The Role of Feedback in Learning Form-Meaning Mappings

Patrick Jeuniaux (pjeuniau@memphis.edu)
Rick Dale (radale@memphis.edu)
Max M. Louwerse (mlouwers@memphis.edu)
Department of Psychology
Institute for Intelligent Systems
Memphis, TN 38152 USA

Abstract

Word learning is a complex activity whose mechanisms are not fully understood. For instance, there is an ongoing debate whether feedback is fundamental to learning the meaning of words. Three experiments and a simulation aimed to investigate the importance and role of feedback in learning a six-word vocabulary of an artificial language. In the experiments, participants had to guess what object in a visual array of objects was referred to by an utterance composed of unknown words. The words referred to objects characteristics like color, size or shape. Results showed that feedback led to a superior learning rate. However, some learning did occur without feedback. A simulation further illustrated possible mechanisms responsible for learning with and without feedback. These data are consistent with the idea that the mechanism by which feedback particularly helps learning consists of reducing the amount of ambiguity in the environment.

Keywords: language acquisition, artificial word learning, feedback, ambiguity, computational modeling, psycholinguistics

Introduction

Language acquisition involves, at least as a substantial part, the mapping between word forms and their meaning. There remain considerable questions about how this mapping unfolds, and specifically, how feedback works to help learning, if it does at all. The current study investigates this issue in three experiments and a computational model whereby participants learn a six-word vocabulary of an artificial language.

Learning Form-Meaning Mappings

The notion of meaning can be thought of as a reference to the perceptual world (Smith & Yu, 2008) or as an association to other words (Landauer & Dumais, 1997). Although these accounts are not mutually exclusive, the grounding in the perceptual world is generally considered as a more basic and early learning process than the associations of word forms to each other (Werker, Cohen, Lloyd, Casasola, & Stager, 1998).

The task that the child faces in acquiring a language is to establish the appropriate mapping between forms and meaning. The acquisition of that mapping, as well as the development of other aspects of language (e.g., syntax), seems to emerge seamlessly given that most children come to have a good command of their native language. It is however not clear by what mechanisms children learn that mapping.

Like all inductive problems, language learning cannot be purely based on experience but requires a bias that constrains the alternatives under consideration (Mitchell, 1997). Theories of language acquisition locate the bias either in the individual, in the environment, or in both.

At the individual level, the child has been described as being equipped with innate linguistic knowledge (Anderson & Lightfoot, 2000), learning preferences (Markman, 1991) or special cognitive characteristics (Kersten & Earles, 2001). At the environmental level, the child is believed to benefit from explicit or implicit cues helping the learning process. Explicit cues consist of the information that the caregiver intentionally provides to the child in order to improve his or her language abilities. Implicit cues are the information that the child, unbeknownst to the caregiver, can use to infer knowledge of language (Regier, 1996).

The Role of Feedback

Feedback is one of the explicit cues caregivers can use in the word-learning process. It is a reaction directed to the child on the basis of one of the child's behaviors, and can be either positive or negative. A parent approving a child's production is an instance of positive feedback (e.g., *yes, doggie!*), while a parent correcting a child's mistake (e.g., calling a dog a *cat*) is negative feedback (e.g., *no, this is not a cat, it's a dog*).

The role of feedback in language acquisition is still a matter of some dispute (cf. Chouinard & Clark, 2003). Two broad types of theoretical postures regarding feedback can be contrasted: either considering feedback as irrelevant to language acquisition or considering feedback as relevant but sometimes present in hidden forms in the child's environment. On the feedback-is-irrelevant side, it has been argued that children generally learn language without feedback (Brown & Hanlon, 1970; cf. Regier, 1996). On the feedback-is-relevant-but-hidden side, it has been argued that feedback is helpful but generally delivered in implicit forms (Chouinard & Clark, 2003).

In the domain of learning in general, the mechanisms by which explicit feedback may be helpful are not clear (Pashler, Cepeda, Wixted, & Rohrer, 2005). Presenting feedback has been thought of as potentially competing with the working memory resources, thereby reducing long-term retention (Scholer & Anderson, 1990), while withholding it has been thought of as a way of forcing the learner to engage in deeper processing, thereby increasing long-term retention (cf. Pashler et al., 2005). On the other hand,

feedback has been presented as an efficient way of focusing attention on the relevant statistical patterns in the raw input (Goldstein & Schwade, 2008).

Most studies on the importance and role of feedback have either been observational (e.g., Brown & Hanlon, 1970) or computational (e.g., Regier, 1996). An intermediate approach is to experimentally investigate phenomena which reveal important aspects of the acquisition process in children and adults (e.g., Goldstein & Schwade, 2008; Pashler et al., 2005).

The current study attempts to shed light on the role of feedback by examining its impact in three experiments, as well as its potential mechanisms in a simulation. Similar to Dale and Christiansen (2004) and Yu and Smith (2007), learning is examined here as the mapping between visual referents and word forms learned by adults.

Experiment 1

In Experiment 1, participants attempted to learn the meaning of words, by guessing what object in the visual field an artificial utterance is referring to.

Method

Participants 37 students from the University of Memphis participated for course credit.

Procedure Participants were randomly assigned to one of two conditions: feedback ($n = 22$) or no-feedback ($n = 15$). A session had two phases: a training phase and a testing phase. In the training phase, participants played 140 trials of the guessing game, in which they learned an artificial vocabulary composed of six words. A scene made of objects appeared, followed by an utterance referring to one of the objects shown on the screen (the topic). The task of the participant was to click on the object which he or she believed was the topic of the utterance. After giving their choice, the participants in the feedback condition were told if they were correct or incorrect, and the appropriate answer was highlighted in case of an incorrect answer. Even though the correct answer was identified, the learner still needed to infer which feature of the object corresponded to which word. In the no-feedback condition the participant's choice was primarily a guess over time possibly becoming a more educated guess because of the repetition of words and objects. For both the feedback and the no-feedback participants the participants' knowledge of the vocabulary was formally evaluated during a testing phase at the end of the experiment. The testing phase required participants to choose the correct word from six options.

Materials Each sequence corresponded to the pairing between the display of an utterance made of one, two or three artificial words, and a scene made of minimum two and maximum eight objects (see Fig. 1). The six artificial words to be learned were non-words based upon Creel, Aslin, and Tanenhaus (2006). These words corresponded to six visual features of the object organized according to three

dimensions: size (*small*, *large*), color (*white*, *black*), and shape (*circle*, *star*). The number of words in the utterance was just sufficient to uniquely refer to the topic, hereby relying on Grice's (1975) cooperation principles. The number and nature of the objects for each scene, and the choice of topic were randomly generated for each participant. The mapping between the words and the visual features was randomly determined at the beginning of the experiment.
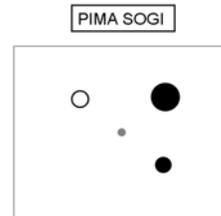

Figure 1: display of the visual scene with four objects and a two-word utterance

Results

The training performance was calculated as the proportion of trials in which the participant's guess was correct. The feedback condition showed a higher performance ($M = .67$, $SD = .18$) than the no-feedback condition ($M = .32$, $SD = .08$), $F(1, 35) = 49.16$, $MSE = .02$, $p < .001$, $\eta_p^2 = .58$. In other words, feedback increased the training performance.

The testing performance was calculated as the proportion of correct answers. The feedback condition showed a higher performance ($M = .70$, $SD = .36$) than the no-feedback condition ($M = .24$, $SD = .22$), $F(1, 35) = 19.68$, $MSE = .10$, $p < .001$, $\eta_p^2 = .28$.

The probability of choosing the correct word in the testing phase if the participants' guess was random was computed for both feedback and no-feedback conditions. Random behavior was modelled by assuming a Binomial distribution with probability $p = 1/6 = .17$. That probability of obtaining the experimental data under the assumption of random behavior was unlikely ($p < .001$) in the feedback condition but more likely in the no-feedback condition ($p = .10$), casting doubt whether the participants were learning words in the absence of feedback.

Discussion

The results from Experiment 1 show that feedback has a considerable impact on learning the meaning of words in a guessing game situation. For the no-feedback condition no evidence was obtained for learning.

A reason for the absence of learning in the no-feedback might be the ambiguity of the utterance-scene pairs, which has shown to be important in form-meaning mappings (Yu & Smith, 2007). A reduction of ambiguity would then increase the likelihood of learning word-meaning mappings, possibly in the no-feedback condition, and shed further light on the role of feedback in forming these mappings. This idea is pursued in Experiment 2, in which the ambiguity of the scene is reduced.

Experiment 2

The possible mappings between the word forms and their mappings can be represented by a 6 x 6 matrix with cells representing associations between form and meaning. In Experiment 1, three dimensions give rise to 66% spurious associations. By reducing the number of dimensions to two, this proportion drops to 50%. Under these more auspicious learning circumstances an increase in training performance is expected in both conditions.

Method

Participants 43 students from the University of Memphis participated for course credit.

Materials The material was generated in the same way as in Experiment 1, except that the visual features belonged to two rather than three dimensions: size (*small*, *medium*, *large*), and color (*white*, *grey*, *black*). The same six-word vocabulary had to be learned by the participants. Due to the change of dimensionality, the maximum number of objects in a scene now increased to nine.

Procedure The procedure was identical to the one used in Experiment 1. The participants were either assigned to the feedback ($n = 25$) or the no-feedback ($n = 18$) condition. An initial recording error reduced the sample size in feedback ($n = 23$) and no-feedback ($n = 17$) conditions during testing.

Results

As in Experiment 1, the feedback condition showed a higher training performance ($M = .75$, $SD = .15$) than the no-feedback condition ($M = .36$, $SD = .20$), $F(1, 41) = 55.33$, $MSE = .03$, $p < .001$, $\eta_p^2 = .57$. Moreover, the feedback condition showed a higher testing performance ($M = .74$, $SD = .29$) than the no-feedback condition ($M = .35$, $SD = .29$), $F(1, 38) = 16.99$, $MSE = .09$, $p < .001$, $\eta_p^2 = .31$.

A 2 x 2 ANOVA was conducted on the training performance of the last third of the trials with the feedback condition as a first factor, and the experiment (Experiment 1 vs. Experiment 2) as a second factor. Experiment 2 ($M = .59$, $SD = .26$) showed a slightly higher performance than Experiment 1 ($M = .53$, $SD = .22$), $F(1, 75) = 5.01$, $MSE = .05$, $p = .03$, $\eta_p^2 = .06$. The interaction was not significant. The Experiment effect was neither significant for testing performance, nor for training performance, if the full range of trials was considered.

However, unlike the test results in Experiment 1, the probability of test performance on the basis of a Binomial distribution model (as chance) was $p < .001$ both in feedback and no-feedback conditions.

Discussion

Overall, a dimensionality reduction from three to two dimensions shows only a small benefit in training performance, but one that seems promising. Subsequent research will need to look into other manipulations to exhibit a clearer impact of referential ambiguity. Finally, consistent with the increase in training performance, it appears that the behavior observed in the no-feedback condition cannot be attributed to random guessing. In a nutshell, although these effects are not as strong as expected, they show that learning improves when ambiguity is reduced. The evidence is therefore consistent with the idea that referential ambiguity should be part of the cognitive model of word learning and feedback. In such a model, feedback is seen as a source of information which decreases ambiguity. Experiment 3 aims at testing the alternative hypothesis that feedback acts as a motivational factor rather than as a source of information.

Experiment 3

There is a possibility that participants in the no-feedback condition became unmotivated in their task of solely guessing form-meaning mappings. Feedback may serve as a motivational function, encouraging the learner to stay vigilant about the mappings. We therefore included what may be coyly termed a "bless-your-heart" condition, in which positive feedback is provided randomly regardless of performance (akin to the "bless-your-heart" encouragements sometimes present in social environments). This condition is referred to as motivational because it provides intermittent positive reinforcement. Furthermore, we augmented the number of scenes to maximize the likelihood of observing occurrences of learning in the no-feedback condition. Indeed, it is hypothesized that likelihood of success during training will increase over time in both conditions, with feedback showing a faster increase.

Method

Participants 117 University of Memphis students participated for course credit.

Materials Four lists were created with predetermined sequences of scenes, randomizing the type and number of objects in the scene, as well as the nature of the utterance topic. These lists had the same characteristics as the sequences generated in Experiment 2, except that the lists were longer (216 trials instead of 140 trials). The four lists contained the same scenes but in a different random order. This feature allowed increasing the control of the stimuli and reducing the inter-individual variability.

Procedure The participants were either assigned to the feedback ($n = 38$), the no-feedback ($n = 39$) or the motivational ($n = 40$) condition. The motivational condition consisted in randomly displaying the statement *good job* in half of the trials, independent of the answer given.

Results

The results are summarized in Fig. 2. Trials are grouped in eight clusters of 27 trials each. The feedback factor had a significant impact on the training performance, $F(2, 114) = 134.73$, $MSE = .17$, $p < .001$, $\eta_p^2 = .74$. In order to identify

the origin of the difference between the three conditions, multiple comparisons were realized with Tukey's HSD test. Because one of the lists did not record all the participants' responses until the end of the sequence due to a technical problem, analyses were limited to the seventh cluster, after learning had in principle occurred. As expected, the training performance in the feedback condition ($M = .93$, $SD = .14$) was higher than in the no-feedback ($M = .39$, $SD = .36$, *Tukey's-HSD* $= .54$, $SE = .07$, $p < .001$) and motivational conditions ($M = .39$, $SD = .31$, *Tukey's-HSD* $= .55$, $SE = .07$, $p < .001$). The no-feedback and motivation conditions exhibited no difference (*Tukey's-HSD* $= -.01$, $SE = .06$, $p = .99$). Parallel results were found for the testing performance.
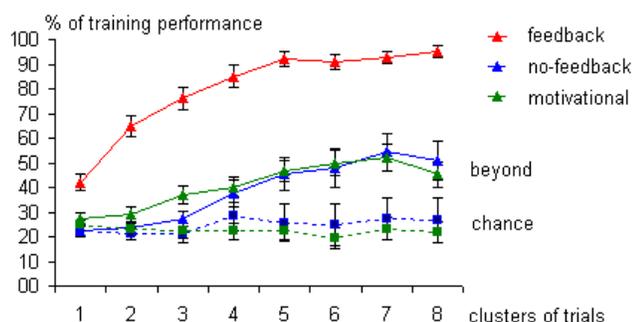

Figure 2: training performance across trials in Experiment 3

The position in the sequence of trials was significant, $F(1, 114) = 46.120$, $MSE = 6.02$, $p < .001$, $\eta_p^2 = .29$, showing that training performance increased from the beginning to the end of the sequence. Moreover, there was a Feedback x Trial interaction, $F(2, 114) = 10.72$, $MSE = .13$, $p < .001$, $\eta_p^2 = .16$, showing that the effect of Trial was different in function of the feedback condition.

In order to determine the conditions in which performance increased, three separate repeated-measure analyses were conducted with Trial as a within-subject factor. These analyses showed that performance increased for all conditions, although it was marginally significant in the motivational condition: feedback ($F(1, 37) = 73.77$, $MSE = .10$, $p < .001$, $\eta_p^2 = .67$), no-feedback ($F(1, 38) = 6.85$, $MSE = .17$, $p = .013$, $\eta_p^2 = .15$), and motivational condition ($F(1, 39) = 3.89$, $MSE = .13$, $p = .056$, $\eta_p^2 = .09$).

In order to track down the existence of learning in the no-feedback and motivational conditions, a subsequent analysis of variance was conducted with training as a repeated measure, and two factors: feedback condition (no-feedback vs. motivational) and whether or not the individuals had a score above chance in the subsequent test. The feedback condition was not included in that analysis because all the individuals in that condition were above chance level at testing. As shown in Fig. 2, individuals above the testing chance level, showed a linear progression of performance at training, while the performance is flat for those at chance level. The effect of being above the testing chance level on the training performance was significant in the no-feedback condition, $F(1, 75) = 23.21$, $MSE = .13$, $p < .001$, $\eta_p^2 = .24$, as well as its interaction with trial, $F(1, 75) = 7.19$, $MSE =$

$.14$, $p = .009$, $\eta_p^2 = .09$.

Discussion

The results again show that feedback considerably increases the training performance compared to a situation without feedback. The no-feedback and motivational conditions could not be reliably differentiated. These findings are consistent with the idea that feedback is a source of information responsible for decreasing ambiguity in the appreciation of form-meaning mappings. An equally consistent account is that the motivational condition is not properly motivational. The proper role of motivation should be examined in further research. Finally, the fact that a subset of individuals show learning gains during training without feedback clearly illustrates that it is possible to keep track of some form-meaning mappings in these circumstances. For simplicity, it is hypothesized that a common learning mechanism relying on statistical regularities found in the environment is active in all conditions. In that framework feedback can be considered as a way to focus attention on these statistical regularities (as argued by Goldstein & Schwade, 2008). More formally, feedback can be seen as a way of reducing the number of spurious correlations between form and meaning. This idea is tested in computational simulations presented below.

Simulations

It has been shown that the environment contains a significant source of cues that facilitate the learning of language (Regier, 1996). By forming a model of the data, the learner can compare the expectations which derive from the model to the incoming data, in order to modify the model, and hence learn the language. One way of representing the data is to track word and feature frequencies across contexts. For instance, the fact that the word *cat* is more frequently associated with cats than with dogs is a source of indirect positive evidence about the correctness of the association *cat*-cat, and of indirect negative evidence about the correctness of the association *cat*-dog.

A variety of computational models have shown word learning without receiving explicit negative evidence (e.g., Dienes, Altmann, & Goa, 1999; Frank, Goodman & Tenenbaum, 2008; Regier, 1996; Siskind, 1996). These techniques have in common the idea that the discovery of word meaning results from cross-situational learning. By keeping track of the frequency of the various possible mappings between sound and vision across situations, the language learner is able to map word forms to their references.

Method

The model implemented a 6 x 6 matrix representing the co-occurrence of word forms and meaning found in the scene. For each scene, the frequency of co-occurrence was incremented between any word form found in the utterance and any visual feature on the display.

The learning mechanism was identical in both feedback and no-feedback conditions. The difference between the two conditions was in the features being considered. That is, without feedback all objects in the scenes had to be considered as potentially referred to by the utterance, whereas in the feedback condition only the features of the right answers needed to be looked at, thereby reducing the number of potential spurious correlations.

For each scene, the model made a decision through a constraints-satisfaction procedure based on the current state of the matrix. The hypothesized topic of the utterance was chosen as the first object in the scene which satisfied the best combination of the strongest mappings between forms and meanings.

The input given to each instantiation of the model was the same as the one given to each participant in Experiments 1-3. Simulations 1-3 are therefore named after their human data counterparts, Experiments 1-3. In order to have a baseline comparison, random data (appearing in green in Figure 3) were generated for each scene (1 for success; 0 for failure), following a Bernouilli distribution with $p = 1 / n$, where n was the number of objects in the scene.

To simulate the individual differences between participants, the trials during which the matrix started to get updated was randomly distributed. This was aimed at reflecting the point at which participants began to pay attention to the form-meaning mapping possibilities. This "Eureka moment" followed a log-normal density with $\mu = 3$ and $\sigma^2 = 1$, manually chosen to maximize the fit between the human data from all experiments and conditions and their simulated counterparts.

Results

Overall, the fit between the data and the outcome of the simulation was acceptable. On Experiments 1 and 2 alone, the fit was significant with $r = .55$, $p < .001$, $R^2 = .31$. The simulation (left panel of Fig. 3) was consistent with the experimental results (right panel of Fig. 3) in two respects. First, feedback had a higher performance than no-feedback in all simulations: Simulation 1 ($F(1, 35) = 49.33$, $MSE =$ .03, $p < .001$, $\eta_p^2 = .59$), Simulation 2 ($F(1, 39) = 13.15$, $MSE = .02$, $p = .001$, $\eta_p^2 = .26$), and Simulation 3 ($F(1, 116) = 31.71$, $MSE = .02$, $p < .001$, $\eta_p^2 = .36$. By the nature of the computational model, the motivational and no-feedback conditions were set to be identical, so their comparison does not constitute an interesting aspect of the analysis. Second, paired-sample t-tests comparing the model performance to the baseline confirmed that the models learned in all simulations except in the no-feedback condition of Simulation 1.

Discussion

The simulation shows that part of the variability between the feedback and no-feedback conditions can be explained by a simple model portraying feedback as a way of reducing the number of meaningful features under consideration. The model, however, does not account well for the large inter-individual variability, simply because little is known about what determines these individual differences. Likely candidates are differences in memory and attention in performing the task, but the exact details require further research.

General Discussion

This study provides support for the thesis that feedback helps artificial word learning by reducing the ambiguity in the scenes to be processed. This idea is consistent with the proposition that feedback is important to language learning because it helps learner to focus their attention on statistical regularities found in the environment (Goldstein & Schwade, 2008).

In Experiment 1, feedback induced a clear increase in learning. The no-feedback actually did not show sign of learning at all. The same conclusion could be drawn from the simulation.

In Experiment 2, the vocabulary to be learned was organized according to a simpler bidimensional structure, supposed to reduce the ambiguity of the learning situations.
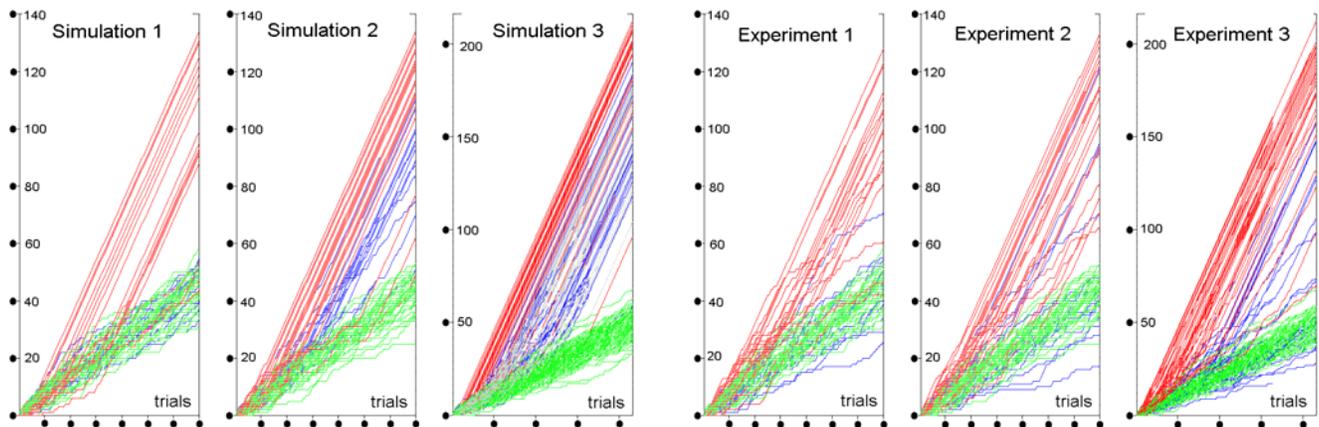

Figure 3. Cumulative training performance (y-axis) in function of the trial (x-axis) for Simulations 1-3 and Experiments 1-3, distinguishing feedback (red) and no-feedback (blue) condition from random behavior (green).

Whereas Experiment 2 showed a small improvement in training, Simulation 2 exhibited a clear-cut improvement in the no-feedback condition. These observations converge towards the idea that the problem of ambiguity is a key issue in learning the words in that situation, although individual differences are likely responsible for a great deal of variance in performance.

With its longer sequence of trials, Experiment 3 illustrated that learning can successfully occur without feedback (both in the human data and in the simulation). Feedback, however, kept having a clear advantage in word learning. The fact that the no-feedback and motivational conditions displayed equivalent performances is consistent with the idea that feedback supports learning by its informational nature, and not simply by encouraging the participants to learn playing the game.

Overall, this research exhibits three features. First, feedback serves as a powerful cue during word learning. Second, under specific conditions, learning occurs without such feedback. Third, feedback helps learning by reducing the amount of ambiguity in the environment. We cannot argue that individuals adopt the same strategy in feedback and no-feedback conditions. Rather, it may be that the reduction of ambiguity is an important factor in explaining the performance in feedback and no-feedback conditions.

Acknowledgments

References

Anderson, S. R., & Lightfoot, D. (2000). The human language faculty as an organ. *Annual Review of Physiology*, *62*, 697-722.

Brown, R., & Hanlon, C. (1970). Derivational complexity and the order of acquisition in child speech. In J. R. Hayes (Ed.), *Cognition and the development of language* (pp. 11-54). New York: Wiley.

Chouinard, M., & Clark, E. V. (2003). Adult reformulations of child errors as negative evidence. *Journal of Child Language*, *30*, 637-669.

Creel, S. C., Aslin, R. N., & Tanenhaus, M. K. (2006). Acquiring an artificial lexicon: Segment type and order information in early lexical entries. *Journal of Memory and Language*, *54*, 1-19.

Dale, R., & Christiansen, M. (2004). Active and passive statistical learning: Exploring the role of feedback in artificial grammar learning and language. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual meeting of the cognitive science society* (pp. 262-267). Mahwah, NJ: Lawrence Erlbaum Associates.

Dienes, Z., Altmann, G. T. M., Gao, S.-J. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science*, *23*, 53-82.

Frank, M. C., Goodman, N. D., & Tenenbaum, J. (2008). A Bayesian framework for cross-situational word-learning. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in Neural Information Processing Systems* (pp. 457-464). Cambridge, MA: MIT Press.

Goldstein, M. H., & Schwade, J. A. (2008). Social feedback to infants' babbling facilitates Rapid Phonological Learning. *Psychological Science*, *19*, 515-523.

Grice, H. P. (1975). Logic and conversation. In P. Cole, & J. Morgan (Eds.), *Syntax and Semantics*, (pp. 41-58). New York: Academic Press.

Kersten, A. W., & Earles, J. L. (2001). Less really is more for adults learning a miniature artificial language. *Journal of Memory and Language*, *44*, 250-273.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211-240.

Markman, E. M. (1991). The whole object, taxonomic, and mutual exclusivity assumptions as initial constraints on word meanings. In J. P. Byrnes & S. A. Gelman (Eds.), *Perspectives on language and cognition: Interrelations in development* (pp. 72-106). Cambridge: Cambridge University Press.

Mitchell, T. (1997). *Machine learning*. McGraw Hill.

Pashler, H., Cepeda, N. J., & Wixted, J. T. (2005). When does feedback facilitate learning of words? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 3-8.

Regier, T. (1996). *The human semantic potential*. Cambridge, MA: MIT Press.

Schooler, L. J. & Anderson, J. R. (1990). The disruptive potential of immediate feedback. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 702-708, Hillsdale, NJ: Erlbaum.

Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, *61*, 39-61.

Smith, L. B., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*, 1558-1568.

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*, 1289-1309.

Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, *18*, 414-420.