Title
Cross-Situational Statistical Learning: Implicit or Intentional?

Authors
Kachergis, George
Yu, Chen
Shiffrin, Richard

Cross-Situational Statistical Learning: Implicit or Intentional?

George Kachergis, Chen Yu, and Richard M. Shiffrin
{gkacherg, chenyu, shiffrin}@indiana.edu
Department of Psychological & Brain Science / Cognitive Science Program
Bloomington, IN 47405 USA

Abstract

For decades, implicit learning researchers have examined a variety of cognitive tasks in which humans seem to automatically extract structure from the environment. Similarly, statistical learning studies have shown that humans can use repeated co-occurrence of words and referents to build lexicons from individually ambiguous experiences (Yu & Smith, 2007). In light of this, the goal of the present paper is to investigate whether adult cross-situational learners require an explicit effort to learn word-object mappings, or if it may take place incidentally, requiring merely attention to the audiovisual stimuli. In two implicit learning experiments with incidental tasks that direct participants' attention to different aspects of the stimuli, we found evidence of learning, suggesting that cross-situational learning mechanisms can be incidental without explicit intention. However, learning was superior under explicit study instructions, indicating that strategic inference may also play a role.

Keywords: implicit learning; language acquisition; cross-situational statistical learning; automaticity

Introduction

Humans have a remarkable capacity to adapt to the regularities in our environment, and our everyday actions—from navigating a room to navigating a conversation—are evidence of our learned skills. Often, we adapt without overt effort or even awareness of the regularity or of our changing behavior. Dubbed *implicit learning* (Reber, 1967), this automatic adjustment to the world is typically studied in cognitive experiments using grammaticality judgments or reaction times to stimuli generated by finite state grammars (see Shanks, 2005 for a review).

The burgeoning statistical learning literature has motivations and predictions that significantly overlap with those of the implicit learning literature, as discussed by Perruchet and Pacton (2006). The seminal work on statistical learning (Saffran, Aslin, & Newport, 1996) demonstrated that infants are sensitive to statistical regularities in a continuous stream of an audible artificial language, enabling them to distinguish probable syllable sequences (i.e., words) from improbable syllable sequences. Newport and Aslin (2004) found that infants are also sensitive to temporally distal regularities, which weighs in favor of a more general statistical learning mechanism, rather than a simple mechanism for associating adjacent sounds. Other studies have found that infants can acquire nouns via the repeated co-occurrence of words and their referents across situations containing multiple words and objects, which are thus separately ambiguous (e.g., Smith & Yu, 2008).

As in adult studies of implicit learning, infant statistical learning studies present participants with structured training data but no explicit learning instructions, and find behavioral differences due to the statistical regularities in the training data. Inspired by this, our aim here is to empirically investigate the automaticity of cross-situational statistical word learning in adults, who are typically given explicit instructions to learn the meaning of the words (e.g., Yu & Smith, 2007). In Experiment 1, we presented participants with a set of spoken words and visual objects with one-to-one mappings between them, but framed the task as one of recognition memory for individual stimuli, and not as one of learning word-object mappings. We then gave participants a surprise test: for each of 54 word-object pairings, they were asked to indicate how often the word and object co-occurred. With their attention focused on memorizing individual words or visual objects, would participants unintentionally learn which words and objects co-occurred more frequently? In Experiment 2, we used a signal detection task as another incidental task to direct participants' attention to both auditory and visual streams, but again with no explicit instructions to learn word-object mappings. After that, we gave them a surprise test to assess their knowledge of word-object mappings. In both experiments, after the initial implicit learning blocks, as a measure of their statistical learning capability (to compare with implicit learning), participants also completed blocks in which they were explicitly instructed to either count word-object co-occurrences, or simply to learn the meaning of the words.

The organization of the paper is as follows: we first introduce the cross-situational learning paradigm, and then discuss the possible learning mechanisms and potential contributions of the present implicit learning studies to advance our understanding of statistical learning. We then present two implicit learning experiments and their results. Finally, we conclude by summarizing the results from the two studies and discussing the connection between statistical and implicit learning.

Cross-Situational Statistical Learning

In a typical version of cross-situational learning, adults are asked to learn which word goes with each object, and are then shown a series of training trials, each of which contains four objects (e.g., a sculpture) and four spoken pseudowords (e.g., "manu"). Because correct word-referent pairings are not indicated, learners can utilize only the repeated co-occurrence of words with their intended referents to learn across many trials. In a typical learning scenario (e.g., in Yu & Smith, 2007), participants attempted

to learn 18 pseudoword-object pairings from 27 12-second trials. This design allowed each stimulus (and hence each correct word-referent pairing) to be presented six times. In one form or another, the learning of a pairing involves the accumulation of word-object co-occurrence statistics across the training trials. Participants acquired, on average, nine of the 18 pairs, as measured by a 4-alternative forced choice (4AFC) referent test for each word.

When each trial contains 16 possible word-referent associations, how might learning proceed? There are at least two distinct approaches that learners may apply. First, an ideal associative learner may maintain a word x object co-occurrence matrix M, incrementing the count in cell $M_{w,o}$ whenever word w and object o appear together in a trial. Table 1 shows such a matrix, which represents the training statistics used in the present study. At test, such a learner may choose the most frequently co-occurring referent for each word. Associative models typically approximate this co-occurrence matrix by strengthening a randomly sampled (perhaps according to current association strengths) subset of pairings on each trial. The association of spatiotemporally proximal stimuli could be carried out by automatic processes that require neither strategy nor intent to learn. Modern memory models such as REM (Shiffrin & Steyvers, 1997) even predict such associations by allowing feature values of nearby items to accidentally be recorded in an item's trace.

Another plausible learning approach is implemented in rule- and inference-based models (e.g., Siskind, 1996), which propose and store a number of hypothesized word-object pairings on each trial. Proposals may be made with respect to constraints such as mutual exclusivity, and hypothesized pairings may be confirmed if consistent evidence is presented later or removed from the lexicon if contradictory evidence is observed. This type of learning is more in accord with a deliberative, strategic learning process. If cross-situational learning is largely automatic, one may expect participants to have some knowledge of which words and objects frequently co-occurred during training, even when they were not explicitly trying to learn these relations. On the other hand, if cross-situational learning relies on more strategic, intentional inferences, then participants may perform much worse in such an incidental learning condition. Thus, the results from incidental learning tasks may shed light on the underlying learning mechanisms that learners use.

In particular, the present study will test participants' knowledge not only of the correct pairings (i.e., the diagonal cells of Table 1) as is typically done, but also of the spurious word-object co-occurrences (non-diagonal cells) that appear during training— the sort of detailed and partial information that is stored by associative models (or an ideal learner), but typically not by rule-based models. We do this by asking participants to rate the strengths of co-occurring word-object pairings for both correct and incorrect pairings.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	6	2	0	1	0	0	1	1	0	0	1	1	3	2	1	4	0	1
2	2	6	1	1	1	0	1	0	0	2	0	0	2	1	2	1	2	3
3	0	1	6	1	2	2	0	0	2	2	1	0	2	1	1	0	2	1
4	1	1	1	6	0	0	0	2	1	1	1	4	0	0	4	0	0	2
5	0	0	2	0	6	1	3	0	1	0	1	1	1	4	0	1	1	2
6	0	1	2	0	1	6	0	1	3	2	0	1	2	0	1	1	3	0
7	1	0	0	0	3	0	6	2	0	2	4	1	0	2	0	1	0	2
8	1	0	0	2	0	1	2	6	0	3	2	2	1	0	2	1	1	0
9	0	2	2	1	1	3	0	0	6	1	0	0	1	1	1	1	2	2
10	0	0	2	1	0	2	2	3	1	6	4	1	1	0	0	0	1	0
11	1	0	1	1	1	0	4	2	0	4	6	2	1	0	0	1	0	0
12	1	0	0	4	1	1	1	2	0	1	2	6	0	0	3	1	1	0
13	3	2	2	0	1	2	0	1	1	1	1	0	6	1	0	2	0	1
14	2	1	1	0	4	0	2	0	1	0	0	0	1	6	0	3	1	2
15	1	2	1	4	0	1	0	2	1	0	0	3	0	0	6	0	2	1
16	4	1	0	0	1	1	1	1	1	0	1	1	2	3	0	6	1	0
17	0	2	2	0	1	3	0	1	2	1	0	1	0	1	2	1	6	1
18	1	3	1	2	2	0	2	0	2	0	0	0	1	2	1	0	1	6

Table 1: Word x Referent matrix with the co-occurrences of each word and object accumulated across the 27 training trials used in each condition in both present experiments.

The testing paradigm allows us to both access participants' knowledge of spurious pairs and to compare that with what they know about correct pairs. Previous work has found evidence that people are sensitive to how often words and objects have co-occurred—even when a single object appears with a few words with differing frequency (Vouloumanos, 2008). However, Vouloumanos presented only a single word-object pair per trial, giving participants no choice as to which pairings to attend. In contrast, our paradigm offers 16 possible pairings per trial. Thus, the presence of four concurrent objects and four successive words per trial demands that participants modulate their attention, possibly forming stronger associations between particular words and objects, or perhaps attending only a subset of possible pairings. Thus, it is unclear how well participants' co-occurrence ratings will be correlated with actual stimuli co-occurrences in the explicit conditions, since inference-based learners may only track a lexicon of the most likely pairs (i.e., high co-occurrence stimuli), rather than a full matrix of associations.

Experiment 1

Every participant went through four blocks of training and testing in a fixed order. Training and testing in block 0 was structured differently than the remaining three. Participants were told that they would see multiple objects and hear multiple words on each trial, and that they should remember each object and word because their memory will be tested at the end. After the brief training period in block 0, they were given a recognition memory test: a single stimulus (word *or* object) was presented, and they were asked to label it *old* or *new*. In block 1, participants were told again that they should remember each object and word for a subsequent memory test. However, after this training period, participants were given a surprise test of their knowledge of stimuli co-occurrences. In block 2, participants were explicitly asked to remember how many times each word and object appeared together during training. They were not

told what type of test to expect, but the co-occurrence rating test given was exactly the same as in block 1. In block 3, participants were simply asked to learn the meanings of the words—explicit learning instructions like those given in previous cross-situational word learning studies.

Subjects

Participants were 35 undergraduates at Indiana University who received course credit for participating. None had participated in other cross-situational experiments.

Stimuli

Verbal stimuli were 72 computer-generated pseudowords that are phonotactically-probable in English (e.g., "bosa"), and were spoken by a monotone, synthetic female voice. Objects were 72 photos of uncommon, difficult-to-name objects (e.g., strange sculptures). Of these sets of objects and words, 54 were randomly assigned to three sets of 18 word-object pairings; one set for each study condition. The remaining 18 words and 18 objects were used for an initial recognition memory test.

In block 0, each trial presented three unusual objects concurrently and three pseudowords heard in succession. Block 0's training consisted of only three 11-second trials, displaying in total nine unique words and objects. After these trials, participants were given a recognition test for each trained word and object, as well as nine new words and objects. On each test trial, a single stimulus (word *or* object) was presented, and participants were asked to indicate if it was *old* or *new*.

In blocks 1-3, each training trial consisted of a display of four objects and four pseudowords were played in succession, and 27 such trials were in each block. Each training trial began with the appearance of four objects, which remained visible for the entire trial. After 2 seconds of initial silence, each word was heard (randomly ordered, duration of one second) followed by two additional seconds of silence, for a total duration of 14 s per trial.

After each training period, participants were tested for knowledge of stimuli co-occurrences. One word and one object were presented on each trial, and participants were asked to indicate how many times [0-6] the given word-object pairing had appeared during training. Each of the 18 words and objects appeared in three test trials, for a total of 54 randomly-ordered trials. The correct (6-co-occurrence) pairings comprised 18 of the test trials (Table 1's diagonal). The remaining 32 trials tested cells in the matrix with 0 (14 trials), 1 (14), 2 (12), 3 (8), and 4 (6) co-occurrences.

Procedure

Condition order was fixed, and each participant took part in all four blocks. Block 0 was a three trial training period with three words and objects per trial, followed by a recognition test of every individual stimulus presented, and nine new words and objects. In block 1, participants were instructed to study individual stimuli for a memory test. However, following the 27 training trials, participants were instead asked to indicate how many times [0-6] each of 54 specific word-object pairings appeared during training. In block 2,

participants were asked to track how often each word co-occurred with each object. After the 27 training trials—which had the same co-occurrence statistics as in block 1, albeit different stimuli—participants were again given test trials asking them to rate the same 54 pairings. Finally, in block 3 participants were simply instructed to learn the meanings of the words, given cross-situational training (statistically identical to blocks 1 and 2), and again tested on the same 54 pairings.

Results & Discussion

In block 0, participants recognized a mean of 96% of the objects and 90% of the words, with a low false alarm rate (8%). In both word and object recognition, every participant was at least 77% accurate. It is notable that memory is imperfect for the stimuli, since many models of cross-situational learning assume that learners can absolutely identify each stimulus, which is evidently not the case.

To determine how related participants' co-occurrence ratings were to the actual number of times the tested word-object pairings actually appeared together during training, Kendall's rank correlation coefficient (tau) was calculated for each participant's 54 test trials in each condition. The mean tau values for each condition are shown in Figure 1. In block 1, when participants were studying individual words and objects (but not attending to co-occurrences), their responses in the surprise rating task showed a small but significant positive correlation with the actual number of times the presented pairings co-occurred during training ($M = .04$, one-sided $t(34) = 1.90$, p<.05). In comparison, in the explicit learning conditions in blocks 2 and 3, when participants were respectively told to track all word-object co-occurrences and to learn the meaning of the words, their ratings were significantly more positively correlated than in block 1 (block 2 $M = .15$, paired $t(34) = 3.82$, $p<.001$; block 3 $M = .17$, paired $t(34) = 3.86$, $p<.001$). Moreover, the strength of correlations in the two explicit conditions is not significantly different (paired $t(34) = 0.66$, p>.05).

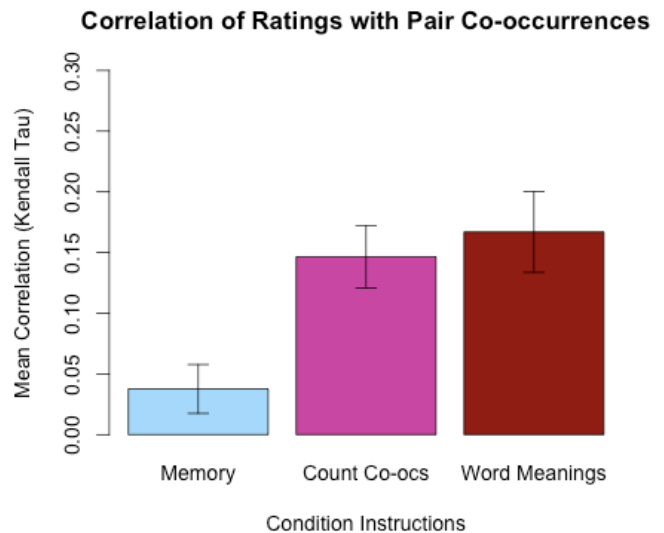**Correlation of Ratings with Pair Co-occurrences**

Figure 1: Mean correlation of participants' responses with the actual pair co-occurrences in Exp. 1. Error bars: +/-SE.

Positive correlations between ratings and a broad sample of the actual co-occurrence statistics from training indicate that participants are sensitive to arbitrary stimuli co-occurrences when explicitly told to attend to such correspondences. However, one could imagine that the positive correlations could be due largely to knowledge of some particular subsets of the co-occurrences: e.g., perhaps learners are sensitive to words and objects that never co-occurred, and thus rated these pairings very low, and all others high. To examine performance in more detail, we calculated each participant's d-prime (d')[1] for the most extreme pairings tested in each condition: stimuli that co-occurred 0 or 6 times. Positive d' shows sensitivity resulting from a high hit rate and low false alarm rate. As shown in Figure 2, participants only had significant sensitivity for 6-co-occurrence ('correct') pairings in the explicit learning conditions (count co-occurrences $M = 0.64$, one-sided $t(34) = 4.92$, $p<.001$; word meanings $M = 0.81$, one-sided $t(34) = 5.08$, $p<.001$).

Two patterns from this study are noteworthy. First, based on both d' analysis and correlation measures, the learning that results from the counting co-occurrences condition and the word learning condition were similar. Although not conclusive, this may suggest that participants in the word learning condition may have used an associative learning strategy based on counting word-object co-occurrences.
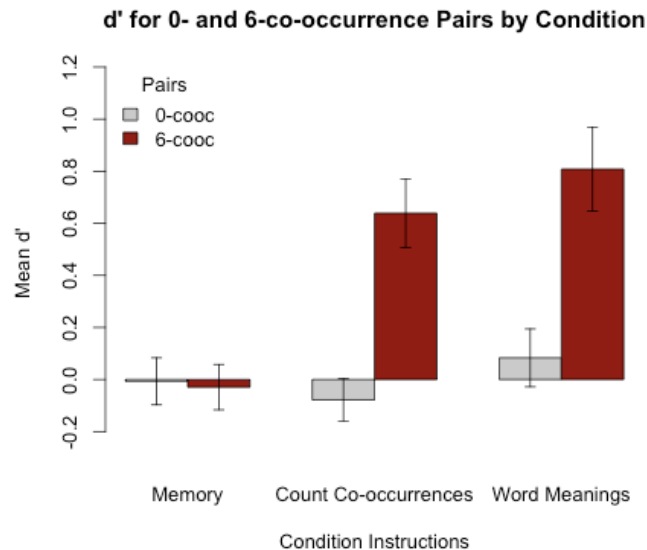

Figure 2: Mean d' for 0- and 6-co-occurrence word-object pairings in Experiment 1, by condition. Error bars are +/-SE.

Second, despite having good recognition memory for individual words and objects that were presented during training, word learning was very poor in the implicit learning condition, as measured both by correlation of their responses with actual pair co-occurrences, and by d' for

---

[1] For example, hits for 0-co-occurrence pairings are responses of 0, and false alarms are responses of 0 for pairings that co-occurred more than never. $d' = Z(p(\text{hit})) - Z(p(\text{false alarm}))$, where Z is the inverse of the cumulative Gaussian distribution.

correct pairings and stimuli pairings that never co-occurred. Nonetheless, the positive correlations found in every condition—although smaller in the implicit condition—show that participants do, on average, absorb some of the stimulus co-occurrences in all conditions. However, this sensitivity is not enough to support implicit word learning in our study, as much stronger correlations are shown when learners are instructed to count co-occurrences or learn word meanings. Under these instructions, participants become sensitive to words and objects that frequently co-occur.

Experiment 2 investigates whether a different incidental task, which may direct attention to word-object co-occurrences, rather than the stimuli themselves, may yield automatic word learning.

Experiment 2

Experiment 1 showed that an incidental memory task results in some implicit knowledge of word-referent co-occurrences, but that explicit instructions to learn word-object co-occurrence or to learn word meanings resulted in much greater knowledge. In Experiment 2, we use a different task in the implicit learning condition: instead of asking participants to remember individual stimuli for a later memory task, we give participants a signal detection task to carry out during training. This task—detecting visual noise added to visual objects, and louder auditory stimuli (words)—directed participants to pay attention to both visual and auditory stimuli simultaneously, but gave no directions to engage in learning of word-object pairings.

Subjects

37 undergraduates at Indiana University received course credit for participating. None had participated in previous cross-situational experiments.

Stimuli

The sets of pseudowords and referents for Experiment 2 were identical to those used in Experiment 1. Training trials were the same as those used in Experiment 1, and had the same co-occurrence statistics (shown in Table 1). However, on each training trial in blocks 1 and 2, a random number [0-4] of the words were louder than others, and Gaussian pixel noise was momentarily added to a single object *during a word presentation* a random number of times [0-4] each trial. Thus, for 6.3% of audio stimulus presentations during training, that word would be loud and *one* of the objects would simultaneously have noise added, highlighting a pairing—but only the correct pairing in 25% of these cases.

Procedure

In block 1, participants were told that they would be presented with artificial words and objects on a series of slides, on which some words would be louder than the others and some objects would have multicolored speckles (noise). Their task was to quickly press the mouse button each time a loud word or noisy object was presented. However, after the 27 training trials, participants were given a surprise test, and asked to indicate how many times [0-6]

each of 54 specific word-object pairings appeared during training. In block 2, participants were asked to track how often each word co-occurred with each object, and were also told to do the same signal detection task during training. Instructions for block 3 asked participants to track word-object co-occurrences without doing the signal detection task, and in block 4 participants were simply told to learn the meanings of the words. The same 54 rating test trials of specific pairings followed the training periods of blocks 2, 3, and 4, though with different stimuli for each block.

Results & Discussion

Experiment 2 used a signal detection (SD) task that required participants to attend to both auditory and visual stimuli, but did not mention that they would need to remember the stimuli later. However, as in Experiment 1, after this first training block participants were given a surprise test for incidental learning. In successive learning conditions, participants were instructed to do both the SD task and to count word-object co-occurrences (SD+CC), to count co-occurrences (with no other task; CC), and finally, to simply learn the meanings of the words (Word Meanings). As in Experiment 1, Kendall's tau was calculated for each participant's 54 test trials in each condition to measure how related their ratings were to the actual number of word-object co-occurrences. As shown in Figure 3, although the SD task resulted in significantly positive correlations ($M =$ .10, one-sided $t(36) = 3.75$, $p<.001$), the explicit learning conditions showed significantly more correlated responses (CC $M = .21$, paired $t(36) = 3.57$, $p<.01$; SD+CC $M = .25$, paired $t(36) = 5.12$, $p<.001$; Word Meanings $M = .29$, paired $t(36) = 4.97$, $p<.001$). Thus, as found in Experiment 1, participants show sensitivity to stimuli co-occurrences in every condition, but greater sensitivity in the explicit learning conditions than in the implicit learning condition.

As in Experiment 1, we calculated d′ for maximal and minimal co-occurrence pairings by condition to gain insight into the kind of pairings to which participants in Experiment 2 were sensitive. As shown in Figure 4, participants had significant sensitivity for 6-co-occurrence pairings in the implicit learning condition (SD $M = .19$, one-sided $t(36) = 1.81$, $p<.05$) as well as the explicit conditions, but showed significantly greater sensitivity in the explicit conditions (CC $M = .61$, paired $t(36) = 2.97$, $p<.01$; SD+CC $M = .61$, paired $t(36) = 3.44$, $p=.001$; word meanings $M = .81$, paired $t(36) = 3.58$, $p<.001$). In the explicit conditions, d′ for 0-co-occurrence pairings was significantly positive (CC $M = .49$, one-sided $t(36) = 1.50$, $p=.07$ (marginal); SD+CC $M = .32$, one-sided $t(36) = 2.66$, $p<.01$; word meanings $M = .30$, one-sided $t(36) = 2.20$, $p<.05$), but not in the implicit condition (SD $M = .07$, one-sided $t(36) = .87$, $p=.19$). Thus, although participants given SD instructions did show some implicit learning of 6-co-occurrence pairings, they were more sensitive to these pairings under explicit instruction.

There are a few intriguing results from this experiment. First, performance in the SD+CC condition was at least as good as CC alone. Thus, participants could handle the two tasks concurrently without hindering performance. We suspected that the signal detection task might encourage participants to attend to both auditory and visual streams simultaneously, perhaps increasing storage of cross-modal associations. Possibly as a result of this focus, in contrast to Exp. 1, participants in Exp. 2 showed significant sensitivity to 0-co-occurrence pairings in the explicit conditions.

Second, word-learning instructions yielded performance as high as found in other explicit instructions (SD+CC and CC). This confirmed our finding from Experiment 1: both counting co-occurrences—as an ideal associative learner might do—and attempting to learn words result in similar performance in humans, both for correct pairs and for spurious co-occurrences.
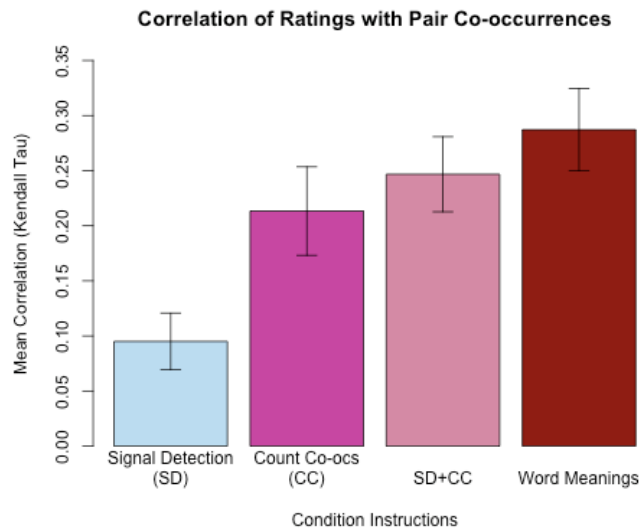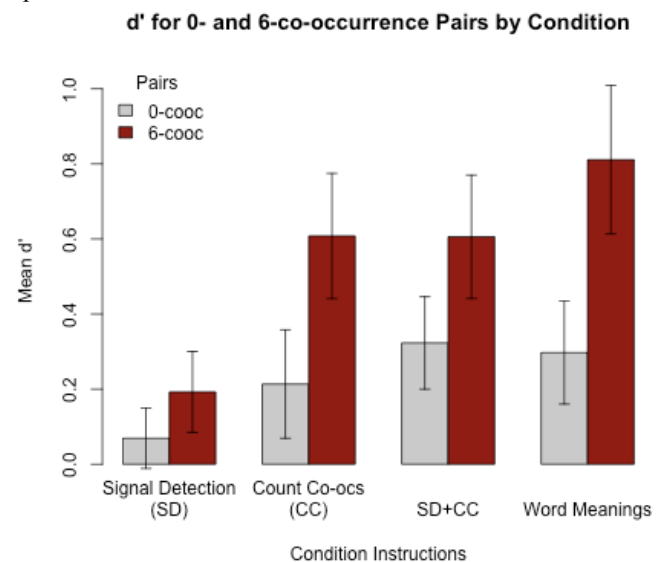

Figure 3: Mean rank correlation of each participant's responses with the actual number of pairing co-occurrences in Experiment 2. Error bars show +/-SE.

Figure 4: Mean d′ for 0- and 6-co-occurrence word-object pairings in Experiment 2, by condition. Error bars are +/-SE.

General Discussion

Implicit learning and statistical learning both describe an agent's adaptation to regularities in its environment. We set out to determine whether cross-situational word learning can be accomplished by mere exposure to the same type of training used in intentional settings. In Experiment 1's implicit learning condition, participants attempted to remember individual stimuli. In a surprise test of knowledge for word-object co-occurrences, participants' ratings were correlated with the actual number of co-occurrences, meaning that learners had acquired a rough approximation of the real-world statistics, much like associative models predict. However, a signal detection analysis showed no sensitivity to correct word pairings. Moreover, in subsequent explicit conditions, participants showed stronger correlations, as well as sensitivity to correct pairings. Using a signal detection task rather than a memory task in the first block, thus encouraging concurrent attention to both words and objects, Exp. 2 asked again whether participants acquire cross-situational co-occurrence statistics automatically. Participants demonstrated some implicit knowledge as in Exp. 1, but also showed some sensitivity for correct word-referent pairs. However, in explicit conditions participants showed greater sensitivity to such frequently co-occurring stimuli, as well as significant knowledge of spurious co-occurrences. Furthermore, we found that participants' learning when instructed to count co-occurrences looks similar to learning under instructions to merely learn words, which we speculate may mean that participants utilize a similar strategy in both conditions. By asking participants to perform slightly different tasks with the same input and then comparing their resulting learning, it will be possible to determine which regularities are automatically acquired and which must be explicitly attended or inferred.

What do the present results tell us about cross-situational statistical learning? They seem to contradict simple hypothesis-testing mechanisms, which would typically not maintain information about spurious co-occurrences, and which may not operate automatically. However, the results also contradict a strong associative account: learning was greater in explicit conditions than in implicit conditions, suggesting that learning may be in part strategic, or at least modulated by attention. Thus, we may say that cross-situational statistical word learning is neither wholly implicit, nor wholly explicit: some statistics are acquired automatically, and the learning system indubitably uses this information during explicit study, as well. Moreover, the fact that the explicit conditions always produced greater sensitivity for the correct pairings than for pairings that never co-occurred suggests that some mechanism for highlighting stimuli that frequently co-occur is at work.

In summary, although the implicit learning we observed was inferior to the explicit learning, its presence indicates that knowledge of co-occurrence statistics can be acquired incidentally. Since implicit learning requires few resources, it can be carried out minute-by-minute, hour-by-hour and day-by-day. Hence, in the long run, cumulative implicit learning may still play an important role in human language acquisition. Overall, our work suggests that neither simple associative models that approximate ideal observers, nor hypothesis-testing models relying on explicit inferences capture both the implicit and intentional aspects of cross-situational word learning. We hope that this work will motivate researchers to consider hybrid models that include both strategic, inference-based mechanisms as well as automatic, associative ones. Finally, we believe this work represents an early step in linking the implicit learning and statistical learning literatures.

Acknowledgments

References

Kachergis, G., Yu, C., & Shiffrin, R. M. (2009). Frequency and contextual diversity effects in cross-situational word learning. In N. Taatgen, H. van Rijn, J. Nerbonne, & L. Schomaker (Eds.) *Proceedings of CogSci 31*. Austin, TX: Cognitive Science Society.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, *48*(2), 127–162.

Perruchet, P. & Pacton, S. (2006). Implicit learning and statistical learning: one phenomenon, two approaches. *Trends in Cognitive Sciences*, *10*, 233-238.

Reber, A. S. (1967). Implicit learning of artificial grammars. *J. Verbal Learn. Verbal Behav*., *6*, 855-863.

Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month old infants. *Science*, *274*, 1926-1928.

Shanks, D.R. (2005). Implicit learning. In *Handbook of Cognition* (Lamberts, K. & Goldstone, R., eds), pp. 202–220, Sage Publications.

Shiffrin, R. M. & Steyvers, M. (1997). A model for recognition memory: REM: Retrieving effectively from memory. *Psychonom. Bulletin and Review*, 4(2), 145-166.

Siskind, J. M. (1996) A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, *61*(1-2):1-38.

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

Vouloumanos, A. (2008). Fine-grained sensitivity to statistical information in adult word learning. *Cognition, 107,* 729-742.

Yu, C. (2008). A statistical associative account of vocabulary growth in early word learning. *Language Learning and Acquisition*, *4*(1), 32-62.

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

Yurovsky, D. & Yu, C. (2008). Mutual exclusivity in cross-situational statistical learning. *Proceedings of CogSci 30.* Austin, TX: Cognitive Science Society.