



Brief article

Infants rapidly learn word-referent mappings via cross-situational statistics

Linda Smith *, Chen Yu

*Department of Psychological and Brain Sciences, and Cognitive Science Program,
Indiana University, Bloomington, IN 47405, USA*

Received 12 February 2007; revised 26 June 2007; accepted 27 June 2007

Abstract

First word learning should be difficult because any pairing of a word and scene presents the learner with an infinite number of possible referents. Accordingly, theorists of children's rapid word learning have sought constraints on word-referent mappings. These constraints are thought to work by enabling learners to resolve the ambiguity inherent in any labeled scene to determine the speaker's intended referent at that moment. The present study shows that 12- and 14-month-old infants can resolve the uncertainty problem in another way, not by unambiguously deciding the referent in a single word-scene pairing, but by rapidly evaluating the statistical evidence across many individually ambiguous words and scenes.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Language acquisition; Word learning; Statistical learning; Development; Infant learning

1. Introduction

The pairing of a word and a scene is not enough to determine the meaning of the word. To illustrate this point, Quine (1960) famously imagined a stranger who hears a native say “gavagai” and points to a scene. To what does “gavagai” refer – a rabbit,

* Corresponding author.

E-mail addresses: smith4@indiana.edu (L. Smith), chenyu@indiana.edu (C. Yu).

the grass, a tree, the rabbit's ears, or perhaps the beauty of the whole? Even if one assumes a perceptual system that segments the scene into separate objects and an attentional system biased towards objects, the intended referent is indeterminate from this one experience. Infants are like strangers who do not know the native language, yet they solve this indeterminacy problem. This paradox – the uncertainty of the referent in word-scene associations and the fact that infants learn object names nonetheless – is a core theoretical problem in the study of early word learning. For the past 30 years most research on children's word learning has concentrated on how the learner resolves the ambiguity *at the moment the novel word is first encountered*. Experimental studies leave no doubt that by the time they are 2 years old children do this at least for object names. That literature points to attentional (Smith, 2000), social (Baldwin, 1993; Tomasello, 2000), linguistic (Gleitman, 1990) and representational (Markman, 1990) constraints as crucial to children's ability to resolve referential ambiguity and fastmap a word to its intended referent.

There are two reasons to suspect that this one-encounter solution to referential uncertainty is not the only (or even the most important) mechanism of early word learning. First, not all opportunities for word learning are as uncluttered as the experimental settings in which fast-mapping has been demonstrated. In everyday contexts, there are typically many words, many potential referents, limited cues as to which words go with which referents, and rapid attentional shifts among the many entities in the scene. It is possible that young learners just ignore the information in such highly ambiguous learning contexts and wait for contexts in which the referents of heard words are more certain (Brent & Siskind, 2001). However, a more optimal learner might be expected to make use of all the available data.

Second, the evidence indicates that 9-, 10-, and certainly 12-month-old infants are accumulating considerable receptive lexical knowledge (Fenson, Dale, Reznick, & Bates, 1994; Swingley & Aslin, 2000). Yet many studies find that children even as old as 18 months have difficulty in making the right inferences about the intended referents of novel words (e.g., Katz, Baker, & Macnamara, 1974; Hirsh-Pasek, Golinkoff, & Hollich, 1999; Moore, Angelopoulos, & Bennett, 1999; Pruden, Hirsh-Pasek, Golinkoff, & Hennon, 2006). There are studies showing that infants as young as 13 or 14 months (Woodward, Markman, & Fitzsimmons, 1994; Woodward & Hoyne, 1999; Schafer & Plunkett, 1998; but perhaps not younger, Werker, Cohen, Lloyd, Casasola, & Stager, 1998) can link a name to an object given repeated unambiguous pairings in a single session. Overall, however, these effects are fragile with small experimental variations often leading to no learning (see especially, Woodward & Hoyne, 1999; Werker et al., 1998; also Oviatt, 1980, 1982 & Bloom, 2000 for a discussion). This raises the possibility that there might be some other way that young children learn word-referent mappings.

The experiment reported here shows for the first time that infants *rapidly* learn *multiple* word-referent pairs by accruing statistical evidence across multiple and individually *ambiguous* word-scene pairings. The indeterminacy problem is solved not in a single trial but across trials, not for a single word and its referent but for a data set of many words and referents. This learning is shown to be sufficiently rapid and robust that it could play a significant role in early lexical learning.

Fig. 1 illustrates how cross-trial statistics might work. The learner hears the unknown words “bat” and “ball” in the context of seeing a BAT and BALL. Without other information, the learner cannot know whether the word form “ball” refers to one or the other visual object. However, if subsequently, while viewing a scene with the potential referents of a BALL and a DOG, the learner hears the words “ball” and “dog” and if the learner can combine the co-occurrence frequencies from the two streams of data across trials, the learner could correctly map “ball” to BALL. This example represents the simplest case of cross-situational statistical learning – two words, two objects, two adjacently informative trials.

Several formal simulations of word-referent learning suggest the plausibility of cross-situational word learning in much more complex situations with many words, many possible referents, highly ambiguous individual learning trials, and the statistical resolution of the ambiguities only through the accumulation and evaluation of information over many word-referent pairings and many trials (Siskind, 1996; Yu, Ballard, & Aslin, 2005). Consider the more complex case in Table 1. On trial 1, a learner could mistakenly link word A to referent b. On trial 4, the mistake could be corrected, *if* the system registers that word A occurred on trial 4 *without* possible referent b, *if* the cognitive system remembers the prior word-referent pairing, *if* it registers both co-occurrences and non co-occurrences, and *if* it calculates the right statistics. Can babies do this?




Fig. 1. Associations among words and referents across two individually ambiguous scenes. If a young learner calculates co-occurrences frequencies *across* these two trials, s/he can find the proper mapping of “Ball” to BALL.

Table 1
An example of cross-situational learning

Trial	Words	Potential referents in scene
1	AB	ba
2	CD	dc
3	EF	ef
4	GA	ga

There is evidence in such phenomena as the mutual-exclusivity effect and contrast that 2- to 3-year-old children combine information across two adjacent naming events, using, for example, knowledge of the just-heard name of one thing to infer the object to which a subsequent name must apply (Akhtar, 2002; Akhtar & Montague, 1999; Markman, 1990; Namy & Gentner, 2002.) However, there is no evidence as to whether young learners can combine and evaluate information from highly ambiguous contexts over many trials. Until recently, there was no evidence as to whether even adult learners were capable of this, although Yu and Smith (2007; Yu, Smith, Klein, & Shiffrin, 2007) have now shown that this form of learning is rapid and robust in adults even in situations of high uncertainty.

In the following experiment, 12- and 14-month-old infants were taught 6 word-referent pairs via a series of individually ambiguous trials. On each trial, two word forms and two potential referents were presented with no information about which word went with which referent. Although word-referent pairings were ambiguous within individual trials, they were certain across trials. For example, for a particular infant, whenever the form *tobi* occurred its assigned referent always occurred. After training, infants were presented with a single word and two potential referents, the cross-trial correct referent and a foil. Past research (e.g., Golinkoff, Hirsh-Pasek, Cauley, & Gordon, 1997; Swingley & Aslin, 2000) shows that within this kind of preferential looking task, infants look longer at the labeled test object. Thus if infants have calculated the statistics appropriately, despite the uncertainty on individual learning trials, they should look longer at the correct referent of the word form.

2. Method

2.1. Participants

The participants, drawn from a working and middle-class population of a midwestern college town, were 28 12-month-old infants (range – 11 mo 17 days to 13 mo 0 days; mean – 12 mo 7 days; 13 males, 15 females) and 27 14-month-old infants (range –14 mo 2 days to 15 mo 14 days; mean –14 mo 12 days; 14 males, 13 females). Two additional children began but did not finish the experiment.

2.2. Stimuli

The 6 “words” – *bosa*, *gasser*, *manu*, *colat*, *kaki* and *regli* – followed the phonotactic probabilities of English and were recorded by a female speaker in isolation and were presented to infants over loudspeakers. The 6 “objects” were drawings of novel shapes, shown in Fig. 2; each was a unique bright color. On each trial, two objects (12 by 14 in. in projected size and separated on the screen by 30 in.) were simultaneously presented on a 47 by 60 in. white screen.




Fig. 2. The six stimulus shapes.

2.3. Procedure

Infants sat (on their mother's lap) 3.5 feet in front of screen with the mother's chair set at the center of the screen. Infants' direction of eye gaze was recorded from a camera centered at the base of the screen and pointed directly at the child's eyes. Parents were instructed to keep their own eyes shut through the entire procedure so as to not to influence their infant's behaviors. A camera directed on the parent through out the procedure confirmed their adherence.

There were 30 training slides. Each presented two objects on the screen for 4 s; the onset of the slide was followed 500 ms later by the two words – each said once with a 500 ms pause between. Across trials, the temporal order of the words and spatial order of the objects were varied such that there was no relation between temporal order of the words and the spatial position of the referents. Each correct word-object pair occurred 10 times. The two words and two objects appearing together on a slide (and creating the same within trial ambiguities and possible spurious correlations) were randomly determined such that each object and each word co-occurred with every other word and every other object at least once across the 30 training trials. The first four training trials each began with the centered presentation of a Sesame Street character (3 s) to orient attention to the screen. After these first four trials, this attention grabbing slide was interspersed every 2–4 trials to maintain attention. The entire training – an effort to teach six word-referent pairs – lasted less than 4 min (30 training slides and 19 interspersed Sesame Street character slides).

There were 12 test trials, each 8 seconds. This duration was chosen from pilot studies to optimize the number of participants able to complete all 12 test comparisons (2 per target word). Each test trial presented one word, repeated 4 times with 2 objects – the target and a distracter – in view. The distracter was drawn from the training set. Each of the 6 words was tested twice. The distracter for each trial

was randomly determined such that each object occurred twice as a distracter over the 12 test trials.

There were 2 unique sets of training slides with different orderings of objects, different mappings of words to the objects, and different combinations of word-referent pairs on the slides. For each set, the left–right locations of objects on the slides and the order with which the names were presented were randomly generated with the constraint that the object on the left was the target referent for the first presented word on half the trials and the target referent for the word presented second on the other half. There were also two unique test orders with unique randomly generated pairings of target and distracter, with the target appearing on the left on half the slides and on the right on the other half. Half the infants at each age level were randomly assigned to each slide set.

Two coders naïve to condition and trial type coded direction of eye gaze from the video recorded from the camera directed at the infant's eyes. They coded, frame-by-frame, all frames from the start to the end (indicated by light on the video) of each training and test trial. The coder's task for each frame was to categorize the direction of look as right, left or away from the screen (hands, ceiling, mother's face, floor, etc.). For reliability, the two coders each coded the same random sample of 25% of the frames. Agreement on these frames was 90.8%.

3. Results

3.1. Training trials

Infants were highly attentive to the training slides, looking (sum of right and left looks) at each 4 s slide on average 3.27 s (12 month olds) and 3.04 s (14 month olds). On average, infants looked at the left and right sides of each training slide for equal durations ($t < 1.00$ for both 12- and 14-month olds). On 87% of all training slides, the infants looked at both sides (both objects) for at least 1 s.

3.2. Test trials

On average, infants looked at each 8 s test slides for a total of 5.6 s for 12 month olds and 6.1 s for 14 month olds. To examine whether infants *preferentially* looked in the direction of the target object, the object that *across trials* was associated with the auditorally presented label, each infant's looking time to target and distracter on each test trial was submitted to a 2(Age) by 2(Target/Distracter) \times 6 (Word) \times 2(Block – first or second test of each target word) analysis of variance for a mixed design. The analysis revealed a highly reliable main effect of looking time to Target/Distracter, $F(1, 54) = 35.32$, $p < .001$ (partial eta squared = .37). As shown in Fig. 3, 12- and 14-month-old children looked reliably longer to the Target than to the Distracter. The analysis also revealed a reliable interaction between Word and Target/Distracter, $F(5, 54) = 3.85$, $p < .05$ (partial eta squared = .19). This result, that infants showed a greater difference in looking time to the target than distracter for




Fig. 3. Mean looking time to target and distracter per 8 s test trial (and standard error of the mean) for younger and older infants.

some words than for others suggests that some word-picture correspondences were learned better than others. Finally, the analysis revealed an interaction between Age and Target/Distracter that approached significance, $F(5, 54) = 3.13$, $p < .08$ (partial eta squared = .04). The older group of children, as can be seen in Fig. 3, showed a bigger preference for the target than did the younger children, although the difference in looking times to target and distracter is individually reliable for both age groups (Tukey's hsd, $p < .05$). No other main effects or interactions approached significance.

Post hoc analyses (Tukey's hsd, $p < .05$) conducted on the difference in looking time to target and distracter for the 6 individual words indicated reliably greater looking time to target than distracter for 4 of the 6 words for the 12-month-old group and for 4 of the 6 words for the 14-month-old age group. (Three of the individual words were the same at the two age levels, one was different; at neither age level were there reliable differences in the wrong direction for the remaining two words.) Since half the children at each age level had different word-object pairings as well as different training orders, and since analyses for effects of slide set yielded no effects or interactions that approached significance, the source of these differences is not obvious. However, the fact that looking times for 4 of the 6 words (67% of the training set) show reliable preferences for the target does indicate that infants can figure out *multiple* word-referent mappings from a *system* of experienced associations. Finally, the group patterns appear to characterize the performance of individual infants in that 46 of the 55 participants infants looked, on average, at the targets more than distracters.

In sum, these results tell us that cross-situational statistical learning is in the repertoire of young word learners. Despite the ambiguity of word-referent mappings on any individual training trial, infants clearly accumulate information across trials and use that information to determine the underlying mappings. In less than four min-

utes, with six different word forms and six different objects, infants learned enough to systematically look longer at the objects more strongly associated with the forms than those more weakly associated.

4. General discussion

Parents, on average, direct between 300 and 400 utterances an hour to their children (Hart & Risley, 1995). Even with social, linguistic and conceptual constraints in play, so many words in so little time seems likely to generate considerable ambiguity about intended referents. These ambiguities are most likely greater than those in this experiment. Nonetheless, the mechanisms responsible for the present results may be relevant to *making use* of the complexity in natural learning environments in that these mechanisms can keep track of multiple word-referent co-occurrences, evaluate the regularities in the data set as a whole, and determine the underlying mappings. Such mechanisms could even benefit from increased complexity in the data set. Consistent with this idea, Yu and Smith (2007; Yu et al., 2007), using a task much like the infant task used here, showed that adults actually learned more word-referent pairs when the set contained 18 words and referents than when it contained only 9. This is because more words and referents mean better evidence against spurious correlations. Although much remains to be discovered about the relevant mechanisms, they clearly should help children learn from the regularities that accrue *across* the many ambiguous word-scene pairings that occur in everyday communication.

The present findings are thus reminiscent of evidence showing that infants' use sequential probabilities to discover segmental units in speech (Gomez & Gerken, 1999; Kirkham, Slemmer, & Johnson, 2002; Saffran, Aslin, & Newport, 1996). The statistical regularities to which infants must attend to learn word-referent pairings are different from those underlying the segmentation of a sequential stream in that word-referent pairings require computing co-occurrence frequencies across two streams of events (words and referents) simultaneously for many words and referents. Nonetheless, the present findings, like the earlier ones showing statistical learning of sequential probabilities, suggest that solutions to fundamental problems in learning language may be found by studying the statistical patterns in the learning environment and the statistical learning mechanisms in the learner (Newport & Aslin, 2004; Saffran et al., 1996).

There are several possible learning mechanisms that could accomplish the cross-situational learning of word-referent mappings. One is the formulation and evaluation of hypotheses (e.g., Tenenbaum & Xu, 2000). Building on the "ball/bat" example in Fig. 1, the learner could, for example, wrongly hypothesize on the initial trial that "ball" refers to BAT but correct that hypothesis on trial 2, which presents disconfirming evidence. Given enough data across individually ambiguous trials, the co-occurrence frequencies would support the "right" hypotheses for the language over others. The outcome of this learning would seem to be a list of confirmed hypotheses, each specifying a word and its referent.

However, statistical learning need not be the result of highly specialized statistical learning mechanisms (e.g., Perruchet & Vinter, 1998). The learner could solve this learning task via simple (or not so simple, see Kruschke, 2001; Kruschke, Kappenman, & Hetrick, 2005; Yu & Smith, 2007; Yu et al., 2007) associative learning mechanisms. Across trials, the learner could accumulate associations between words and potential referents by strengthening and weakening associative links with each co-occurrence or non-co-occurrence (see, Plunkett, 1997). Building on the example in Fig. 1, the learner could equally associate “ball” with BALL and BAT but after the experience of “ball” in the context of BALL and DOG, the association between “ball” and BALL would be stronger than that between “ball” and BAT. Over enough trials, these association strengths would converge on the real world statistics. The outcome of this learning, unlike the hypothesis testing account, might not be knowledge that an individual word refers to one thing, but may only be stronger correct associations than spurious ones.

The present results cannot distinguish these possibilities. Perhaps early associative learning lays the ground work in infancy for more rapid (and perhaps more hypothesis-testing like) processes in later word learning. A recent simulation study by Yu (in press) makes this point. That study examined a probabilistic associative learning mechanism that learns a *system* of associations (Yoshida & Smith, 2003). In such a system, a single word-referent pairing is correlated with all the other pairings that share the same word and all the other pairings that share the same referent, which are in turn correlated with more word-referent pairs, yielding a system of correlations. Such large systems of associations create system-wide accelerations of word-referent learning even when the individually contributing associations are partially learned. This is because a system of even partially learned associations yields latent structure that can be used to guide subsequent learning. In this way, the lexical knowledge of 12- to 14-month olds, even if based on associations, could contribute significantly to the later, more rapid, and more seemingly sophisticated one-trial word learning of older children. Regardless of which kind of mechanism proves right, the present results suggest that the relevant mechanisms may be best conceived as not being about the learning of individual words and referents – not about the testing of individual hypotheses or the learning of single associations – but rather as being about processes that evaluate the regularities in data sets of many words and referents (a point originally made by Billman & Knutson, 1996). The human learning environment is data rich. If human learners possess the right learning mechanisms, they may mine this complexity and in so doing solve the problem of referential uncertainty.

Acknowledgments

We thank Ling Jiang, Lisa Cantrell and Char Wozniak for collection of the data and Rich Shiffrin for insightful discussions. This research was supported by National Science Foundation Grant BCS0544995.

References

- Akhtar, N. (2002). Relevance and early word learning. *Journal of Child Language*, 29(3), 677–686.
- Akhtar, N., & Montague, L. (1999). Early lexical acquisition: The role of cross-situational learning. *First Language*, 19(57, Pt. 3), 347–358.
- Baldwin, D. A. (1993). Early referential understanding: Infants' ability to recognize referential acts for what they are. *Developmental Psychology*, 29(5), 832–843.
- Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(2), 458–475.
- Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA, USA: The MIT Press.
- Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, 81(2), 33–44.
- Fenson, L., Dale, P. S., Reznick, J. S., & Bates, E. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, 59(5), v–173.
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1, 1–55.
- Golinkoff, R. M., Hirsh-Pasek, K., Cauley, K. M., & Gordon, L. (1997). The eyes have it: Lexical and syntactic comprehension in a new paradigm. *Journal of Child Language*, 14(1), 23–45.
- Gomez, R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70(2), 109–135.
- Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore, MD: Brookes Publishing.
- Hirsh-Pasek, K., Golinkoff, R. M., & Hollich, G. (1999). Trends and transitions in language development: Looking for the missing piece. *Developmental Neuropsychology*, 16(2), 139–162.
- Katz, N., Baker, E., & Macnamara, J. (1974). What's in a name? A study of how children learn common and proper names. *Child Development*, 45(2), 469–473.
- 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.
- Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, 45(6), 812–863.
- Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(5), 830–845.
- Markman, E. M. (1990). Constraints children place on word learning. *Cognitive Science*, 14, 57–77.
- Moore, C., Angelopoulos, M., & Bennett, P. (1999). Word learning in the context of referential and salience cues. *Developmental Psychology*, 35(1), 60–68.
- Namy, L. L., & Gentner, D. (2002). Making a silk purse out of two sow's ears: Young children's use of comparison in category learning. *Journal of Experimental Psychology: General*, 131(1), 5–15.
- Newport, E. L., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48, 127–162.
- Oviatt, S. L. (1980). The emerging ability to comprehend language: An experimental approach. *Child Development*, 51(1), 97–106.
- Oviatt, S. L. (1982). Inferring what words mean: Early development in infants' comprehension of common object names. *Child Development*, 53(1), 274–277.
- Perruchet, P., & Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language*, 39, 246–263.
- Plunkett, K. (1997). Theories of early word learning. *Trends in Cognitive Sciences*, 1, 146–153.
- Pruden, S. M., Hirsh-Pasek, K., Golinkoff, R. M., & Hennon, E. A. (2006). The birth of words: Ten-month-olds learn words through perceptual salience. *Child Development*, 77(2), 266–280.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928.

- Schafer, G., & Plunkett, K. (1998). Rapid word learning by fifteen-month-olds under tightly controlled conditions. *Child Development*, 69(2), 309–320.
- Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, 61(1-2), 1–38.
- Smith, L. B. (2000). How to learn words: An associative crane. In R. Golinkoff & K. Hirsh-Pasek (Eds.), *Breaking the word learning barrier* (pp. 51–80). Oxford: Oxford University Press.
- Swingle, D., & Aslin, R. N. (2000). Spoken word recognition and lexical representation in very young children. *Cognition*, 76(2), 147–166.
- Tenenbaum, J. B., & Xu, F. (2000). Word learning as Bayesian inference. In L. Gleitman & A. Joshi (Eds.), *Proceedings of the 22nd annual conference of the cognitive science society*. Hillsdale, NJ: Erlbaum.
- Tomasello, M. (2000). Perceiving intentions and learning words in the second year of life. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 111–128). Cambridge University.
- Werker, J. F., Cohen, L. B., Lloyd, V. L., Casasola, M., & Stager, C. L. (1998). Acquisition of word-object associations by 14-month-old infants. *Developmental Psychology*, 34(6), 1289–1309.
- Woodward, A. L., & Hoyne, K. L. (1999). Infants' learning about words and sounds in relation to objects. *Child Development*, 70(1), 65–77.
- Woodward, A. L., Markman, E. M., & Fitzsimmons, C. M. (1994). Rapid word learning in 13- and 18-month-olds. *Developmental Psychology*, 30(4), 553–566.
- Yoshida, H., & Smith, L. B. (2003). Shifting ontological boundaries: How Japanese- and English-speaking children generalize names for animals and artifacts. *Developmental Science*, 6(1), 1–17.
- Yu, C. (in press). A statistical associative account of vocabulary growth in early word learning. *Language Learning and Development*.
- Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18(5), 414–420.
- Yu, C., Ballard, D. H., & Aslin, R. N. (2005). The role of embodied intention in early lexical acquisition. *Cognitive Science*, 29, 961–1005.
- Yu, C., Smith, L. B., Klein, K., & Shiffrin, R. M. (2007). Hypothesis testing and associative learning in cross-situational word learning: Are they one and the same? In *Proceedings of the 29th cognitive science society conference*, August, 2007.