Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Statistical Learning of Visuomotor Sequences: implicit Acquisition of Sub-patterns

Permalink

https://escholarship.org/uc/item/6ct525m1

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 20(0)

Authors

Hunt, Ruskin H.
Aslin, Richard N.

Publication Date

1998

Statistical Learning of Visuomotor Sequences: Implicit Acquisition of Sub-patterns

Ruskin H. Hunt (ruskin@bcs.rochester.edu)
Dept. of Brain and Cognitive Sciences, Meliora Hall
University of Rochester, Rochester, NY 14627

Richard N. Aslin (aslin@cvs.rochester.edu)
Dept. of Brain and Cognitive Sciences, Meliora Hall
University of Rochester, Rochester, NY 14627

Abstract

A visuospatial reaction time task was used to gain an on-line measure of learning as subjects responded manually to strings of stimuli containing embedded transitional probabilities. We hypothesized that items within a stimulus sequence that have low transitional probabilities will be learned more slowly than items that have high transitional probabilities. Subjects were instructed to make button press responses to stimulus strings composed of sequences of lights. Items in the strings were organized into triplets, with a low average transitional probability for the first item in a triplet, and transitional probabilities of 1.0 for the second and third items. Results indicate that learning is poorer for stimulus items with low transitional probabilities than for stimulus items with high transitional probabilities. This work ties together a number of previous investigations of sequence learning, and has implications for how more complicated, hierarchically structured sequential input, such as language, may be learned.

Introduction

Sequence learning is of interest to cognitive scientists because of the wide variety of human experiences and behaviors that are serial in nature, such as language production and comprehension, complex motor tasks, and perception of sequential regularities in the environment. Despite the obvious conceptual links between studies of sequential motor learning and language, these literatures have rarely intersected. Studies of artificial language learning (e.g., Braine, 1963; Morgan, Meier & Newport, 1987) have typically involved complex grammatical rules and short strings of elements (visual and/or auditory). Subjects are exposed to grammatical strings and then tested for their knowledge of the underlying rules. In contrast, studies of sequential motor learning (e.g., Cleeremans & McClelland, 1991) have utilized reaction time (RT) measures to investigate improvements over time and exposure, as well as the relationship between sequential dependencies and motor performance. Of special interest is the *implicit* nature of some forms of sequence learning, in which subjects are unable to describe the underlying structure of a complex pattern of serial inputs despite clear evidence that this structure has been learned.

In an example of learning based on a fairly arbitrary, artificial grammar (i.e., one that does not reflect many of the subtle complexities of natural languages), Reber (1967) presented subjects with strings of visual stimuli organized according to a finite state grammar. The strings consisted of printed items presented simultaneously, rather than sequentially, with each string generated by a single pass through the grammar. Reber demonstrated that subjects exposed to strings generated by such a finite state grammar can discriminate grammatical from non-grammatical test strings. Note that Reber's use of a finite state grammar to generate these strings introduced variability in the learning set that was considerably greater than that presented in some serial RT tasks (e.g., Cohen, Ivry & Keele, 1990), but considerably less than the variability found in natural languages or the stimuli used by other researchers (e.g., Morgan, Meier & Newport, 1987).

Subsequently, Reber (1969) found that changing the specific symbols used in test strings produced little or no interference in a transfer task, whereas changing the "rules" of the grammar did. He argued that what subjects had become sensitive to was not the order of specific sequences of symbols, but rather their underlying grammatical structure. Finally, Reber (1976) found that learners who were informed that there was some kind of structure in the input and were encouraged to find it performed less well on a grammaticality judgment task than did learners who were given neutral instructions. These results were interpreted as evidence of an interference effect, produced by an explicit search for rules, which led to the formation of false hypotheses about the underlying structure of the input.

Other researchers have utilized stimuli containing far less internal variability than stimuli generated from finite state grammars of the type used by Reber. For example, Cohen, Ivry, and Keele (1990) used a RT paradigm in which short, fixed sequences were cycled repeatedly to assess the subject's ability to form associations between items in the sequence, as indicated by faster RTs, while they simultaneously performed a distraction task. Subjects used the fingers of one hand to press buttons corresponding to locations on a computer screen whenever an indicator appeared at a given button's location on the screen. Individual stimuli were presented serially as rapidly as the subject responded.

Cohen, Ivry, and Keele (1990) reported that when these fixed, repeated sequences were constructed with unique associations between sequence elements (e.g., 15243), subjects could not only learn the sequences, but could also do so in the presence of a distraction task. Moreover, subjects developed only an incomplete explicit representation of the sequential structure of the stimuli. However, when sequences contained multiple associations among elements because of repetition and changes of position within a sequence (e.g., 132312), they were not learned in the presence of a distraction task. Thus, it seems clear that in *simple* serial RT tasks (i.e., those with unique associations between sequence elements) involving a small number of items, learning can occur despite distractor interference and with little conscious awareness of the sequential structure.

Cleeremans and McClelland (1991) combined the use of finite state grammars and RT measures to evaluate the implicit learning of complex visuospatial sequences. As in Cohen, Ivry, and Keele (1990), subjects were presented with a position indicator on a computer screen and responded by pressing the appropriate key on a keyboard. Stimuli were generated from a finite state grammar, with the added feature that there was a 15% chance of a randomly selected successor item being substituted at any point in a given sequence. This source of noise served to prevent memorization of the input sequences (which were maximally 5 items in length) and allowed for analyses of multiple sequential dependencies (e.g., RT to item 5 given item 4, given items 3 and 4, given items 2, 3, and 4, etc.). Subjects produced faster RTs to grammatical sequences than to non-grammatical sequences, and RTs benefited from the cumulative predictiveness of the preceding context. As in Cohen, Ivry, and Keele's (1990) experiment, subjects reported only limited knowledge of the sequential structure of the stimuli. They also indicated that they did not try to use that knowledge during training because doing so tended to produce more errors and slower responses. The authors concluded that explicit knowledge, if present, played a minimal role in the learning process.

One of the advantages of using *probabilistic* sequences, such as those generated by a finite state grammar, is that they are sufficiently variable to minimize the likelihood of memorization. In contrast, *deterministic* sequences, which have a fixed order, may be memorized if they are fewer than ten items long, even without the massive amount of exposure required to exhibit learning of completely probabilistic sequences. An alternative is to generate stimuli that are simple enough to learn in fewer than the 60,000 trials used by Cleeremans and McClelland (1991), but complex enough that subjects are not likely to develop explicit representations of the sequence structure when there is no distraction task.

Recent work by Saffran, Newport, and Aslin (1996) and Saffran, Aslin, and Newport (1996) used sequences with both deterministic and probabilistic properties to study word segmentation from fluent speech. Continuous sequences of consonant-vowel (CV) syllables, ranging in duration from 2-7 minutes, were constructed by concatenating synthetic speech syllables according to a simple set of deterministic and probabilistic "rules." The entire sequence consisted of random orderings of four or six "words", where each word was composed of three CV syllables. Each word consisted of a unique trisyllabic string, although some of the syllables occurred in more than one word in the Saffran, Newport, and Aslin (1996) study of adults. In the Saffran, Aslin, and Newport (1996) study of 8-month-old infants, each syllable occurred in only one word. The coherence of syllable strings within words and the unpredictability of word orderings enabled both adults and infants to group syllables and segment words from the continuous stream of speech.

Using similar artificial language corpora, subsequent research has shown that word segmentation can proceed implicitly in adults and 7-8 year old children (Saffran, Newport, Aslin, Tunick & Barrueco, 1997), and that 8-month-old infants can segment words from fluent speech based solely on the transitional probabilities of successive CV syllables within the language corpus (Aslin, Saffran & Newport, 1998). Transitional probabilities, defined as the frequency in the corpus of item Y given X divided by the frequency of item X, were high for syllable-pairs within words and low for syllable-pairs spanning a word boundary. Sensitivity to the peaks and troughs in syllable-pair transitional probabilities is apparently sufficient for both infants and adults to group successive syllables and recognize them as relatively familiar in a post-exposure test. One limitation of this post-familiarization test paradigm, however, is that it does not provide an on-line measure of the learning process.

In the present study, we used a serial RT task to assess the time-course of learning when stimuli were structured like the trisyllabic words used in the studies of Saffran and colleagues. The fundamental difference between our paradigm and that of Cleeremans and McClelland (1991) is that our sequences are generated using both probabilistic *and* deterministic mechanisms, whereas theirs are entirely probabilistic. Our visuospatial stimulus sequences were composed of units analogous to words, within which transitional probabilities were perfectly predictive (1.0), whereas the random ordering of these words rendered the stimulus that followed each word much less predictable (transitional probabilities averaged 0.37). Our use of a RT methodology provides us with an on-line metric of improvement throughout the learning phase of the experiment for each type of transitional probability, thus avoiding the limitations of a post-familiarization testing phase. Moreover, our stimulus sequences were more complex than those from the studies that used fixed stimulus sequences, but not as variable as those from the studies that generated stimulus sequences using finite state grammars. Finally, the use of a visuospatial RT task of moderate complexity, with no instructions to the subject about the "rules" of the embedded sequences, suggested that any learning would likely be implicit. We predicted that if our assumptions about the nature of learning in this paradigm were correct, RTs would decline faster for stimulus elements with high transitional probabilities than for stimulus elements with low transitional probabilities.

Method

To obtain an on-line measure of sequence learning, we constructed a button-box capable of recording RT responses to visuospatial sequences that varied in transitional probability.

Apparatus

The button-box was built around an Apple Quadra 650 microcomputer and a hardware/software interface (designed by James R. Sawusch at the Speech Research Laboratory, State University of New York at Buffalo) for presenting computer controlled stimuli and recording responses. Special hardware timers in the interface allowed responses to be recorded with millisecond accuracy.

To allow subjects to interact with the computer, we designed a response box with lighted buttons. The buttons were housed in an aluminum box whose dimensions were 10 x 17 x 2 inches. The buttons were 3/4 inch in diameter, red, and the entire button surface was illuminated when stimuli were presented. Seven buttons were arranged in a semicircle with a five-inch radius, with an eighth button positioned at the center of the base of the semicircle. The eighth button served as a location for a "home" orienting response prior to the onset of each stimulus.

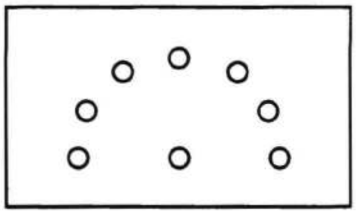

Figure 1: Button-box layout.

Stimuli

Stimuli consisted of *pairs* of simultaneously illuminated lights. We used pairs of lights because the seven buttons in the semi-circle did not provide a sufficient number of unique visuospatial stimuli to create a complex sequence with highly differentiated transitional probabilities. By using pairs of the seven non-"home" buttons, twenty-one unique pairs of lights could be presented in the semicircle.

Note that for convenience, we will adopt terminology analogous to that used by Saffran, Aslin, and Newport (1996) to describe a continuous sequence of trisyllabic words. Thus, in the present study, "syllable" refers to a single pair of lights, and "word" refers to a fixed sequence of three pairs of lights.

By convention, the buttons in the semicircle were numbered from left to right, and Figure 2 illustrates the button-pair that was illuminated simultaneously for a particular syllable (2, 4).

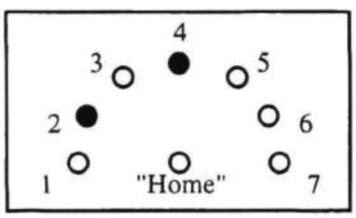

Figure 2: Syllable stimulus pair (2, 4).

The twenty-one unique "syllables" (pairs of lights) were organized into seven unique "words" (triplets of pairs). Within each word no individual button was lighted more than once (to control for repetition effects). Thus, within a word, six of the seven buttons were lit at some point as part of the word's component syllables. Words were randomly sequenced into blocks of seven, with the constraint that each word had to occur once within any given block (to equate frequency of exposure across blocks). Ten blocks were then concatenated into a single 70-word session, with the constraint that words could not repeat at the edges of blocks (again, to control for repetition effects). The stimuli were then presented as a single continuous sequence of events. Subjects took part in eight 70-word sessions during the course of one, hour-long training period, with 1-2 minute rest breaks between sessions. Six 8-session training periods were conducted over the course of six consecutive days, yielding a total of 48 sessions for each subject.

The organization of words used in the experiment is shown in Table 1. Each subject was assigned one of these "codes" for syllables and words for all sessions of the experiment. Words were then randomly sequenced into blocks, and blocks concatenated to form sessions separately for each subject, as outlined earlier.

Table 1: Unique triplets of button pairs.

Word	Syllable-Button Pairs
Assignment One	
A	(2, 4) (3, 7) (5, 6)
B	(4, 7) (1, 6) (3, 5)
C	(1, 5) (3, 4) (2, 6)
D	(6, 7) (2, 5) (1, 3)
E	(4, 6) (1, 2) (5, 7)
F	(1, 4) (2, 7) (3, 6)
G	(2, 3) (1, 7) (4, 5)
Assignment Two	
A'	(3, 4) (1, 7) (2, 5)
B'	(3, 6) (2, 4) (1, 5)
C'	(5, 6) (4, 7) (2, 3)
D'	(5, 7) (2, 6) (1, 4)
E'	(2, 7) (1, 3) (4, 6)
F'	(1, 6) (4, 5) (3, 7)
G'	(1, 2) (3, 5) (6, 7)

Using one of these two schemes to create a continuous sequence of button-pairs, the transitional probabilities within words (i.e., between the first and second syllables and between the second and third syllables) are 1.0. That is, because both the identity and order of syllables within words are fixed, the last two syllables of any word are completely predictable after the first syllable of a given word has been presented. The second and third syllables always, and only, follow the first syllable of a given word whenever it appears in the sequential corpus. However, at the boundaries between words, the transitional probabilities are much lower because of the random ordering of words within blocks, so that there is an average probability of 0.37 that the third syllable of a word will be followed by the first syllable of another word.

Subjects

Subjects were nine University of Rochester undergraduates (three male, six female). Subjects were paid for their participation, although payment was not used as a motivator to improve performance during the experiment. One subject was left-handed, and all subjects either had good, uncorrected vision or wore corrective lenses during the experiment.

Procedure

Subjects were comfortably seated in a sound attenuated room, with the button-box placed on their lap. Subjects were instructed to push buttons only with their preferred hand and to steady the box with their other hand. At the start of an experimental session, all of the buttons were illuminated for one second to indicate that the session was about to begin, after which all the lights were extinguished except for the "home" button (see Figure 2). Subjects were instructed to press and hold the "home" button until a pair of buttons (corresponding to a syllable) in the semicircle was illuminated. When the syllable button-pair was illuminated, the "home" button was simultaneously extinguished.

Subjects were informed that the experiment involved RTs, but no mention was made of learning or of patterns embedded in the sequential stimuli. Subjects were instructed simply to press the illuminated pair of buttons as rapidly as possible while maintaining high accuracy. They were told that the order in which they pressed the button-pairs was unimportant, but that the buttons should be pressed in succession (to avoid any tendency to simultaneously press two adjacently lit buttons). Both buttons in the pair remained illuminated until two buttons in the semicircle had been pressed. At that point, regardless of which two buttons had been pressed, the illuminated button-pair was extinguished and the "home" button was simultaneously illuminated. The subject returned to the "home" button by pressing it and holding it down until the onset of the next stimulus pair (syllable). After returning to the "home" button, there was a 100ms delay before the onset of the next stimulus pair. This delay period was included to prevent subjects from leaving the "home" button prematurely, which would have resulted in artificially short RTs. If the subject anticipated the next stimulus by releasing the "home" button before illumination of the stimulus pair, the stimulus pair was not illuminated. In such cases, the subject was required to return to the "home" button, the delay period was reinitiated, and if the subject correctly delayed release of the "home" button, the same stimulus pair that would have been presented during the previously aborted event was illuminated. No feedback was given to subjects at any point during the six-day training period. Subjects were given a free recall test prior to being debriefed after the final session.

Data Collection

RT data were collected throughout each session by recording the sequence of all button presses. Both button position (1-7) and RT were recorded for the first response to a stimulus pair (syllable), whereas only button position was recorded for the second response. A total of 10,080 RT measurements were possible across the entire six days of the experiment for each subject. If the subject pressed an unilluminated button, either as a first or second button press, the response was counted as an error and the RT for that stimulus presentation was discarded. Subjects' error rates were below three percent.

Subjects were classified as implicit learners, explicit learners, and non-learners based on RT and free recall performance. First, for individual subjects, mean reaction times were calculated for each syllable within each session for the last eight sessions. Subjects who demonstrated no difference in reaction time (paired t-test, $p < .05$) within the last eight sessions between syllable 1 and syllable 2 or between syllable 1 and syllable 3 were deemed not to have learned the underlying structure of the input, were classified as non-learners, and were excluded from the results ($N = 1$). Second, when subjects had completed the final (48th) session, but prior to being debriefed about the details of the experiment, they were asked if they had any general impressions about the nature of the task. All subjects who showed RT evidence of learning indicated that they thought they occasionally noticed patterns in the way the pairs of lights were presented. They were asked to write down any patterns they thought they could remember, and were given free access to the button-box while doing so. This was used as a measure of explicit knowledge. We tallied each subject's recall of correct sequences of syllables for part words (sequences of two button-pairs) and full words (sequences of three pairs). Subjects who demonstrated explicit knowledge of more than 3 part words or more than 1 full word were considered likely to have developed a strong explicit representation of the underlying structure of the input, were classified as explicit learners, and were excluded from the results ($N = 3$). Subjects who passed both the RT criterion for learning and the free recall criterion for non-explicit knowledge were classified as implicit learners and were included in the results ($N = 5$; three subjects trained on button assignment one, and two subjects trained on button assignment two).

Results

We hypothesized that with extended training, subjects who become implicitly sensitive to the structure within the sequence of stimulus syllables (button-pairs), should reflect this sensitivity through differentially faster RTs. Specifically, we predicted that RTs for first syllables, which have low average transitional probability, would begin to plateau before RTs for second and third syllables, which are maximally predictive. As shown in Figure 3, our results support this hypothesis. Across the first few sessions, there was general improvement in RTs for all three syllables, which likely reflects the subjects' adaptation to the general task of making rapid button-press responses from the "home" location. Subsequently, RTs for the second and third syllables continued to decrease, while RTs for the first syllable began to level off. By session forty-eight, the gap between RTs for first versus second and third syllables was on the order of 60ms.

An ANOVA, with task exposure (sessions) and syllable (1, 2, or 3) as repeated measures, indicated a significant linear main effect of exposure, $F(1,4) = 44.45$, $p < .005$, $MS_e = 17497.46$, and a significant interaction of linear trend

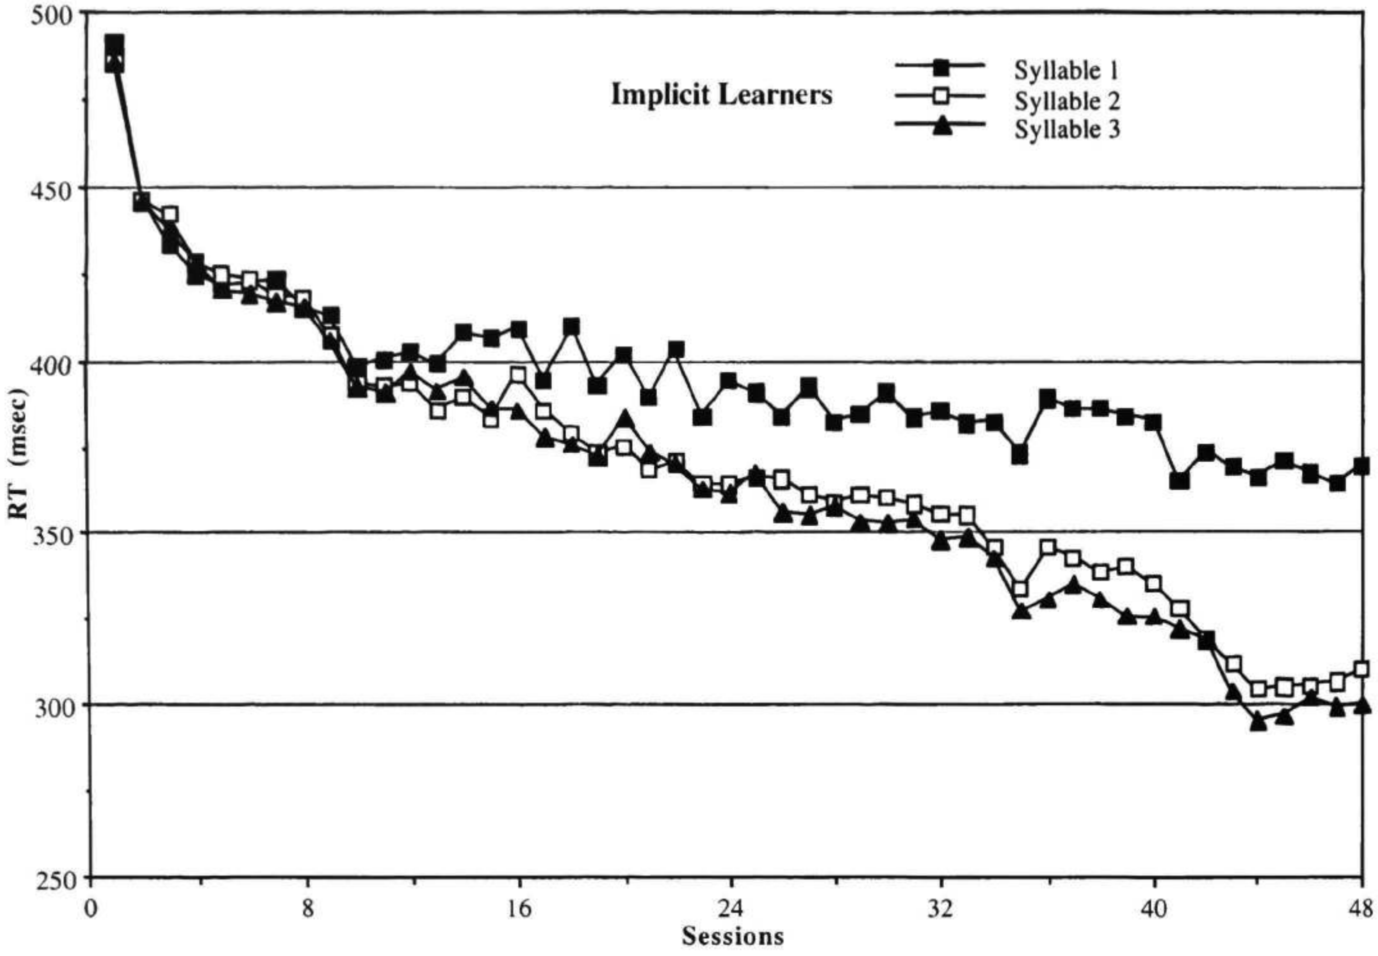

Figure 3: Mean reaction times for syllables 1, 2, and 3 for implicit learners (N = 5).

of exposure with syllable, $F(2, 8$, Greenhouse-Geisser adjusted to $1.18, 4.74) = 11.67$, $p < .05$, $MS_e = 2044.54$. This interaction indicates not only that there was an overall linear decrease in RT with exposure to the task, but also that this change was not the same for all syllables. The ANOVA also revealed significant simple effects of linear trend of exposure for each of the different syllables: syllable one, $F(1,4) = 104.80$, $p < .001$, $MS_e = 1343.83$; syllable two, $F(1,4) = 506.98$, $p < .001$, $MS_e = 896.18$; syllable three, $F(1,4) = 133.53$, $p < .001$, $MS_e = 133608.44$. Finally, there was a significant interaction between linear trend of exposure for syllable one and linear trend of exposure for syllables two and three combined, $F(1,4) = 12.60$, $p < .05$, $MS_e = 3739.26$, but not between linear trend of exposure for syllable two versus linear trend of exposure for syllable three, $F(1,4) = 1.68$, $p > .1$, $MS_e = 349.82$. These last results indicate that the linear improvement in RT with increased exposure to the task is different for syllable one than for syllables two and three.

It is clear from Figure 3 that the rate of improvement for syllable one is slower than for syllables two and three, and that the improvement for syllables two and three is roughly equivalent. A final post-hoc, non-trend comparison revealed that the effect of exposure on RT for syllable two versus syllable three was not different, $F(47, 188$, Greenhouse-Geisser adjusted to $3.08, 12.34) = .91$, $p > .1$, $MS_e = 47.08$, consistent with the fact that the transitional probabilities predictive of these syllables were both 1.0.

Discussion

The results of this study support the hypothesis that with increasing amounts of exposure to sequential stimuli, low probability transitions are learned more poorly than high probability transitions. The reaction time methodology provides a detailed on-line measure of learning, as indicated by improvements in RT, throughout a multi-day, multi-session learning process. This study paves the way for work on the learning of stimulus sequences with more complicated embedded patterns of transitional probabilities.

From the summary data shown in Figure 3, a number of points should be emphasized. First, although our statistical analysis revealed a linear effect of sessions of exposure, it is likely that the data reflect the initial portions of non-linear (exponential) learning curves. We suspect that with more exposure to the task (i.e., a greater number of sessions), the rate of RT improvement would level off for each of the syllables. Anecdotally, a pilot subject (not included in the study) who was trained to 58 sessions did show a leveling off of RT for all three syllables, suggesting that it may be profitable to extend the length of task exposure beyond 48 sessions.

Second, in addition to the RT improvement resulting from subjects learning the embedded syllable-structure of words within the stimulus sequence, there was some non-specific (generalized) learning based on adaptation to the physical requirements of the task. Although we did not assess non-specific adaptation directly, it seems likely that it was confined primarily to the initial part (~10 sessions) of the experiment. Certainly it is reflected in the large improvement seen from session one to session two. Work in progress incorporates methods to help evaluate non-specific improvement in RT.

Third, for most of the sessions after session twenty-four, RTs were faster for syllable three than for syllable two. Although this difference was not statistically significant, it is tempting to speculate that syllable three may enjoy an enhanced advantage in predictability despite transitional probabilities equivalent to syllable two. This might occur if subjects come to use the combined occurrence of syllables one and two to predict the occurrence of syllable three. More research will be necessary to test this hypothesis, possibly involving greater numbers of syllables or involving stimuli designed so that prediction of syllable three is necessarily dependent on the combined sequence of syllables one and two.

Fourth, the significant separation of RTs based on transitional probabilities demonstrates that subjects have learned some aspect of the sequential structure of the visuospatial stimuli. The research by Saffran and colleagues (Aslin, Saffran & Newport, 1998; Saffran, Aslin & Newport, 1996; Saffran, Newport & Aslin, 1996; Saffran, Newport, Aslin, Tunick & Barrueco, 1997) on word segmentation from fluent speech, which used similar sequential structures, employed a two-alternative forced-choice test after familiarization to the speech corpus. This post-familiarization test compared words with non-words and part-words, and clear evidence of discrimination suggested that syllable-pairs with high transitional probabilities were grouped into words. That is, statistical learning resulted in segmentation of words from fluent speech. In the present study, the pattern of RTs was analogous to word segmentation in that motor responses were facilitated for predictive sequences of visuospatial stimuli. While there is no direct evidence that the visuospatial sub-patterns making up words in the present study were implicitly *segmented*, it seems likely that evidence for segmentation would be found in a post-familiarization test. Such a test would be necessary to draw this kind of inference.

The results of this preliminary experiment have implications for more complex stimulus sequences. One hundred years ago, in their studies of telegraph operators training to use Morse code, Bryan and Harter (1897) recognized that the timecourse of learning often reflects the underlying hierarchical structure of the sequential "code." The present study, despite theoretical linkages to language learning, was modeled after a word-segmentation task in which the sequential structure was not hierarchical. We are currently expanding this line of work to determine, on-line, how subjects learn embedded structures that have multiple hierarchical levels of organization. The serial RT paradigm is particularly well suited to studies aimed at understanding the timecourse, and variability between subjects, of implicit learning at different levels of structural complexity.

Extensions of this work to studies of transfer of learning may also prove fruitful. For example, subjects who learn one specific organization of syllables within words may readily transfer that learning to a different organization, provided that the same embedded structure of transitional probabilities is maintained. These empirical measures of learning are also eminently suitable for computational modeling, using a variety of architectures, including the simple recurrent networks implemented by Cleeremans and McClelland (1991).

Acknowledgments

This research was supported by NSF research grant SBR-9421064 to RNA and NIH training grant DC00035 to RHH.

References

Aslin, R. N., Saffran, J. R. & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science, 9*, 321-324.

Braine, M. D. S. (1963). On learning the grammatical order of words. *Psychological Review, 70*, 323-348.

Bryan, W. L., & Harter, N. (1897). Studies in the physiology and psychology of the telegraphic language. *Psychological Review, 4*, 27-53.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General, 120*, 235-253.

Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 17-30.

Morgan, J. L., Meier, R. P., & Newport, E. L. (1987). Structural packaging in the input to language learning: Contributions of prosodic and morphological marking of phrases to the acquisition of language. *Cognitive Psychology, 19*, 498-550.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior, 6*, 855-863.

Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology, 81*, 115-119.

Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory, 2*, 88-94.

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

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language, 35*, 606-621.

Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological Science, 8*, 101-105.