

How Misleading Cues Influence Referential Uncertainty in Statistical Cross-situational Learning

Yayun Zhang and Chen Yu

1. Introduction

In early language learning, children have to learn what object a word refers to. Given that there are seemingly infinite number of possible word-object mappings when a word is heard for the first time, the question is how young learners can solve this puzzle. This is the problem of referential uncertainty famously framed by Quine (1960).

Many recent studies claim that when children learn to find what object a new word refers to, they may use very different learning strategies, such as the whole-object assumption (Markman, 1991) and the mutual-exclusivity assumption (Mather & Plunkett, 2012). Children may also benefit from a number of different sources of information, including social and pragmatic knowledge (Baldwin, 1991; Bloom, 2000; Tomasello, 2003) and knowledge of syntactic structure (Fisher, Klinger, & Song, 2006; Yuan & Fisher, 2009). In addition to these useful strategies that give children a better idea of what to guess, recent evidence argues that both adult and infant learners can correctly identify a word's referent by using cross-situational information (e.g., Yu & Smith, 2007; Smith & Yu, 2008). For example, Yu and Smith (2007) presented adults a series of trials in which they saw four novel objects and heard four novel words on every training trial. Although each novel object has a fixed label associated with it, participants have no information indicating which object mapped on to which label on any given trial. They found that even though participants could not find the correct word-object mappings initially, they were able to gradually learn which novel label consistently co-occurs with which object across multiple training trials. This finding suggests that human learners are capable of keeping track of multiple possible word-object pairings simultaneously, and they use aggregated knowledge to inform later decisions (Yu & Smith, 2007).

However, other studies suggest that it is impossible for human learners to keep track of multiple co-occurrences of object-label mappings in one naming situation as the real-world is too noisy that infinite referents can potentially be treated as the named object (Medina, Snedeker, Trueswell, & Gleitman, 2011). Their argument is that learners only form one hypothesis or conjecture about an object-label pairing in one learning moment. They only retain this hypothesis and

* Yayun Zhang and Chen Yu, Indiana University-Bloomington. Corresponding author: yayzhang@indiana.edu.

discard all possible alternatives. If this hypothesized referent is present in the next learning instance, it will be confirmed and treated as learned knowledge. However, if the hypothesized referent is absent, then the hypothesis is discarded. The learners generate a new hypothesis based on information available from the current referential scene. This hypothesis testing model suggests that children begin word learning by making an initial fast mapping between a new word and its likely meaning (Carey & Bartlett, 1978). Although they also modify the guesses as more input comes in, they do not modify their guesses based on past experience or knowledge as shown in the cross-situational learning model.

There are different theories of early word learning. Empirical studies to examine different accounts using similar method have yield mixed results. In one recent study, researchers followed young children around their home and recorded videos of natural interactions between parents and their children (Medina, Snedeker, Trueswell, & Gleitman, 2011). Using the “Human Simulation Paradigm (HSP)” (Gillette, Gleitman, Gleitman & Lederer, 1999), they cut the original videos into 40-second vignettes of parents uttering labels to their infants. They muted the sound of the video and inserted a beep at the onset of the label when parent named the object. Adult participants were asked to watch the video and guess which object parents intended to label at the moment of the beep. They found that 90% of the natural learning instances are quite uninformative, meaning that participants’ identification accuracy is low. Only a small percentage (7%) of naming instances are considered highly informative. They then explored how providing several contexts for new words influences learning or whether cross-situational information helps learning. They showed participants 5 learning instances, in which 4 are low informative (LI) and 1 is high informative (HI). They controlled the serial ordering of high and low informative vignettes by placing the HI trial at different places: at the beginning, in the middle, at last, and no HI. Participants made one guess after viewing each vignette and also provided a final conjecture at the end of the experiment. They found that participants were not able to aggregate information and learn the correct word-referent mapping across trials if HI trials were not presented first. These results suggest that learners only memorized their previous guess. If that guess was disconfirmed, participants had little to no memory of alternative pairings that they could return to. In their case, guesses on the first trial determines later learning accuracy, which is in line with the “fast mapping” or hypothesis testing model (Medina et al., 2011). If only high informative cases help learning and most of the naturalistic learning instances children experience are low informative, is it the case that most of these ambiguous moments do not contribute to real-world word learning?

To answer this question, Yurovsky, Smith and Yu (2013) did another study using the same HSP paradigm. In addition to recording parent-child natural interaction from a third-person perspective, they also attached a small camera to toddlers’ foreheads in order to get the child’s first-person view when naming occurs. Similar to the previous study, they showed adult participants a series of 5-second naming events from either the first- or the third-person view, and participants were instructed to guess the referent the mother labeled. The result

shows that guess accuracy varied considerably across vignettes in the corpus. Specifically, target objects were correctly identified 60% of the time from both views, and about half of the naming instances were highly ambiguous or highly unambiguous. Knowing the bimodal distribution of ambiguity in natural naming events, Yurovsky et al. (2013) then tested whether learners can extract useful information from LI naming instances. They showed participants 5 LI vignettes in random orders and participants were told to make a guess after each video. They found significant learning across instances from the first-person view but not the third-person view. The hypothesis testing model suggests that learning only occurs after initial successful guesses. However, results generated from first-person view data show that participants made a significant progress even after an incorrect initial guess. This finding supports the cross-situational learning view that word learning not only emerges from highly informative learning events, but also from aggregating information from less informative cases. In a related study using the same set of stimuli, Zhang, Yurovsky and Yu (2015) presented adult participants a mixture of high and low informative first-person videos and assessed learning performance on a trial-by-trial basis. They found that even when participants failed to find the correct target for the previous trial, their current trial accuracy was still significantly above baseline, indicating that learners do use their previous knowledge to guide current decision making and word learning is more likely to be a slow and continuous process rather than a fast mapping one.

Although both aforementioned studies show that word learning moments in which an object is labeled vary in informativeness (Medina et al., 2011; Yurovsky et al., 2013), the question of whether learners aggregate information across learning instances is still difficult to reconcile. There is no doubt that HI moments in which parents name the only dominant object in child's view helps learning, but not all parent-generated labels refer to the dominated objects in children's view (Yu & Smith, 2012). Even though LI moments are generally considered as having high referential uncertainty, they may nonetheless contain useful information if learners can accumulate such information across multiple learning situations.

The present paper focuses on the question of whether low-informative instances contribute to learning. To address this question, we propose that the general concept of low-informativeness can be decomposed into two factors: uncertainty and correctness. Previous studies focus on the first factor – uncertainty, that is, how certain a learner is about the referent of a word at a learning moment, but our recent studies (Yurovsky et al., 2013, Zhang et al., 2015) suggested another important factor that may have been overlooked previously: correctness – whether the referent is correct. If the referent is incorrect, even with certainty, that learning moment would be low-informative for learners. Thus, with the two factors, low-informative learning instances can be categorized into two cases: 1) low certainty and low correctness, in which learners are uncertain about the correct referent among several possible candidates; 2) high certainty but low correctness, in which learners guess the wrong referent with high confidence.

The two primary aims of the present study were to examine 1) whether these two different kinds of low-informative learning moments occur in natural naming

instances in parent-child interactions. Different from previous studies, we not only assessed participants' guess accuracy but also their level of confidence as we expect that not all LI learning moments are equally ambiguous. We hypothesize that learners will be more confident regarding their guesses in some cases compare to others even though their accuracies are all considered low; and 2) whether those two kinds of learning instances contribute to statistical learning and if so whether they contribute in different ways. One hypothesis is that learning moments with high uncertainty contain partial information that could be useful for statistical learning. Taken this view, those learning moments with high certainty but low correctness contain misleading information – pointing to the wrong referent with certainty, which may confuse learners and hurt learning.

2. Experiment 1: Individual naming moments

Because previous studies do not distinguish LI trials with high uncertainty from LI trials with low uncertainty, they found mixed results supporting different theoretical ideas behind word learning. Experiment 1 thus provided a measure to test whether LI trials contain different information.

2.1. Participants

Forty-seven adults (28 females and 19 males, $M_{\text{age}} = 24.52$, $SD_{\text{age}} = 7.01$) at Indiana University participated for either course credit or payment. None had done similar experiments before.

2.2. Stimuli

A total of 96 naming-moment vignettes (36 HI, 60 LI) were selected from a video corpus collected by Yurovsky et al (2013) for their original study. The videos included play sessions from 8 mother-child dyads. Parent-child dyads were asked to play naturally with 25 toys for about 10 minutes while their interactions were recorded by a tripod-mounted camera and a head-mounted camera in order to get both the observer's view and the child's view at each naming moment. The current study only used videos from the child's view.

Using HSP, the original sound of each naming instance used in the study was muted and the toy name was replaced by a beep at the onset of the label. Most vignettes were 5 seconds long, with the name's onset occurring at exactly the third second (Figure 1). Two more seconds were added to the vignettes if mothers said the toy name again within 2 seconds after the first naming instance.

2.3. Procedure

Participants were told to watch vignettes of mothers playing with their children and to guess which object the mothers were naming by choosing some likely answers. After watching each vignette, participants would see all 25

possible objects on the screen, and they needed to make 3 explicit choices from the array with the first choice being the most likely referent, the second choice being the next possible referent, and the third choice being the third likely referent. They were allowed to choose the same object more than once if they were absolutely certain that that object was the correct target. We expected that for HI trials, participants would consistently pick the correct target. For LI trials with high certainty, we expected to see a similar choice pattern found in HI trials, such that participants would consistently pick one target but it happened to be the wrong one. For LI trial with low certainty, we expected to see more diverse choice patterns because participants might treat multiple objects as potential targets and they were uncertain which one might be correct. We allowed any combination of choices, meaning that participants could decide to switch to another choice or stay with the current choice at any given trial. No feedback was given after each test and participants were not allowed to go back and change their answers from previous trials. Participants had no knowledge of the ambiguity levels of these vignettes.

In addition, four sample trials with varying difficulties were given before the real trials to ensure participants' clear understanding of the instruction. HI and LI trials were tested separately. All participated completed the LI session first and then HI session. Trial order was randomized within each session.

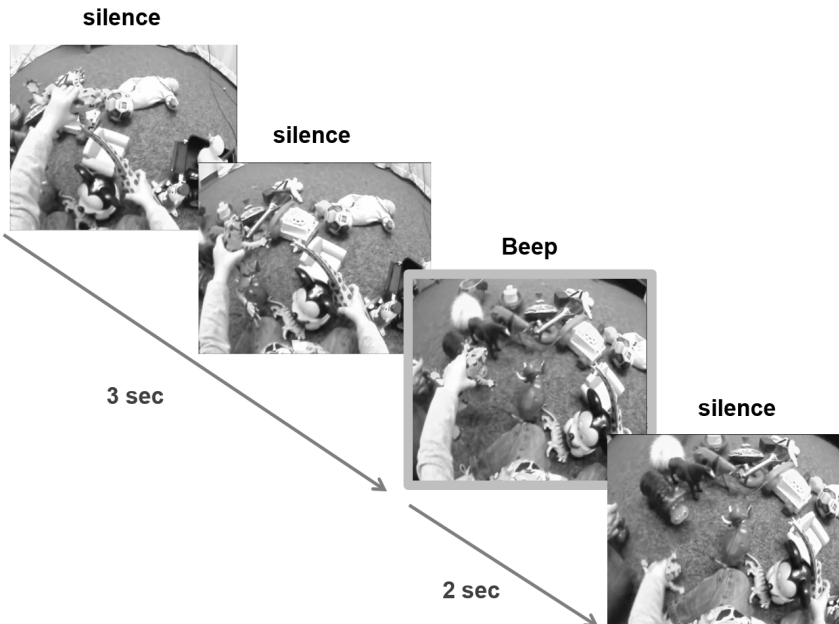


Figure 1. Human Simulation Paradigm

2.4. Results

First, we confirmed the ambiguity level of the 96 trials (36 HI, 60 LI) by calculating the average accuracy using participants' top 1 choices for every naming event. Accuracy of HI trials is indeed high ($M = 0.92$, $SD = 0.01$) while accuracy of LI trials is low ($M = 0.13$, $SD = 0.01$). Figure 2 shows the ambiguity distribution of all naming instances. This finding shows that there is a high degree of variation in natural learning instances. Some learning moments are highly informative (e.g. about 30% learning instances have 100% accuracy), while others are extremely difficult for learners to guess (e.g. about 12% instances have 0% accuracy). Between those two extreme cases, there is a spectrum of instances, which reflects the variation in when parents choose to name objects in natural parent-infant toy play (Tamis-Lemonda, Kuchirko, Luo, Escobar & Bornstein, 2017).

Given such various degrees of informativeness among different parent naming instances, the converging evidence from previous studies suggests that word learners, both adults and children, are likely to identify the correct referent in HI instances (Medina et al., 2011), but there is only a small proportion of those instances in naturalistic interactions (Yurovsky et al., 2013). The debate in the literature lies in low-informative instances – in particular, whether and in what ways LI instances may (or may not) facilitate learning.




Figure 2. Ambiguity distribution of all naming events used in Experiment 1.

To answer this question, we further examined participants' choice patterns. We found three different types of vignettes: (1) HI with high certainty: only one object is dominant and it is the correct referent (Figure 3A); (2) LI with low certainty: multiple possible target options in view and participants are uncertain which one is correct (Figure 3B); (3) LI with high certainty, only one object is dominant but it is not the correct referent (Figure 3C). As shown in Figure 3C, most participants consistently picked the horse as target instead of the correct referent duck due to the fact the horse was the most dominant object in view during naming. For naïve participants, LI trials with high certainty and HI trials contain similar information as both kinds of trials have cues pointing to one particular target. The only difference was that cues in LI trials consistently lead to wrong targets while cues in HI trials always point to the correct one. Participants are certain of their choice but the cues happen to be misleading, thus lead to low learning performance that might take longer to "recover" if the cues are consistently wrong. Another type of LI trial is shown in Figure 3B, the top four most popular choices (hippo, duck, tiger and horse) were all likely to be selected as target, suggesting that learners were unsure which one was correct, thus they tended to switch choices more and keep all four as potential options.

Based on our observation, we categorized a subset of Experiment 1 trials into three types:

- **High informative (H)** trials: above 90% accuracy
- Low informative (L) trials: below 50% accuracy
 - **Ambiguous (A)**: the target toy is among the top 3 choices and accuracy of the top 1 choice does not exceed 50%
 - **Misleading (M)**: the target toy is not the top 1 choice and the possibility of choosing the wrong top 1 choice exceeds 50%

In total, 72 trials were selected and categorized, among which 32 are defined as high-informative, 12 are defined as ambiguous, and 28 are defined as misleading. Detailed descriptive data are shown in Table 1 below.

Table 1. Average accuracies of 3 types of trials

| | 1 st choice | 2 nd choice | 3 rd choice |
|-----------------|--|---|---|
| HI | $M = 0.98, SD = 0.01$ $Min = 0.93, Max = 1$ | $M = 0.97, SD = 0.01$ $Min = 0.93, Max = 1$ | $M = 0.91, SD = 0.01$ $Min = 0.70, Max = 1$ |
| LI - Ambiguous | $M = 0.28, SD = 0.02$ $Min = 0.05, Max = 0.5$ | $M = 0.27, SD = 0.02$ $Min = 0.07, Max = 0.49$ | $M = 0.26, SD = 0.02$ $Min = 0.12, Max = 0.49$ |
| LI - Misleading | $M = 0.08, SD = 0.01$ $Min = 0, Max = 0.33$ | $M = 0.13, SD = 0.01$ $Min = 0, Max = 0.37$ | $M = 0.20, SD = 0.02$ $Min = 0, Max = 0.44$ |

Experiment 1 result confirmed our hypothesis that there are different types of low informativeness within learning moments that could lead to low guessing performance. With more fine-grained 3-choice measures, we were able to see how

certain participants were when making their choices, and then use that information to separate ambiguous LI trials from misleading LI trials.

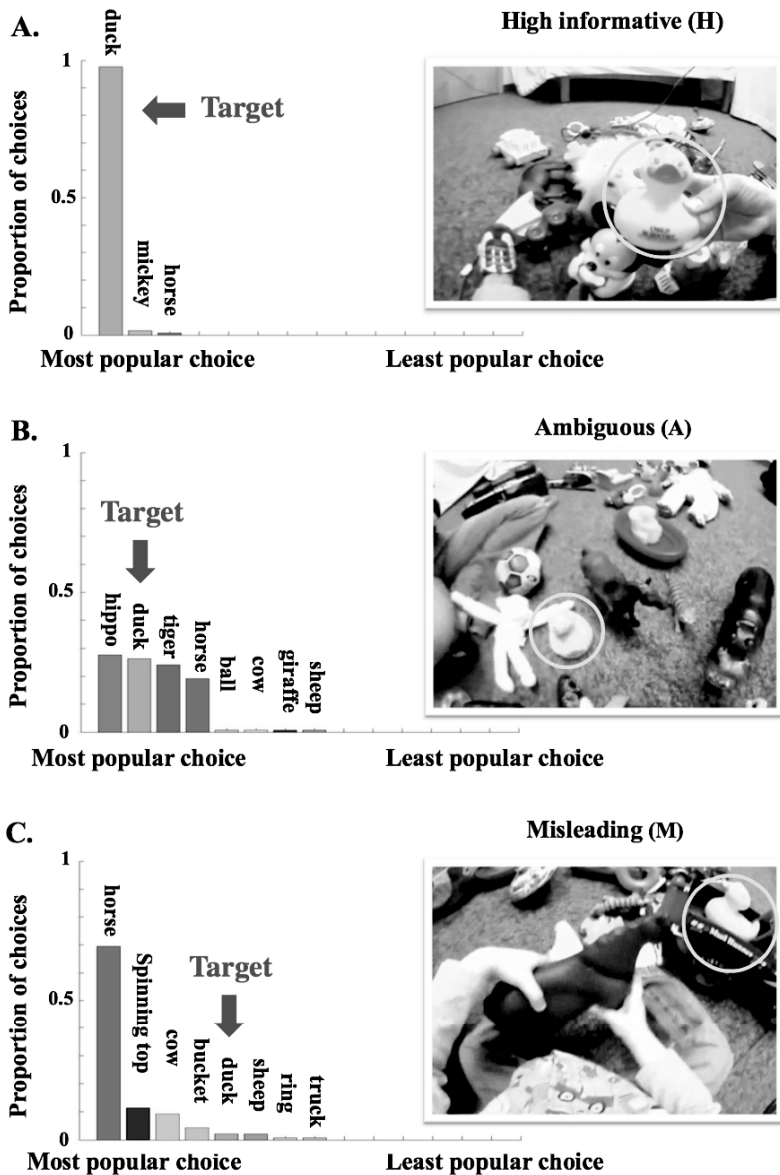


Figure 3: Examples of three different types of learning trials: (A) HI trials with high certainty; (B) LI trials with low certainty; (C) LI trials with high certainty. Arrows indicate the correct referent.

3. Experiment 2: Real-time trial-by-trial integration

Experiment 1 demonstrates that even though LI moments are generally considered as having high referential uncertainty, they can include very different information and may contribute differently to word learning. Thus, in Experiment 2, we asked participants to identify target object by observing a set of learning instances with varying ambiguity. We measured their trial-by-trial guesses to test whether they aggregate evidence across multiple ambiguous and unambiguous learning events and whether learning improves over time. We treated participants' average first-choice accuracies from Experiment 1 as baseline and compared those with participants' cumulative learning results.

3.1. Participants

Fifty-four adults (42 females and 12 males, $M_{\text{age}} = 21.04$, $SD_{\text{age}} = 3.67$) participated for either course credit or payment. None had participated in Experiment 1 or related experiments.

3.2. Stimuli

In this experiment, we used the categorized 72 trials from Experiment 1. Each trial belongs to one of the 3 categories: High informative (H), Ambiguous (A) or Misleading (M). These vignettes were grouped into 12 blocks (6 vignettes per block referring to the same referent). In each block, H, A and M trials were mixed in unique ways. All 6 trials are referring to the same toy within a block. In Experiment 2, we mainly focused on two particular orders in which participants watched low informative trials first: Condition 1 started with two misleading trials followed by two ambiguous and two high informative ones (MMAAHH) and Condition 2 started with two ambiguous trials followed by two misleading and two high informative trials (AAMMHH).

3.3. Procedure

The experimental procedure is demonstrated in Figure 4. Participants were instructed to guess the target object from an array of 25 objects by providing one choice after watching each vignette. In addition, participants were told that they would watch several blocks of 6 vignettes and mothers in these 6 vignettes were naming the same object. Throughout the 6 testing trials, they were allowed to change their guess at any given trial. However, if they believed their previous answer was correct, they could choose the same answer again. They were not allowed to go back and change their previous answers and they were not aware of the degree of ambiguity of each vignette. After each block, a prompt would appear to remind them to get ready for the next block of trials. Again, no feedback was given. As indicated by the circle on each trial in Figure 4, the correct referent of this MMAAHH block is Mickey Mouse.

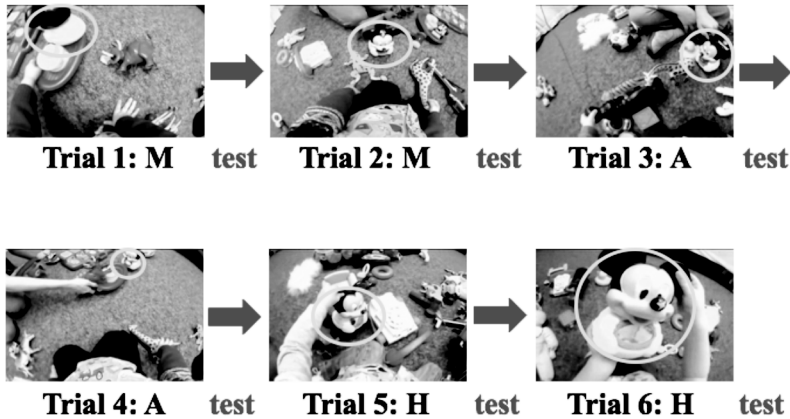


Figure 4. Experiment 2 procedure. Participants provided one choice after watching every trial. Circles (not shown in real trials) indicate the correct referent.

3.4. Results

Target objects' trial-by-trial accuracies for both conditions (MMAAHH, AAMMHH) are shown in Figure 5. Results from Experiment 1 are also plotted on the same graph, which allows us to compare cumulative learning performance with individual trials' baseline measures. From our data, we observed several interesting patterns below:

3.4.1. Learners utilize both types of LI trials in statistical learning

Accuracy on the first trial was similar to the result found in Experiment 1 (Condition 1: $M_{\text{baseline_trial1}} = 0.11$, $M_{\text{Exp2_trial1}} = 0.20$, $t(53) = 1.69$, ns; Condition 2: $M_{\text{baseline_trial1}} = 0.45$, $M_{\text{Exp2_trial1}} = 0.44$, $t(53) = 0.93$, ns), which validated our ambiguity measures.

While accuracies of the first trials were comparable to baseline, they diverged significantly with additional trials of different types. In Condition 1, compared with baseline, any learning trials after the first one (either As or Ms) are better than corresponding baseline (Condition 1: $M_{\text{baseline_trial2}} = 0.02$, $M_{\text{Exp2_trial2}} = 0.30$, $t(53) = 4.41$, $p < 0.001$; $M_{\text{baseline_trial3}} = 0.40$, $M_{\text{Exp2_trial3}} = 0.63$, $t(53) = 3.46$, $p < 0.01$; $M_{\text{baseline_trial4}} = 0.23$, $M_{\text{Exp2_trial4}} = 0.59$, $t(53) = 5.37$, $p < 0.001$). Similar result was found in Condition 2 (Condition 2: $M_{\text{baseline_trial2}} = 0.32$, $M_{\text{Exp2_trial2}} = 0.69$, $t(53) = 5.72$, $p < 0.001$; $M_{\text{baseline_trial3}} = 0$, $M_{\text{Exp2_trial3}} = 0.74$, $t(53) = 12.30$, $p < 0.001$; $M_{\text{baseline_trial4}} = 0.02$, $M_{\text{Exp2_trial4}} = 0.61$, $t(53) = 8.82$, $p < 0.001$). These results replicate Yurovsky et al.'s (2013) findings, showing evidence of continuous learning across ambiguous instances.

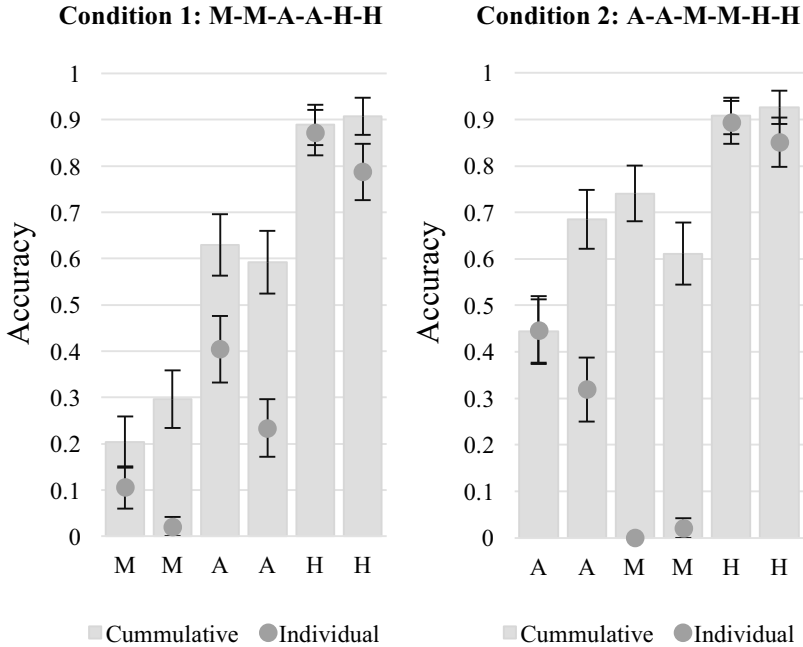


Figure 5. Cumulative and individual trial accuracy.

3.4.2. Effective information integration in learning

For Condition 1, accuracy increased marginally from the first to the second vignette ($t(53) = 1.70, p = 0.09$), but was significantly higher on the third vignette. ($t(53) = 5.15, p < 0.001$). In addition, vignette number and guess accuracy were significantly correlated ($r = 0.96, p < 0.01$). For Condition 2, accuracy increased significantly from the first to the second vignette ($t(53) = 3.74, p < 0.001$), but not from the second to the third vignette ($t(53) = 1.35, ns$). Vignette number is also positively correlated with guess accuracy ($r = 0.86, p = 0.03$). This finding demonstrates trial-by-trial information integration, and ambiguous trials are better than misleading trials for statistical learning as they can provide useful partial knowledge for information aggregation.

3.4.3. Multiple LI trials in a row (ultimately) hinder learning

In both conditions, after the first 3 trials, we observed no improvement on the 4th trial. In Condition 1, the two A trials do not significantly differ ($t(53) = 1.00, ns$). In Condition 2, there is even a significantly decrease of accuracy from the first M trial to the second M trial ($t(53) = 2.44, p = 0.02$). This finding may explain the results found in Medina et al. (2011) showing no accumulative learning when

learning instances are low informative. However, this result could still support the statistical learning view as it could be the case that participants also accumulated negative evidence over time that hurts learning.

3.4.4. HI trials facilitate statistical learning

Although we see a decrease of accuracy on the 4th trial, learners recovered from LI trials by watching HI trials and reached high accuracy performance at the end. Consistent with previous literature, as shown in both conditions, after watching an HI trial (5th trial), participants were very certain about their choices and less likely to switch in the last trial (Condition 1: $M_{\text{trial}_5} = 0.89$, $M_{\text{trial}_6} = 0.91$; Condition 2: $M_{\text{trial}_5} = 0.91$; $M_{\text{trial}_6} = 0.93$). This result supports previous finding that HI trials are helpful for learning word-object mappings (Medina et al., 2011).

In summary, we found converging evidence supporting the statistical learning view that learners are able to accumulate information across trials even when they are misleading or ambiguous and they are able to learn from multiple instances with varying degrees of informativeness. We found that while accumulating correct partial knowledge could facilitate learning, accumulating negative knowledge could also potentially hinder learning. All the evidence supports a gradual learning process in which learners are sensitive to the information contained in different types of learning trials and they integrate information to guide real-time decision making.

4. Discussion

The present study found evidence that word learning moments in which an object is labeled vary in informativeness. Parent may name an object that happens to be the only dominant object in the child's view, which makes the learning moment highly informative or parent may name objects that cannot be easily identified or are not even present in the child's view. Low informative trials contain different kinds of information: 1) Misleading information, learners consistently map the label to the wrong target with high confidence; 2) Ambiguous information, learners select multiple potential targets, and they are unclear which one is correct. Although low informative trials are typically considered as offering ambiguous information, it could be the case that they offer misleading information that takes participants longer to "recover".

We further argue that participants' ability to learn by gradually accumulating partial knowledge from LI moments is sensitive to the types of the information embedded in those naming moments. Thus, in Experiment 2, we examined how a mixture of 3 types of trials (high informative, ambiguous, misleading) influences information aggregation on a trial-by-trial basis. Our finding suggested that cumulative learning performance was affected by trial properties. Although learners achieve high learning performance when the naming moments are unambiguous, learners do not have to rely on those "perfect" moments. Their

learning performance improved even when they watched misleading or ambiguous trials first.

However, even though low informative trials are not useless, too many of them may hurt learning to some degree. This could possibly due to two misleading trials having the same object as the likely target, which may cause learners to switch to the wrong choice. If learners treat consistent misleading trials the same way they treat HI trials, then it may take them multiple HI trials to recover from this mistake. Whether misleading trials in the current design are consistent or not is another interesting question to investigate more in details as they may suggest different levels of ambiguity (within trial vs. between trial ambiguity). Because real-life learning contains all 3 types of learning moments, it is very likely that word learning is a gradual and prolonged process that involves multiple skills that rely on semantic cues, cross-situational statistics, pragmatic and social clues, etc. Learning word-object mappings from low informative contexts is an emerging skill that does not immediately lead to word learning (Bion, Borovsky & Fernald, 2013).

The disagreement from previous studies using similar HSP paradigm can be explained by misleading but not ambiguous learning moments. It could be the case that previous studies (Medina et al., 2011; Yurovsky et al., 2013) used a mixture of misleading and ambiguous trials as their LI stimuli assuming participants treat them in the same way. Thus, we see very different learning outcomes possibly due the amount of misleading information learners gathered from the learning input. The current preliminary results demonstrate that participants are able to learn by gradually accumulating partial knowledge from LI moments and they treat misleading and ambiguous trials differently. This finding may help explain the debate on whether word learning is a “fast mapping” procedure or a gradual statistical one.

To further understand the learning mechanism behind this task, we are currently building a computational model based on the learning patterns observed in our behavioral data. The model assumes learners pick a subset of word-object pairs on each trial, and their confidence about their choices is based on current trial’s level of ambiguity and how consistent it is with previous trials. We hope to find converging evidence to support our behavioral data.

We believe that information contained in naturalistic learning moments vary in terms of their informativeness. Because word learning is a continuous process, different learning instances would greatly influence how learners select information to build word-object mappings.

References

- Baldwin, Dare A. (1991). Infants' contribution to the achievement of joint reference. *Child development*, 62, 874-890.
- Bion, Ricardo AH., Borovsky, Arielle, & Fernald, Anne. (2013). Fast mapping, slow learning: Disambiguation of novel word-object mappings in relation to vocabulary learning at 18, 24, and 30months. *Cognition*, 126, 39-53.

- Bloom, Paul. (2000). *How children learn the meanings of words* (Vol. 377). Cambridge, MA: MIT press.
- Carey, Susan, & Bartlett, Elsa. (1978). Acquiring a single new word.
- Fisher, Cynthia, Klingler, Stacy L., & Song, Hyun-joo. (2006). What does syntax say about space? 2-year-olds use sentence structure to learn new prepositions. *Cognition*, *101*, B19-B29.
- Gillette, Jane, Gleitman, Henry, Gleitman, Lila, & Lederer, Anne. (1999). Human simulations of vocabulary learning. *Cognition*, *73*, 135-176.
- Mather, Emily, & Plunkett, Kim. (2012). The role of novelty in early word learning. *Cognitive Science*, *36*, 1157-1177.
- Markman, Ellen. M. (1991). The whole-object, taxonomic, and mutual exclusivity assumptions as initial constraints on word meanings. *Perspectives on language and thought: Interrelations in development*, 72-106.
- Medina, Tamara N., Snedeker, Jesse, Trueswell, John C., & Gleitman, Lila R. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences of the United States of America*, *108*, 9014-9019.
- Quine, Willard V. (1960). *Word and Object*. MIT Press. Cambridge, MA.
- Smith, Linda, & Yu, Chen. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*, 1558-1568.
- Smith, Linda B., & Yu, Chen. (2013). Visual attention is not enough: Individual differences in statistical word-referent learning in infants. *Language Learning and Development*, *9*, 25-49.
- Tamis-LeMonda, Catherine S., Kuchirko, Yana, Luo, Rufan, Escobar, Kelly, & Bornstein, Marc H. (2017). Power in methods: language to infants in structured and naturalistic contexts. *Developmental Science*.
- Tomasello, Michael. (2003). *Constructing a language: A usage-based approach to child language acquisition*. Cambridge, MA: Harvard University Press
- Yu, Chen, & Smith, Linda B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, *18*, 414-420.
- Yu, Chen, & Smith, Linda B. (2012). Embodied attention and word learning by toddlers. *Cognition*, *125*, 244-262.
- Yuan, Sylvia, & Fisher, Cynthia. (2009). 'Really? She blicked the baby?': Two-year-olds learn combinatorial facts about verbs by listening. *Psychological Science*, *20*, 619-626.
- Yurovsky, Daniel, Smith, Linda B., & Yu, Chen. (2013). Statistical word learning at scale: The baby's view is better. *Developmental Science*, *16*, 959-966.
- Zhang, Yayun, Yurovsky, Danial, & Yu, Chen. (2015). Statistical Word Learning is a Continuous Process: Evidence from the Human Simulation Paradigm. *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*.

Proceedings of the 41st annual Boston University Conference on Language Development

edited by Maria LaMendola
and Jennifer Scott

Cascadilla Press Somerville, MA 2017

Copyright information

Proceedings of the 41st annual Boston University Conference on Language Development
© 2017 Cascadilla Press. All rights reserved

Copyright notices are located at the bottom of the first page of each paper.
Reprints for course packs can be authorized by Cascadilla Press.

ISSN 1080-692X

ISBN 978-1-57473-076-0 (2 volume set, paperback)

ISBN 978-1-57473-176-7 (2 volume set, library binding)

Ordering information

To order a copy of the proceedings or to place a standing order, contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, sales@cascadilla.com, www.cascadilla.com