Title
Behaviorism Reborn? Statistical Learning as Simple Conditioning

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 29(29)

ISSN
1069-7977

Authors
Gureckis, Todd M.
Love, Bradley C.

Publication Date
2007

Peer reviewed

Behaviorism Reborn? Statistical Learning as Simple Conditioning

Todd M. Gureckis (tgurecki@indiana.edu)
Department of Psychology, Indiana University, 1101 E. 10th Street, Bloomington, IN 47405 USA

Bradley C. Love (brad_love@mail.utexas.edu)
Consortium for Cognition and Computation, University of Texas at Austin, Austin, TX 78712 USA

Abstract

In recent years, statistical learning (SL) has emerged as a compelling alternative to nativist theories of language acquisition (i.e., Chomsky, 1980). However, in many ways the framework of statistical learning echoes aspects of classic behaviorism by stressing the role of associative learning processes and the environment in shaping behavior. How far backwards has the needle swung? In this paper, we show how a subset of behaviors studied under the rubric of SL are in fact entirely consistent with a simple form of conditioned priming inspired by models from the behaviorist tradition (i.e. a variant of Rescola-Wagner which learns associative relationships through time). **Keywords:** Statistical Learning; Simple Recurrent Network; Behaviorism; Language; Serial-Reaction Time Task

In recent years, statistical learning (SL) has emerged as a compelling alternative to nativist theories of language acquisition (e.g., Chomsky, 1980). For example, young infants are able to learn the statistical patterns in short segments of artificial speech and to use these dependencies to isolate word-like units (Saffran, Aslin, & Newport, 1996). Results such as these appear to challenge arguments concerning the "poverty of the stimulus" by demonstrating how the environment may provide subtle sources of structure which language learners can use to bootstrap into adult capabilities.

The implications of this work revisit the tension in Chomsky's (1959) attack on the behaviorist research program. Much like current SL research, Skinner's (1957) research program placed a strong emphasis on the role of the environment and associative learning processes in shaping behavior. Interestingly, many experimental inquiries into SL have relied on simple tasks which share much in common with stimulus-response or stimulus-stimulus conditioning. For example, in tasks such as the serial reaction time (SRT) task (Nissen & Bullemer, 1987), the syllable task (Saffran et al., 1996), and artificial grammar learning (A. Reber, 1967) subjects are presented with cued response sequences, auditory sequences, or letter strings with the goal of assessing learned associations between stimulus elements. Given these aspects of common theoretical alignment between SL literature and classic behaviorism, we reasonably might wonder the degree to which SL research reflects a return to behaviorist principles.

In this paper, we will argue that an important subset of the learning behaviors studied under the rubric of SL are, in fact, consistent with a simple form of conditioned priming well captured by models inheriting from behaviorist tradition (Rescorla & Wagner, 1972). Central to our analysis is a detailed comparison between human learning and two computational models of statistical sequence learning. While closely related, these models differ in the types of associative learning processes they utilize. The first model we consider, the simple recurrent network (SRN, Elman, 1990), is a recurrent network architecture which learns via back-propagation to predict successive sequence elements on the basis of the last known element and a representation of the current context. This mechanism has been shown to posses a number of powerful computational properties. For example, the SRN can act as a "universal approximator" of arbitrary finite state machines (Cleeremans, Servan-Schreiber, & McClelland, 1989).

We compare the SRN to a simple conditioning model based on direct stimulus-response mappings: the linear associative-shift register (LASR) model of sequence learning (Gureckis & Love, 2005). Like the SRN, LASR is an associative model of sequential processing. However, instead of relying on recurrent patterns of activation, memory in the model takes the form of a simple shift-register memory (similar to the Buffer network of Cleeremans, 1993). Unlike the SRN or the Buffer network, this memory is directly associated with output responses without mediation by an internal abstraction or recoding using hidden units. In this way LASR represents a variant of Rescorla-Wagner which learns statistical relationships through time. Due to limited space we refer the reader to Gureckis & Love (2005) for the mathematical details of the model. A critical difference between these two models concerns their assumptions about the stages of processing that mediate stimulus and response. In LASR, sequence learning is driven through direct stimulus-stimulus or stimulus-response associations while the SRN includes a set of internal "transformative" representations which mediate behavior.

In this paper, we focus our analysis on two empirical studies of SL behavior in a standard serial reaction time (SRT) task. On each trial of our experiments, subjects were presented with a set of six response options, one of these options was cued, and subjects were instructed to quickly and accurately press a corresponding button. Unknown to the subject, the pattern of cued-responses followed simple sequential patterns which we manipulated. The primary variable of interest was the time between presentation of the cued stimulus and the subject's

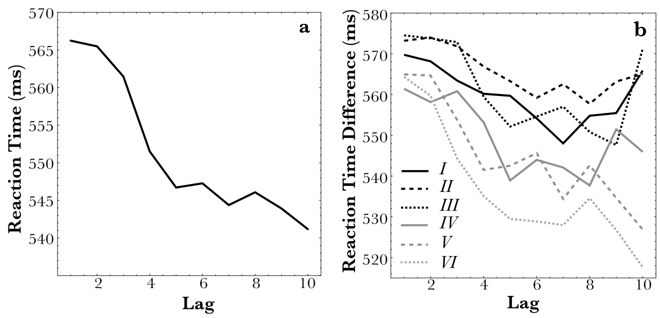

Figure 1: *Panel A*: Reaction time as a function of the lag between two repeated stimuli (Exp. 1). Reaction time is faster as lag increases. *Panel B*: The evolution of the lag effect over the course of six blocks the experiment.

response under the assumption that faster RTs reflect a learned anticipation of the next sequence element.

In Experiment 1, the pattern of cues followed a simple pattern defined by a negative recency relationship. In Experiment 2, we created a sequence which required the computation of higher-order statistical relationships (such as the second order conditional probability). Both tasks were otherwise identical. Our goal was to evaluate the speed at which learning unfolds for each type of sequence relative to the predictions of our benchmark models. To foreshadow, in Exp. 1 we find that human subjects quickly adapt to the structure of the sequence. In contrast, learning of higher-order statistical relationships (Exp. 2) is found to be considerably slower.

Experiment 1

In our first study, the pattern of cues in a standard SRT task was determined by a simple rule: each of the six choice options had to be visited once in a group of six trials in a random order (with the additional constraint that the last element of one group could not be first element of the next to prevent direct repetitions) (Lee, 1997; Boyer, Destrebecqz, & Cleeremans, 2004). Examples of legal six-element sequence groups are 132546, 432615, and 246153 which were concatenated into a continuously presented series. Beside these six-element groups, there is another subtle source of structure. Since the same cued response cannot repeat on successive trials, repeated cues were separated by a minimum of 1 trial. The longest possible lag separating two repeating cues is 10 which occurs when the first item of one six-element sequence group is repeated as the last item of the following six-element group. Overall, the probability of encountering any particular cue is an increasing function of how many trials since it was last encountered. In other words, if you haven't been cued to press a particular button in a few trials, it becomes increasingly likely that you will soon. Boyer, et al. found that when engaged with this particular sequence subjects responses are driven primarily by the lag between repeated cues. In order to examine more closely the early blocks of learning and to collect our own data against which to compare models we ran a conceptual replication of Boyer, et al (Exp. 1) using the same sequence and procedure.

Design and Procedure Forty-six Indiana University undergraduates participated and were paid \$8 for their time. Each block consisted of 180 trials using the sequence structure described above and in Boyer, et al. (2004).

Results Data was analyzed by finding the median RT for each subject for each block of learning. Any trial in which subjects responded incorrectly were dropped from the analysis (mean accuracy was 96.9%). The basic results are shown in Figure 1. Panel A replicates the lag effect first reported by Boyer, et al. averaged over all subjects. Participants were faster to respond to an item the longer it had been since it was last visited, $F(9, 396) = 13.477$, MSe = 3983, $p < .0.001$. A trend analysis on lag revealed a significant linear effect ($t(44) = 4.99$, $p < .001$) and a smaller but significant quadratic effect ($t(44) = 3.241$, $p < .003$). Like the original study, our subjects showed evidence of learning with little training (only six blocks).

However, looking closer at the evolution of the lag effect over the course of the experiment (Figure 1B) reveals that the difference in RT between recently repeated events and events separated by many intervening elements increases consistently over the course of the experiment. Early in learning (blocks 1 and 2), subject show about a 10ms facilitation for lag-9 or lag-10 items over lag-1 responses while by the end of the experiment (blocks 5 and 6), this facilitation increases to about 45ms. These observations were confirmed via a two-way repeated measures ANOVA with lag (10 levels) and block (6 levels) which revealed a significant effect of lag ($F(9, 387) = 11.496$, MSe = 19992.2, $p < .001$), block ($F(5, 215) = 11.167$, MSe = 52464.2, $p < .001$), and a significant interaction ($F(45, 1935) = 1.494$, MSe = 1726.9, $p < .02$).

Simulation Results

In order to evaluate the ability of both the SRN and LASR to account for the results of Experiment 1 we tested each model under conditions similar to those of human subjects. Each model was initialized with 6 input units and 6 output units which corresponded to the six choice options in the task. On each trial, the activation of one of the six input units was set to 1.0 (corresponding to the cued response shown to subjects) and the model's task was to predict the next sequence element. The resulting pattern of activation across the six output units in response to this input pattern was converted into a choice probability. Human reaction time in the experi-

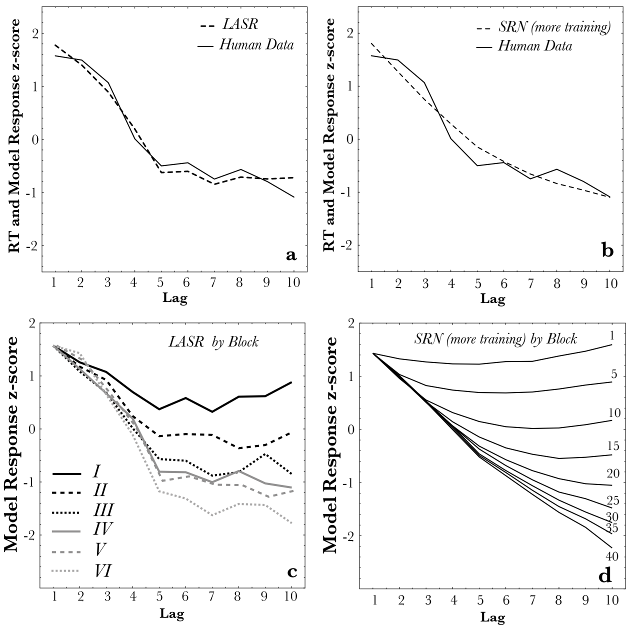

Figure 2: *Panel A&C*: Average results from best fit parameters for LASR. *Panel B&D*: The results for the SRN after extensive training (number of blocks is shown on the right in panel D).

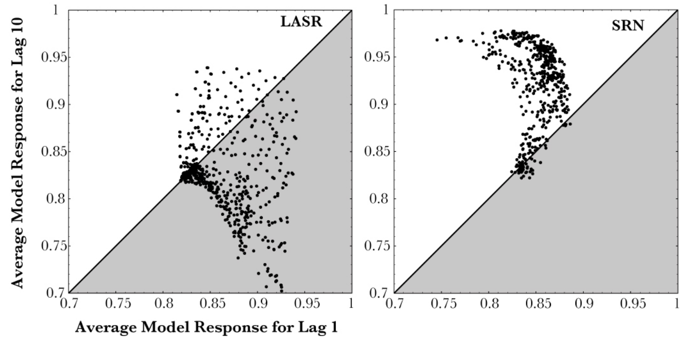

Figure 3: Explorations of the parameter space for LASR and the SRN in Experiment 1. Each model's average response for lag-1 are plotted against the average response lag-10. The division between the grey and white regions represent the line y=x. Each point in the plot represents the performance of the respective model with a particular setting of the parameters. If the point appears below the y=x line in the grey area, it means the model predicts faster responding to lag-10 events than to lag-1 (the correct pattern).

ment was assumed to inversely relate model's response probability for the correct response on that trial (reflecting correct anticipation). Each simulation consisted of running each model 200 times over each of the same sequences given to subject, each time with a different random initial setting of the weights (sampled from a uniform distribution between -1.0 and 1.0).

LASR Figure 2A&C shows LASR's response at each of the 10 levels of lag along with the evolution of this lag effect over the course of learning. Data in panel C was recoded in terms of the amount of RT facilitation over lag-1 responding, thus RT to stimuli at lag-1 was always scored as 0 ms with increasingly negative values for longer lags. This allows us to measure the changes in learning to respond to stimuli with greater lag independent of unspecific practice effect over the course of learning blocks. In addition, all human and model responses have been converted to average z-scores for comparison.

Starting from a random initial setting of its weights the model very quickly adapts to the lag structure of the stimuli. Like the human data, LASR shows a strong lag effect even in the first block of training. Furthermore, the strength of this learning effect continues to increase until the end of the experiment. Indeed, the model provides a very close quantitative fit of the data (the average correlation between the model and human data shown in the left panel of Figure 2 was 0.981 (SD = 0.005) and the mean RMSE was 0.181 (SD = 0.0026).

However, the overall pattern of results was not specific to this particular setting of the parameters. In the left panel of Figure 3, we show the model's predicted relationship between lag-10 and lag-1 responses over a large number of parameter combinations. Each point in the plot represents the performance of the LASR with a particular setting of the parameters (simulations were averaged over a smaller set of 90 independent runs of the model due to the size of the parameter space). Also plotted is the $y < x$ region. If a point appears below the $y = x$ line (in the gray region of the graph), it means the model predicts faster responding to lag-10 events than to lag-1 (the correct pattern). The left panel of Figure 3 shows that over the majority of parameter combinations LASR predict faster responding to lag-10 events (67% of the parameter combinations we tried captured the correct qualitative pattern).

SRN The SRN describes learning in a much different way. Despite extensive search, under no circumstances could we find parameters which allowed the SRN to learn the negative recency effect in the same number of trials as human subjects. A similar exploration of the parameter space as was conducted for LASR is shown in the right panel of Figure 3. As shown in the figure, very few parameter combinations predict the correct ordering of lag-1 and lag-10 responses (i.e., very few points appear below the y=x line). In fact, of all the combinations evaluated only 8% predicted the correct qualitative pattern. However, manual examination of these rare cases revealed that in these situations the model failed to capture the overall pattern of successively faster responses across all 10 lags demonstrated by human subjects in Figure 1B. In contrast, following Boyer, et al. (1998), when the SRN is given more extensive exposure to the material such that training lasts for 30,240 trials (30 times the training that subjects in Experiment 1 experienced) the model is able to eventually adapt to the lag structure of the material, although it does so in a more uniform way than do human subjects. For example, notice in Figure 2D how the lag effect predicted by the SRN with extensive training is almost linear.

Discussion The results of Experiment 1 show that even when faced with a complex and highly variable environment, subjects are able to quickly pick up on inhibitory relationships between their own actions and presented stimuli. This agility contrasts with the predictions of the SRN, which requires considerable training in order to elaborate its representation of the task. In Experiment 2 we examine learning in a similar task which stresses a different type of learning. In particular, we consider learning of sequences where prediction-relevant information is composed entirely of second-order conditional (SOC) relationships. The critical question is how readily human subjects learn about these sequence patterns compared to those used in Experiment 1 and how these differences are captured by existing models.

Experiment 2

In the original paper on the SRN, Elman (1990) demonstrated how the network can learn to predict elements of a binary sequence defined by the following rule: every first and second element of the sequence was generated at random, while every third element was a logical XOR of the previous two. In this sequence the outcome on every 2nd trial is paired an equal number of times with each successor outcome on every 3rd trial (i.e. 1 is equally likely to be followed by a 0 or a 1). Likewise, each outcome on every 1st trial is also paired an equal number of times with each outcome on every 3rd trial, and so on. Thus, unlike the sequence in Exp 1, no predictive information can be obtained using first-order transition probabilities alone. Indeed, the only predictable component of this sequence requires learners to integrate information from both time step $t-2$ and $t-1$ simultaneously to represent the rule (or higher-order conditional probability) "if the last two events are the same the next response will be 0, otherwise it will be 1." In Experiment 2, we consider human learning in a similar task where the prediction-relevant aspects of the sequence were restricted to higher-order conditional relationship.

Design and Procedure Fifty-two University of Texas undergraduates participated for course credit and a small cash bonus tied to accuracy. Subjects were evenly divided between one of two conditions: a second-order conditional (SOC) or a first-order conditional (FOC). The training sequence was the same for both groups and was constructed by randomly sampling sequence elements from the following list of triplets [0, 2, 4], [0, 3, 5], [1, 2, 5], and [1, 3, 4] and presenting each element one at a time in a continuous stream. This sequence is similar to the XOR task because every third element (either 4 or 5) is uniquely predicted based on the previous two elements. After completing 10 training blocks (consisting of 90 trials each), participants were given two transfer blocks (180 trials total). After the two transfer blocks in each condition, the sequence returned to the original training sequence for an additional 3 blocks for a total of 15 blocks in the experiment (or 1350 trials).

The structure of the transfer blocks depended on which condition the subject was assigned to. In the SOC condition, participants were transfered to a new sequence which (in addition to the triplets from the training sequence described above) included the following triplets [0, 2, 5], [0,3,4], [1,2,4], and [1,3,5]. Notice that in these new items the third element of each subsequence is flipped (thus, the transfer set includes both 134 and 135). Thus, only the critical predictable sequence element (position 3) is changed during transfer. In the FOC condition, participants were transferred to a different sequence which included the triplets [0,4,2], [1,4,3], [0,5,3], and [1,5,2] (in addition to the triplets used in training). These 4 new items shuffle two of the columns (i.e. positions 2 and 3). This transfer condition tests for learning of first-order information such as the fact that 5 always followed either a 2 or a 3 but never a 0 or 1 in the training set. In contrast the SOC transfer sequence isolates learning to the higher-order non-linear component.

Results For each subject, the median RT was computed for every block of learning. Any trial in which subjects responded incorrectly were dropped from the analysis. Overall accuracy was 97.8% in the SOC condition, and 96.4% in the FOC condition. Figure 4 shows the mean of median RT for each block of learning. One way to assess learning in this task to examine if RT increased during the transfer blocks (blocks 11 and 12, where the sequence structure changes) relative to the surrounding sequence blocks. In order to assess this effect, we computed a pre-transfer, transfer, and post-transfer score for each subject by averaging over blocks 9 and 10 (pre-transfer), 11 and 12 (transfer), and 13 and 14 (post-transfer).

In the SOC condition, we found no significant difference between the pre-transfer and post-transfer RT compared with RT during the transfer phase, $t(25) = 1.62$, $p > .1$. Reaction time values between the pre-transfer

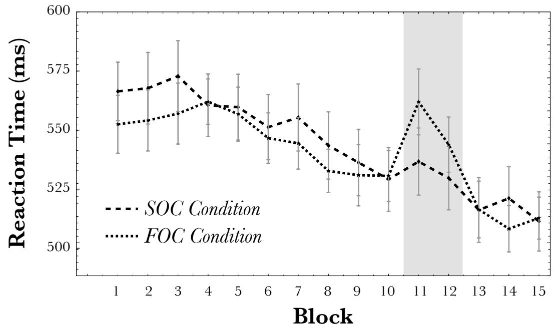

Figure 4: Mean of median reaction times for Experiment 2 as a function of training block. Error bars are standard errors of the mean. Transfer blocks are highlighted by the grey strip covering blocks 11 and 12.

and transfer block scores also did not reach significance, $t(25) = 1.26$, $p > .2$ (M= 514 ms, 510 ms, and 499 ms, respectively). We found a significant linear trend across the three phases, $t(25) = 3.60$, $p < 0.002$, while the quadratic relationship failed to reach significance, $t(25) = 1.62$, $p > 0.1$. Thus, we failed to find evidence that subjects slow their responses during the transfer blocks relative to the surrounding blocks.

However in the FOC condition, we found a highly significant difference between the pre-transfer and post-transfer RT compared with RT during the transfer phase, $t(25) = 7.16$, $p < .001$ and between the pre-transfer and transfer score, $t(25) = 4.28$, $p < .001$ (M= 510 ms, 532 ms, and 491 ms, respectively). Both the linear and quadratic trends were significant ($t(25) = 4.51$, $p < 0.001$ and $t(25) = 7.16$, $p < 0.001$, respectively). Subjects in this condition *did* slow down during the transfer blocks relative to the surrounding blocks (by about 22 ms on average).

Discussion The results of Experiment 2A failed to find strong evidence that subjects learn the higher-order statistical patterns in the sequence. For example, we found no evidence of differentially faster responses to predictable versus unpredictable sequence elements in the transfer block after 900 trials of learning, and subjects show no evidence for differential RT to the SOC transfer blocks. This result was surprising given previously published results showing that subjects can learn higher order sequence structures (P. Reber & Squire, 1994; Remillard & Clark, 2001; Fiser & Aslin, 2002). However, unlike at least some of these previous reports, the task we utilized carefully restricts the statistical information available to subjects. In addition, many previous studies used extensive training regimes which took place over a number of days (Cleeremans & McClelland, 1991; Remillard & Clark, 2001).

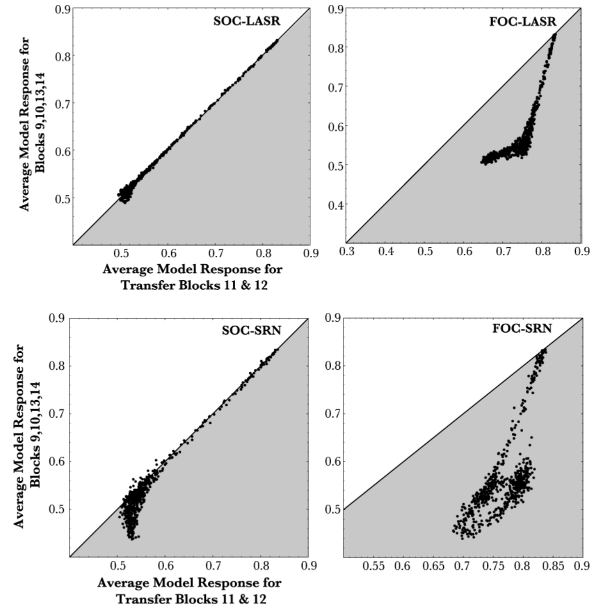

Figure 5: Explorations of the parameter space of the LASR and SRN in Experiment 2. Each point in the plot represents the performance of the LASR with a particular setting of the parameters. If the point appears below the y=x line (in the grey region), it means the model predicts slower responding during the transfer blocks (and thus a learning effect).

Despite the fact that few, if any, of our subjects demonstrated any evidence of learning the structure of the XOR task within the 1350 trials of the experiment, our conclusion is not that human subjects cannot such sequences. Instead, we think that our results simply show that learning of these relationships in isolation is considerably slower than learning for other types of information. In fact, in other experiments (not reported) here, under longer, multi-day training regimes our subject eventually showed strong evidence of learning.

Simulations

Our goal in the following simulations was to evaluate the overall ability of each model to learn the sequential XOR tasks rather than to precisely fit the results of Experiment 2 which showed little evidence of SOC learning. Thus, we simulated both models over a large range of their parameters and considered what each set qualitative predicted. Of particular interest was a comparison between this analysis and the parameter space analysis reported for Experiment 1.

LASR LASR predicts that subjects will only slow their responses during transfer blocks in the 6Cb condition. This is because LASR, lacking hidden units, is unable to learn the higher-order statistical component of the sequence. Instead, it is limited to learning sequences with first order relationships. The results of a parameter space analysis with LASR for Exp. 2 confirms these intuitions. The model was evaluated on the same ranges of parameters used in Experiment 1. Figure 5 plots the average model responses during the transfer blocks (11 & 12) versus the surrounding training blocks (9, 10, 13, & 14) for each parameter combination. As is clearly visible, LASR predicts a strong difference between the transfer and surrounding block only in the FOC condition. Interestingly, this is exactly the pattern which our subjects showed.

SRN In contrast to LASR, the SRN predicts that subjects will slow their responses during the transfer blocks in both experimental conditions. This is confirmed in Figure 5. Overall, 83% of the parameter combinations show learning in the SOC condition. The magnitude of the effect was on average significantly different from zero, M=.01, SD=.02, t(649)=20.82, $p < .001$. Finally, like LASR the SRN predicts that virtually all parameter combinations (99%) show learning in the FOC condition (M=.2, SD=.07, t(649)=20.82, $p < .001$). In contrast to the SRN simulations of Experiment 1, it was much easier to find parameters which allowed the SRN to learn in all three conditions of Experiment 2.

General Discussion

These results place the SRN in a difficult position. In Experiment 1, the model fails to learn the material at the rate that human subjects do while predicting that learning in Experiment 2 should actually be more robust. In contrast, human learning follows the opposite pattern. Learning to inhibit recent actions is quite easy for subjects, while learning sequences composed of higher-order relationships is more difficult. Our analysis shows that at least a subset of results in the SL literature may be better accounted for with these direct stimulus-response associations without the need for internally mediated representations such as those used in the SRN. In fact, the transformative representational processes in the SRN appear too complex for these types of learning behavior. Instead, a simple more limited process like LASR is at an advantage, as it is able to rapidly adapt to the structure of the environment. Overall our results suggests that some of the methodologies and results in the SL literature may have more in common with the behaviorist research tradition than has been previously recognized.

The success of the SRN across so many domains of cognitive behavior has either tacitly or explicitly led many authors to suggest that a single, domain general learning process may govern sequential processing throughout cognition. In light of our results we suggest that sequential learning behavior is better conceptualized as being organized at many overlapping time scales which differ in complexity. At the lowest level of complexity, such as the kind of behavior we might engage in the course of single experimental session or in many aspects of daily life, learning appears largely consistent with a simple, limited processes based on Rescorla-Wagner (1972). On the other hand, given a lifetime of exposure to the structure of our native language, more complex processes unfold such as those embodied by the SRN. Many of the impressive applications of the SRN arise from the unique computational properties of this system which make it well suited for the domain of language. However, a more careful delineation maybe necessary between certain results in the SL literature and their broader implications for learning.

Acknowledgments This work was supported by NIH-NIMH training grant #:T32 MH019879-12 to T.M. Gureckis and AFOSR grant FA9550-04-1-0226 and NSF CAREER grant #0349101 to B.C. Love.

References

Boyer, M., Destrebecqz, A., & Cleeremans, A. (2004). Processing abstract sequence structure: Learning without knowing, or knowing without learning? *Psychological Research*.

Chomsky, N. (1959). A review of b.f. skinner's verbal behavior. *Language*, *35*(1), 26-58.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*(3), 235-253.

Cleeremans, A., Servan-Schreiber, D., & McClelland, J. (1989). Finite state automata and simple recurrent networks. *Neural Computation*, *1*, 372-381.

Fiser, J., & Aslin, R. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458-467.

Gureckis, T., & Love, B. (2005). A critical look at the mechanisms underlying implicit sequence learning. In *Proceedings of the 27th annual conference of the cognitive science society.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Lee, Y. (1997). Learing and awareness in the serial reaction time task. In *Proceedings of the 19th annual conference of the cognitive science society* (p. 119-124). Hillsdale, NJ: Lawrence Erlbaum Associates.

Nissen, M., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1-32.

Reber, A. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *77*, 317-327.

Reber, P., & Squire, L. (1994). Parallel brain systems for learning with and without awareness. *Learning and Memory*(217-229).

Remillard, G., & Clark, J. (2001). Implicit learning of first-, second-, and third-order transitional probabilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(2), 483-498.

Rescorla, R., & Wagner, A. (1972). A theory of pavolvian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. Black & W. Prokasy (Eds.), *Classical conditioning ii: Current research and theory* (p. 64-99). New York: Appleton-Century-Crofts.

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