

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Multiple Proposal Memory in Observational Word Learning

Permalink

<https://escholarship.org/uc/item/5gp8j5qs>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Koehne, Judith
Trueswell, John C.
Gleitman, Lila R.

Publication Date

2013

Peer reviewed

Multiple Proposal Memory in Observational Word Learning

Judith Koehne (judith.koehne@uni-bamberg.de)

Otto-Friedrich-Universität Bamberg
Markusplatz 3, 96045 Bamberg

John C. Trueswell (trueswel@psych.upenn.edu)

Lila R. Gleitman (gleitman@psych.upenn.edu)

Institute for Research in Cognitive Science, University of Pennsylvania,
3401 Walnut Street, Suite 400A, Philadelphia, PA, USA

Abstract

The temporal co-occurrence of a novel word and a visual referent undoubtedly facilitates establishing the meaning of a word. It is less understood, however, how precisely learners can keep track of the frequencies of these co-occurrences across situations. Observational learning may rely on one or few highly informative exposures (propose-but-verify) or it may be driven by the collection of evidence in a more gradual and parallel manner (multiple-hypotheses tracking). We evaluated both hypotheses within two experiments and found that learners were able to keep track of more than one hypothesis for a novel word. However, this memory was strongly dependent on each learner's individual learning path (i.e., which meanings they had considered before) and influenced by the order of presentation of potential referents. We argue for an account of a multiple-proposal memory rather than a multiple co-occurrence memory.

Keywords: observational word learning; memory; cross-situational analysis; multiple hypotheses tracking; propose-but-verify; individual learning paths

Observational Word Learning

While observing the world can be a very direct path to the meaning of a novel word (*fast mapping*, Carey, 1978), the relationship between both sources of input is often too ambiguous to make a promising immediate guess. The learner could solve this problem in various ways: On each learning instance, she could store multiple possible solutions and then identify the best solution across several learning instances through an intersective process, an assumption that is commonly understood to underlie the idea of *cross-situational word learning* (Quine, 1960; Yu & Smith, 2007). Alternatively, she could make an immediate guess about the word's meaning and wait for confirmation or rejection. In this case, the learner would have no memory for the alternatives that were not guessed, but maximally a memory for the different guesses tried along the way until the correct one is identified.

While experiments reported in Medina, Snedeker, Trueswell, & Gleitman (2011) and Trueswell, Medina, Hafri, & Gleitman (2013) support this latter idea (*propose-but-verify* account), other studies indicate that learners are able to extract multiple hypotheses on each learning instance (Vouloumanos, 2008; Vouloumanos & Werker, 2009; Koe-

hne & Crocker, 2011). An important aspect that is ignored in these studies, however, is the role that each learner's individual learning path plays. It is therefore unclear whether one and the same person in fact stores multiple possibilities for a word and if this is the case, under which circumstances. One factor that has been shown to be relevant to this question is the order in which the language novice has encountered and re-encountered potential referents (Medina et al., 2011).

We evaluate the way learners exploit observational word learning situations within two experiments, employing the standard paradigm of psychologically investigating cross-situational word learning. Importantly, we consider both the learner's individual learning path and the order of exposures and re-exposures of potential referents. We moreover address the possibility that the different outcomes in different studies may be due to the implemented experimental procedure. In particular, we compare a procedure, in which participants make a choice on each learning trial (Exp. 1) to a passive look-and-listen learning phase (Exp. 2).

Learning based on Co-occurrence Frequencies

Trueswell et al. (2013) examined learners' memory in observational learning situations in a series of experiments. During the learning phase, participants were presented two or five visual referents and a spoken sentence containing one novel noun per trial. The task was to choose that referent (by mouse click) in each trial that the learner believed to match with the novel noun. Trueswell et al. found that even if the learning situations were greatly simplified but still ambiguous (just two possible referents), participants later showed no sign of memory for any referent other than the one they had selected. Specifically, when a learner re-encountered a noun (e.g., *mipen*), he was at chance at selecting the correct referent (e.g., bear) if he had made the wrong choice the previous time he had encountered *mipen* (e.g., if he had chosen door rather than the correct bear). Had he remembered that the unselected (but correct) referent (bear) had co-occurred with *mipen*, he could have unambiguously identified it as correct in the current situation.

Interestingly, other studies indicate that learners are able to precisely differentiate the co-occurrence frequencies of different alternatives for one noun. Vouloumanos (2008)

employed a passive look-and-listen learning phase with one referent and one noun per trial. Over the course of the experiment, each noun co-occurred with several referents with varying frequencies. In a final forced-choice vocabulary test, learners could differentiate between these alternatives based on small differences in their co-occurrence statistics. However, since there was only one referent per learning trial, this study does not answer the question whether multiple possibilities are memorized from one situation. Addressing this issue, Koehne & Crocker (2011) integrated a learning procedure with four objects depicted for each novel noun. As in Vouloumanos (2008), nouns co-occurred with objects with different frequencies (83%, 50%, and 17%). Interestingly, when the 83% referent was not available in a final forced-choice test, learners preferred the 50% referent over 17% alternatives. This result suggests sensitivity to differences in co-occurrence statistics even when learning trials are ambiguous.

Differences between Trueswell et al. and Koehne & Crocker could be due to the experimental procedure (forced choice vs. look-and-listen during learning). However, individual learning paths were not considered in Koehne & Crocker: It is unclear whether selecting the 50% referent depended on the choices, or proposals, the learner had made before and whether one and the same learner had stored multiple alternatives for one noun.

Indeed, as noted by Trueswell et al. (2013), the strictest version of a *propose-but-verify* procedure, in which only a single meaning is ever maintained, is inadequate because it fails to explain the learning of ambiguous words. They therefore propose that “when a confirmed (and even re-confirmed) hypothesis for a word is then not supported by a later context, the learner would actively search memory for past rejected hypotheses, and may ... establish a second meaning for the word.” Here we call this *multiple-proposal* memory, in which only previously proposed meanings are available in memory rather than entire referential sets from past learning instances (i.e., the context) as stipulated by the most common cross-situational accounts.

The method used in the two experiments presented here allows us to differentiate between the predictions from the propose-but-verify versus the multiple-hypothesis-tracking account. In particular, it addresses the question whether *one and the same* learner keeps track of more than one hypothesis for a novel word.

Experiment 1

Experiment 1 addresses the questions of whether and how learners track multiple meanings for a novel word and what role both the learning path and the order of (re-) exposures of potential referents play in this process.

Methods

Participants 36 participants were tested, four of which had to be excluded due to technical and eye-tracking problems.

Data of 32 participants (11 Male, average age 22) was analyzed.

Design, Materials, & Procedure The overall task of Experiment 1 was to learn the meanings of 16 novel nouns. Learning trials consisted of one spoken English sentence containing one of the novel words (e.g., *I see a moke!*) and four objects that were depicted on the screen. During training, each noun had six learning trials, intermixed with the other learning trials. Crucially each of the 16 nouns was assigned two meanings with different co-occurrence frequencies: One referent was present whenever the noun was present (six times, *100% referent*, e.g., television), the other referent was present in only half of the cases the noun was present (three times, *50% referent*, e.g., dog). All other objects co-occurred only once with a noun (17%). We manipulated the order in which trials including and excluding the 50% referent were presented within four levels (within participants): Firstly, the 50%-present (P) and 50%-absent (A) trials could be either blocked (**AAAPPP** and **PPPAAA**) or not blocked (**APAPAP** and **PAPAPA**); secondly, the first encounter of a noun could be either an A trial (**AAAPPP** and **APAPAP**) or a P trial (**PPPAAA** and **PAPAPA**).

On each learning trial, participants selected by mouse click the referent they thought belonged to the novel noun. After each response, they gave a confidence rating for their selection (on a scale from 1 to 9). No feedback was given and participants were *not* informed that nouns may have multiple meanings.

After all six learning trials had been encountered for a word, a final test was given for each word, in which eight objects and one spoken word were presented and learners were asked to again select the matching referent and indicate their confidence. The 100% referent, however, was not available which means that the 50% object was the one with the highest co-occurrence rate - all other objects were 0% and 17% referents.

The experiment consisted of two parts: Eight novel nouns were taught and tested (Block 1) before the other eight nouns were taught and tested (Block 2). Order of presentation of learning and test trials was pseudo-randomized: Between two exposures of the same noun, there was always at least one but not more than 8 trials with other nouns. Participants were run individually and the experiment lasted approximately 30 minutes.

Predictions Standard cross-situational accounts (such as Yu & Smith, 2007, Vouloumanos, 2008, and Koehne & Crocker, 2001) predict that learners precisely keep track of the co-occurrence frequencies between nouns and referents. The 50% referent should therefore be chosen at final test above chance in all conditions, independent of both the learning path and the order of (re-)exposures of 50%-present trials.

According to a strict *propose-but-verify* account selection of the 50% alternative at final test would occur if and only if it is the current working hypothesized meaning - that is, if

the 50% referent had been selected on the preceding learning instance. This is impossible when the 50% referent is Absent on the last learning trial, predicting chance performance in conditions PPPAAA and PAPAPA. During the final test in the other two conditions, the 50% alternative would be selected above chance on those rare occasions when the learner had selected the 50% referent on the last learning instance (i.e., when they failed to learn the 100% target by Instance 6).

According to the weaker *propose-but-verify* account, a final test with the 100% referent absent will trigger consideration of all past *proposed* meanings. This means that above-chance performance on the 50% referent is expected if and only if the learner had previously *selected* (clicked on) a 50% referent during the learning phase. One might expect such a memory to have a recency component: More recently proposed meanings will be easier to remember. Moreover, early encounters of a string of 50% referents (i.e., PPPAAA) will increase the probability that this referent will be selected during the learning phase and thus more likely to be recalled at test. Conversely, it is very unlikely during learning that the 50% referent will be selected on any trial when these occurrences are grouped late in the sequence (AAAPPP): Most learners will have already locked onto the 100% item as the referent by Instance 4, and thus rarely select the 50% referent during learning. Therefore, they will not select it at test either.

Data Analysis, Results, & Discussion

The results are most consistent with the weaker *propose-but-verify* account. Across conditions, participants selected the 50% referent in the final test significantly more often than chance (25.4% vs. 12.5% chance; $t(31) = 7.77, p < .001$)¹. Both confidence ratings and reaction times support that this difference is meaningful: Ratings were significantly higher ($\chi^2(1) = 17.87, p < .001$) and reaction times were significantly lower ($\chi^2(1) = 9.36, p < .01$) when the 50% referent was chosen than when it was not. Moreover, it was chosen significantly more often than any other of the seven (0% and 17%) objects.

While this trend holds for all four conditions, differences to chance were significant only in Conditions PPPAAA (34.4%; $t(31) = 4.91, p < .001$), PAPAPA (27.3%; $t(31) = 4.32, p < .001$), and APAPAP (23.4%; $t(31) = 2.80, p < .05$) but not in AAAPPP (16.4%; $t(31) = 1.02, p = .32$; Figure 1). This finding is inconsistent with a standard cross-situational account because all conditions should have been above chance independent of presentation order. It is also inconsistent with the strict *propose-but-verify* account because PPPAAA and PAPAPA ought not be above chance, but they are. Consistent with the weaker *propose-but-verify*, PPPAAA offers the best performance overall whereas AAAPPP offers the worst.

To get insight into the roles of ordering and learning paths on the final test, we analyzed the effects of *Condition* and

Previous Selection of the 50% Referent, that is, whether the 50% referent had been chosen in the previous encounter when it had been present. Note that this trial was in different positions depending on condition: It was the last trial in Conditions AAAPPP and APAPAP, the second to last trial in Condition PAPAPA, and the fourth to last trial in Condition PPPAAA.




Figure 1: Selections in test, Exp. 1

Consistent with the weaker version of *propose-but-verify*, we found that participants were only above chance at selection of the 50% referent at test if they had selected the referent on its last encounter during learning (Figure 2). Note that the number of observations contributing to each proportion differs in the way expected if learners were using the weaker *propose-but-verify* procedure during the learning phase; the 50% referent was selected on the previous encounter during learning only 24 times (out of 128) in AAAPPP, but 56 times in PPPAAA. In APAPAP it was chosen 27 times and in PAPAPA 32 times. If selected during learning however, it was recalled at final test at similar rates regardless of condition (i.e., Figure 2).

To confirm the reliability of the effects in Figure 2, we conducted a multi-level logistic regression using *Condition* and *Previous 50% Referent Accuracy* as predictors of selecting the 50% referent at test, entering both as fixed effects (using the lme4 package in R, Bates, 2005). Random intercepts and slopes of Subjects and Items were integrated. If a model did not converge, random effects were reduced until convergence was reached (always discarding the random effect with the smallest effect). Main effects were tested using model comparison (Chi-Square values are reported; Baayen, Davidson, & Bates, 2008). We found a significant effect of *Previous 50% Referent Accuracy* only ($\chi^2(1) = 80.21, p < .001$) but no effect of *Condition* ($\chi^2(3) = 3.67, p = .30$) and no interaction ($\chi^2(3) = 4.52, p = .21$). T-tests confirm that for that subset of trials for which it was not the case that the 50% referent had been chosen in the previous learning trial in which it had been present, selecting the 50% referent was not above chance ($t(31) = -.68, p = .50$). This reveals that, independent of condition, the 50% referent was only chosen reliably if it had also been chosen in the previous encounter for which it had been present.

¹ All t-tests are two-tailed.

Interestingly, 50% selection was still above chance when it additionally was the case that the 100% referent had been chosen two to five times during learning ($t(25) = 6.43, p < .001$). This means that one and the same learner could consider the 100% referent as the correct referent and still be sensitive to the fact that the 50% referent was a better candidate than the 17% objects as long as the 50% referent, as well, had been considered.

This pattern of results supports the weaker version of the *propose-but-verify* account: While in fact a referent is only stored as the potential meaning if it has been actively considered before, this consideration does not need to happen in the absolutely previous encounter of the noun but only in the last *common* encounter of the noun and that referent. This means that learners do not only memorize the last guess they made for a noun but also less recent guesses. Our results are clearly not in line with the hypothesis that learners are equipped with a general multiple co-occurrence memory.




Figure 2: 50% referent selections in test, Exp. 1

It is possible that the results from Experiment affected by the employed learning procedure: Forcing a selection on each trial may enforce the influence of the learning path (i.e., previous accuracy). We address this possibility in Experiment 2.

Experiment 2

Experiment 2 investigates whether learning path and conditions have the same effect on memorizing potential referents if learners are not forced to make a choice on learning trials.

Methods

Participants 39 participants were tested, seven of which had to be excluded due to technical and eye-tracking problems. Data of 32 participants (16 Male, average age 23) was analyzed.

Design, Materials, & Procedure The learning paradigm, design, materials, and procedure were exactly the same as in Experiment 1 except that participants were asked to simply look and listen during learning trials while trying to figure out what the novel nouns mean. As in Experiment 2, however, trial change was self-paced (elicited by button press). Moreover, participants' eyes were tracked using a Tobii 1750 eye-tracker (sampling rate 50 Hz).

Predictions Hypothesizing that clicking does not influence the learner's behavior predicts that one will find the same results as in Experiment 1. Hypothesizing that clicking enforces previous accuracy to be crucial on the other hand predicts a weaker effect of the learning path on the memory for the 50% referent.

Data Analysis, Results, & Discussion

Selecting the 50% referent in the test again was significantly more frequent than would be expected by chance (22.7% vs. 12.5%; $t(31) = 6.07, p < .001$) and than selecting any of the other candidates. As in Experiment 1, confidence ratings were higher ($\chi^2(1) = 13.12, p < .001$) and reaction times were lower ($\chi^2(1) = 5.12, p < .05$) when the 50% referent was selected than when another object was chosen.

Selection rates were (at least marginally) significantly above chance in all four conditions (PPPAAA: 29.9%, $t(31) = 4.53, p < .001$; PAPAPA: 18.8%, $t(31) = 1.76, p = .09$; APAPAP: 22.7%, $t(31) = 3.13, p < .01$; AAAPPP: 19.5%, $t(31) = 1.83, p = .08$, Figure 3).




Figure 3: Selections in test, Exp. 2

To evaluate the effect of the learning path although no choices were made during learning, we used learners' eye movements as a predictor: Specifically, we coded test trials for the frequency of 50%-present learning trials in which the 50% referent had been fixated more often than any of the three other candidates after the novel noun was presented (i.e., from onset of the noun until the self-paced end of the trial). The rationale of this coding was that looking at a referent most reveals that participants had paid attention to it, indicating that it was selected as the potential referent.

We then included this measurement of *Previous 50% Referent Accuracy* as a predictor, together with *Condition* (Fig-

ure 4). Similar to Experiment 1, we found that choosing the 50% referent at test was not predicted by Condition ($\chi^2(3) = .96, p = .31$) but by *Previous 50% Referent Accuracy* ($\chi^2(1) = 7.49, p < .01$). Again, there was no interaction ($\chi^2(3) = 0.89, p = .83$). And again the number observations across conditions patterned like in Experiment 1 in terms of how often the 50% referent was ‘selected’ (by eye) on its last occurrence during learning (N = 22 for AAAPPP; N = 44 for PPPAAA; N = 24 for APAPAP; and N = 23 for PAPAPA).




Figure 4: 50% referent selections in test, Exp. 2

Interestingly, however, the 50% referent was still chosen significantly more often than chance at test if it was *not* looked at most often in the previous encounter ($t(31) = 2.43, p < .05$). We therefore also coded test trials for whether the 50% referent had been looked at most in any (i.e., at least one) learning trial (*Any Accuracy*). We found that if this was not the case, selecting the 50% referent was not more frequent than chance ($t(21) = -.21, p = .83$). *Any Accuracy* was a marginally significant predictor ($\chi^2(1) = 3.44, p = .06$) whereas *Condition* was not ($\chi^2(3) = 3.45, p = .33$) and both did not interact ($\chi^2(3) = 1.27, p = .74$).

Similar to Experiment 1, having looked at the 100% referent most often in two to five learning trials did not change this pattern: The 50% referent was still chosen significantly more often than chance as long as it was also looked at most at least once ($t(31) = 3.80, p < .001$).

These results suggest that learners’ behavior when choices were not forced during learning was similar to their behavior when they were forced to respond (i.e., as in Experiment 1). While it may be less crucial that the 50% referent was paid particular attention to exactly the last time it was encountered, the data indicates that it is necessary that it at some point in learning it had been attended to. While this difference could suggest that memory in Experiment 2 was better than in Experiment 1 (i.e., that learners stored all proposals rather than only the last one), the different measurements of *Previous 50% Referent Accuracy* cannot be perfectly compared with one another.

Most important, however, is that even if the learner is not forced to make decisions during learning, it is still crucial for a potential referent to be paid particular attention to at some point. We interpret this as a confirmation of our findings from Experiment 1: Learners show no sign of a general multiple co-occurrence memory but they are able to memorize more than one proposal they have made.

Analyses Experiments 1 & 2

In order to evaluate a potential difference between Experiments 1 and 2 regarding the influence of Condition, we entered data from both into one analysis. *Experiment* (Experiment 1: click vs. Experiment 2: no click) and *Condition* were used as fixed factors. We found a marginal effect of *Condition* ($\chi^2(3) = 7.61, p = .06$), no effect of *Experiment* ($\chi^2(1) = 1.59, p = .21$), and no interaction ($\chi^2(3) = 3.22, p = .36$; Figure 5). We then grouped the four conditions into two: 50% present in first trial (PPPAAA & PAPAPA) versus 50% absent in first trial (AAAPPP & APAPAP) and repeated the analysis. While selecting the 50% object was significantly more frequent in the first-trial present than the first-trial absent conditions ($\chi^2(1) = 6.63, p < .05$), still neither an effect of *Experiment* ($\chi^2(1) = 1.04, p = .31$) nor an interaction was found ($\chi^2(1) = 2.31, p = .13$). Within experiments, however, both condition groups differed significantly only for Experiment 1 ($\chi^2(1) = 8.04, p < .01$) but not for Experiment 2 ($\chi^2(1) = 0.51, p = .47$). It is therefore not quite clear whether the order of exposure and re-exposure was equally meaningful to both Experiments. Possibly, it was slightly more important in Experiment 1 than Experiment 2 that a referent’s first encounter happened early, as also indicated by the missing significance of selecting the 50% referent in Condition AAAPPP in Experiment 1 (Figure 1). Either way, for both experiments, the effect of *Previous Accuracy* was a much clearer predictor than *Condition*.

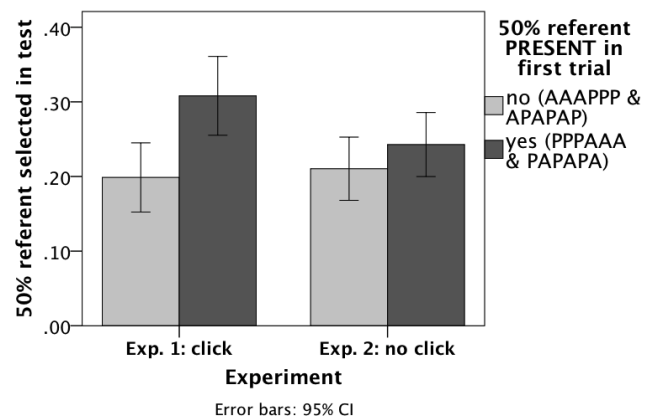


Figure 5: 50% referent selections in test, Exp. 1 & 2

Conclusions & General Discussion

Results from Experiments 1 and 2 reveal that learners successfully learned to differentiate between co-occurrence

frequencies of 50% versus 17% and 0% even though another referent co-occurred perfectly (100%). However, this was only the case if the 50% referent was in the learner's attention at least once before (or if it even was actively selected). Importantly, the 50% referent was also stored even if it was not the only referent that the learner had considered (i.e., when both the 100% referent and the 50% referent were in the learner's focus of attention at some point during learning). These findings clearly reveal that while co-occurrences were not generally all stored, 'multiple proposal' memory is possible in observational word learning. This is not in line with the standard cross-situational account whereas it generally supports a *propose-but-verify* account. Interestingly though, selecting the 50% referent was above chance in conditions PPPAAA and PAPAPA; unlike a strict *propose-but-verify* theory would predict, learners can memorize more than the most recent choice they have made.

Asking participants to select a referent during learning trials did not generally suppress memorizing multiple-proposals. While it may be the case that it is more important for a forced-choice learning procedure than the non-forced choice one that the 50% referent is considered exactly in the previous encounter of it, a clear comparison between choosing and looking is impossible. If the difference is real, it would indicate that forcing a choice enhances the role of previous consideration, possibly because a stronger memory trace is built by actively (and physically) making a selection than by mental consideration.

Our results moreover at least indicate that there is a possible influence of the order in which referents are firstly encountered and re-encountered: Early on, when the hypotheses space is still completely open, learners are more willing to memorize co-occurring objects as potential meanings than later, when other hypotheses (or considerations) have already been made for a novel noun. This may be more strongly the case when selections are forced even early on in learning (in Experiment 1).

Summary

We investigated learners' memory for co-occurrence frequencies in referentially ambiguous observational-word learning situations within two experiments. Our data reveals that while participants were able to recall more than one potential meaning for a noun, this memory was dependent on the person's single considerations during learning: Only if a potential meaning had been proposed before (i.e., selected or paid particular attention to), it was stored. However, learners memorized more than the most recent proposal they had made for a novel word. Moreover, a meaning was more likely to be proposed if it co-occurred with a noun early on the learning path. While this whole pattern was very similar independent of the learning procedure (choice made during learning, Experiment 1, vs. no choices made, Experiment 2), the influence of being proposed early may be enhanced when choices are made. In line with a moderate version of the *propose-but-verify* account (Medina et al.,

2011; Trueswell et al., 2013), our results can be accounted for by a multiple-proposal memory rather than a multiple-co-occurrence memory. Indeed, such a procedure is logically necessary to explain the learning of words with more than one meaning (i.e., homophones). Future research is necessary to explore the conditions under which ambiguous words are successfully learned, taking into account the mutually exclusive occurrence of appropriate referents (Meaning 1 vs. Meaning 2), which was not modeled experimentally here (i.e., the 100% referent was simultaneously present alongside the 50% referent on each "P" learning trial). Moreover, other distinguishing contextual features likely support the differentiation of two meanings for the same word. Finally, future work must examine how well these observations hold for naturally occurring word-learning environments in which referential ambiguity is greater and the contexts of word use are more variable. Artificial stimuli like those used here offer better experimental control and thus allow for closer examination of the learning mechanism but do not address how this mechanism responds to more typical input (Medina et al., 2011; Trueswell et al., 2013).

References

- Baayen, R., Davidson, D., & Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412.
- Bates, D. (2005). Fitting linear mixed models in r. *R News*, (pp. 27–30).
- Carey, S. (1978). The child as a word learner. In J. Bresnan & G. Miller (Eds.), *Linguistic theory and psychological reality* (pp. 264–293). Cambridge, MA: MIT Press.
- Koehne, J. & Crocker, M.W. (2011). The interplay of multiple mechanisms in word learning. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the 33rd Annual meeting of the Cognitive Science Society* (pp. 1930-1936). Austin, TX: Cognitive Science Society.
- Medina, T., Snedeker, J., Trueswell, J., & Gleitman, L. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences of the United States of America*, 108, 914–919.
- Quine, W. (1960). *Word and Object*. Cambridge, MA: Cambridge University Press.
- Trueswell, J., Medina, T., Hafri, A., & Gleitman, L. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66, 126–156.
- Vouloumanos, A. & Werker, J. F. (2009). Infants' learning of novel words in a stochastic environment. *Developmental Psychology*, 45(6), 1611–7.
- Vouloumanos, A. (2008). Fine-grained sensitivity to statistical information in adult word learning. *Cognition*, 107(2), 729–42.
- Yu, C. & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science: A journal of the American Psychological Society*, 18(5), 414–20.