition task as in Experiment 2A.

Results

Sampling choices. Children preferentially selected ambiguous object-label associations during the Sampling Phase, $b = .58$, $z = 2.87$, $p = .004$. Participants chose an object from the ambiguous set on 64.2% of trials (95% CI = [55.0%, 73.3%]) (chance level = 0.5; Figure 5A). The likelihood of children making ambiguous selections increased with age, $b = .49$, $z = 2.42$, $p = .016$ (logistic mixed-effects models; Figure 5B).

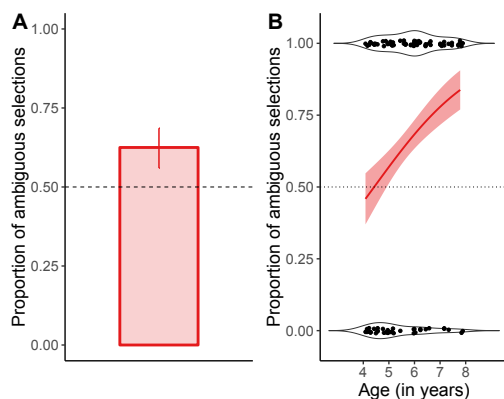


Figure 5. Proportion of ambiguous item selections in Experiment 2B overall (A) and across age (B). Error bars represent 95% CIs and error bands are ± 1 SEs based on model estimates.

Test performance. Overall, participants showed significant learning of the label-object pairs, choosing the correct object to go with a label at above-chance levels (chance selection of novel object = 0.25), $M = 59.9\%$, 95% CI = [50.4%, 69.4%], $t(52) = 7.38$, $p < .001$. Accuracy for mutual exclusivity items ($M = 64.2\%$, 95% CI = [53.2%, 75.1%]) and for the ambiguous items ($M = 55.7\%$, 95% CI = [44.0%, 67.3%]) was similar, $b = .45$, $z = 1.20$, $p = .23$.

Discussion

When given the opportunity to select which object-label pairs they wanted to learn more about, 4-8-year-olds preferentially selected object-label pairs that remained

ambiguous during training over object-label pairs that could be disambiguated through mutual exclusivity. These findings demonstrate that – at least in some ambiguous word learning situations – children prefer to select learning events that aid in reducing referential uncertainty. The tendency to make ambiguity-reducing selections began to emerge around 5 years of age in our sample.

General Discussion

When learning the referents of novel labels in ambiguous contexts, adult learners chose to learn more about object-label associations that remained more ambiguous at the end of training. These choices appear to help learners improve performance: participants' learning performance at test was higher if they had selected more ambiguous items during the Sampling Phase. It is interesting to note the modest magnitude of adults' preference on the task: ambiguous items were selected on slightly less than two-thirds of adults' sampling trials. This may be partly related to the design of the sampling phase, which allowed for a number of potentially successful sampling strategies (e.g., selecting a known word on each sampling trial in order to hear that known word in combination with other words). However, another intriguing possibility is that there are individual differences in how adults organize their learning, and that these differences may lead to distinct learning outcomes.

We find mixed evidence that children spontaneously sample object-label associations that reduce ambiguity. When presented with a similar task, 4-8-year-olds did not choose to learn about object-label associations that remained ambiguous during training. However, we think this result may be partially explained by the fact that word-referent pairs occurring in ambiguous contexts also never provided children with evidence that could disconfirm an existing hypothesis about word reference. In a simplified design that highlighted the ambiguous nature of the trials in which two referents always occurred together, older children in our sample chose to learn about items that reduced uncertainty about the words' referents.

Children have substantial control over their “curriculum” as they learn new words in the world (Smith, Jayaraman, Clerkin, & Yu, 2018), with potentially immense consequences for the difficulty of the learning problem they face (Hidaka et al., 2017). The results from Experiment 2B are consistent with results from domains such as causal learning that suggest that children are motivated to explore novel objects when presented with confounded evidence (e.g., Schulz & Bonawitz, 2007). However, the limits on the extent to which children spontaneously make ambiguity-reducing selections also raise important questions about what sampling strategies children employ when in control of what they learn next. A key question for future research will be investigating how children's sampling strategies and the structure of their environment interact to support learning.

Acknowledgements

This research was supported by NSF-GRFP DGE-1256259 awarded to MZ and grants from the NICHD awarded to JS (R37HD037466) and the Waisman Center (U54 HD090256). We would like to thank Joseph Austerweil for advice on the project and Tess Gapinski, Louisa Forrest, and Andrew Kressin for aiding in data collection.

References

- Bates, D., & Maechler, M. (2009). *lme4: Linear mixed-effects models using Eigen and Eigen*.
- de Leeuw, J. R. (2014). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 1–12.
- Hidaka, S., Torii, T., & Kachergis, G. (2017). Quantifying the impact of active choice in word learning. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 519–525). London, UK: Cognitive Science Society.
- Kachergis, G., Yu, C., & Shiffrin, R. M. (2013). Actively learning object names across ambiguous situations. *Topics in Cognitive Science*, 5, 200–213.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88, 449–460.
- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, 20, 121–157.
- Partridge, E., McGovern, M. G., Yung, A., & Kidd, C. (2015). Young children's self-directed information gathering on touchscreens. In R. Dale, C. Jennings, P. Maglio, T. Matlock, D. Noelle, A. Warlaumont, & J. Yoshimi (Eds.), *Proceedings of the Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- R Development Core Team. (2018). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43, 1045–1050.
- Smith, L. B., Jayaraman, S., Clerkin, E., & Yu, C. (2018). The developing infant creates a curriculum for statistical learning. *Trends in Cognitive Sciences*, 22, 325–336.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94.
- Trueswell, J. C., Medina, T. N., Hafri, A., & Gleitman, L. R. (2013). Propose but verify: fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66, 126–156.
- Yu, C., & Smith, L. B. (2007). Rapid word learning under

uncertainty via cross-situational statistics.
Psychological Science, 18, 414–420.