Journal of Experimental Psychology:
Animal Behavior Processes
2013, Vol. 39, No. 2, 166–173

0097-7403/13/$12.00 DOI: 10.1037/a0031922

BRIEF REPORT

Modeling Human Sequence Learning Under Incidental Conditions

F. Yeates
University of Exeter

F. W. Jones
Canterbury Christ Church University

A. J. Wills
Plymouth University

R. P. McLaren and I. P. L. McLaren
University of Exeter

This research explored the role that associative learning may play in human sequence learning. Two-choice serial reaction time tasks were performed under incidental conditions using 2 different sequences. In both cases, an experimental group was trained on 4 subsequences: LLL, LRL, RLR, and RRR for Group "Same" and LLR, LRR, RLL, and RRL for Group "Different," with left and right counterbalanced across participants. To control for sequential effects, we assayed sequence learning by comparing their performance with that of a control group, which had been trained on a pseudorandom ordering, during a test phase in which both experimental and control groups experienced the same subsequences. Participants in both groups showed sequence learning, but the group trained on "different" learned more and more rapidly. This result is the opposite that predicted by the augmented simple recurrent network used by F. W. Jones and I. P. L. McLaren (2009, Human sequence learning under incidental and intentional conditions, *Journal of Experimental Psychology: Animal Behavior Processes*, Vol. 35, pp. 538–553), but can be modeled using a reparameterized version of this network that also includes a more realistic representation of the stimulus array, suggesting that the latter may be a better model of human sequence learning under incidental conditions.

Keywords: sequence learning, serial reaction time, augmented SRN

Understanding human sequence learning under incidental conditions, whether it involves learning a sequence of events or a sequence of actions, is key to explaining much of human and infrahuman behavior. To learn sequences, people and animals need to cope with information embedded in a temporal context, adding an extra dimension to the more static problems typically studied in research on associative learning, and bringing them closer to those that occur in real situations outside the laboratory. This extra complexity also constrains the modeling of human sequence learning, where it is often addressed by the addition of recursion to otherwise static models, for example, the simple recurrent network (SRN; Elman, 1990) and the augmented SRN (Cleeremans & McClelland, 1991). The question that this article addresses is whether or not these models provide adequate accounts of sequence learning under incidental conditions.

In the experiment reported here, we focused on a very simple task in which sequence learning is known to occur, even though it is not explicitly required, and is hence often cited as a situation in which "implicit" learning occurs. This is the variant of the two-choice serial reaction time (SRT) task recently developed by Jones and McLaren (2009). In this task, participants observe two circle outlines on a screen and are given two response keys, one for each circle. On each trial, one of the circles "fills in," and the participants press the corresponding key as quickly and accurately as possible. Following this, the circle outlines reappear for 500 ms before the next trial starts. Trials come rapidly one after the other, and the experience is of a fast-paced task that emphasizes speed and accuracy in reacting to the stimuli and requires little else.

In fact, for the experimental groups in this task, there is a probabilistic rule governing the sequence of locations in which the circle appears, knowledge of which could enable participants to prepare for the stimulus and so increase the speed and accuracy of their responding. The roles of the two stimulus locations are counterbalanced across participants and henceforth are referred to as X and Y rather than right and left. In our previous work (Jones & McLaren, 2009), we were able to show that the augmented SRN

F. Yeates, R. P. McLaren, and I. P. L. McLaren, School of Psychology, University of Exeter, Exeter, United Kingdom; F. W. Jones, Department of Applied Psychology, Canterbury Christ Church University, Tunbridge Wells, Kent, United Kingdom; A. J. Wills, Department of Psychology, Plymouth University, Plymouth, United Kingdom.

We gratefully acknowledge the support of the Economic and Social Research Council in funding this research (RES-000-22-4036, Dual Process Models of Sequence Learning and Serial Reaction Time Tasks), the University of Exeter for providing an Exeter Graduate Fellowship to Fayme Yeates, an anonymous reviewer for their comments, and two undergraduate students for help with data collection.

Correspondence concerning this article should be addressed to I. P. L. McLaren, School of Psychology, University of Exeter, WA Singer Laboratories, Perry Road, Exeter EX4 4QG, United Kingdom. E-mail: i.p.l.mclaren@exeter.ac.uk

(Cleeremans & McClelland, 1991) could successfully model incidental learning of a sequence that comprised subsequences XXX, YYX, XYY, YXY, which follow the rule that if the first two locations are the same, then the third is an X; if they are different, then it is a Y.

In the current experiment, we varied the subsequences to see whether the augmented SRN could still model the results. Thus, one group in this experiment had XXX, YYY, XYX, and YXY as their subsequences, which follow the rule that the third element is the same as the first. By concatenating these subsequences (e.g., XXXYYYXYX . . . etc.), we can produce a sequential structure that has the property that two thirds of the time a trial is the same as the trial before last. The other group was trained on the complementary set XXY, YYX, XYY, and YXX, where the rule is that the third element is different from the first, so that after concatenation, two thirds of the time the current trial is different from the trial before last. In our experiments, learning was measured relative to pseudorandom control groups. The controls experienced a mixture of all eight subsequences so that the first trial had no predictive value for the third. Our interest, then, was in comparing the differences between experimental and control groups for those participants trained on sequences in which the first trial was different from the third (Group Different) with those trained on sequences in which the first was the same as third (Group Same). The factor of group, denoting the type of subsequences used during training, was a dummy variable for the controls as all of these participants received the pseudorandom mixture of all eight subsequences throughout.

We focused on this comparison because a simple extrapolation from the empirical results of Jones and McLaren (2009) leads one to predict that Group Different should have an advantage. This is because the subsequences XXX and YYY can be expected to be very difficult to learn based on these earlier findings and both of these subsequences fall in Group Same. Intriguingly, when we ran the augmented SRN on this new experiment with the same parameters as those used in Jones and McLaren, the pattern we obtained was actually the reverse, with Group Same subsequences learned better than Group Different subsequences. Thus, evidence-based intuition and the model seem to be in conflict, and an empirical test was needed to resolve the issue. We return to a discussion of the modeling once we have reported the results of our experimental work.

Method

Participants

The study was conducted on 128 participants, randomly divided into four groups (two experimental and two control). There were 32 participants in each of the two experimental conditions and in the two control conditions (both control conditions were actually treated identically and participants were randomly assigned as the control for one of the two experimental conditions). The participants were all students at the University of Exeter, ages from 18 to 35 years. In addition, each of the participants was rewarded for their contribution with £10 at the end of their second session.

Materials

The two-choice SRT task was run on an Apple Mac computer, with the basic display being one of two white outline circles on a black background. The circles were 1.9 cm in diameter and each was positioned 2.2 cm to the right or left of the middle of the screen, which was approximately 0.5 m from the participant. The stimulus was a white filled circle 1.9 cm in diameter that replaced either the right or left outline circle during the trials. The participants were instructed to press the *x* key on a QWERTY keyboard if the target stimulus appeared on the left and the *period* (.) key if the stimulus appeared on the right. These keys were chosen to be spatially compatible with the two stimulus locations.

Design

The experiment consisted of a two-choice SRT task that was conducted over two sessions, each lasting approximately 1 hr. The first session was usually undertaken in the morning, with the second session typically commencing after a 3- to 4-hr break on the same day. Both sessions consisted of 20 blocks of 120 trials, with the last five blocks of Session 2 acting as the test phase. All other blocks acted as the training phase. The blocks for each of the experimental conditions were constructed by concatenating equal numbers of the relevant subsequences, as already described. Thus, during the training phase of the SRT task, experimental participants in Group Different were presented with sequences made up of subsequences in which the location of the third trial was the location opposite to that of the first trial (e.g., XXY). The rule was different for participants in the experimental condition of Group Same; in training, they were presented with sequences made up of subsequences in which the third trial location was the same as the first trial location (e.g., XYX). During training, participants in the control conditions experienced pseudorandom blocks, which were created by concatenating equal numbers of the eight possible triplets in a random order (see Jones & McLaren, 2009, for further details). Note that, for all the conditions and groups, when the subsequences (or triplets) were concatenated, they formed continuous strings of trials, and previous evidence suggests that participants do not learn about the special status of the third trials, but rather learn the contingencies on a trial-by-trial basis (Jones & McLaren, 2009). When training blocks are considered trial-by-trial, trials consistent with the experimental groups' subsequences occur two thirds of the time, with the remaining third of trials being inconsistent (e.g., in the experimental condition of Group Same, XX is followed by X twice as often as it is followed by Y).

For all conditions, the last five blocks of Session 2 acted as the test phase and consisted of pseudorandom blocks only. By comparing experimental and control performance on what are effectively the same types of sequence, we controlled for possible confounds due to sequential effects (Jones & McLaren, 2009).

Procedure

As in Jones and McLaren (2009), the participants were instructed to respond as quickly as possible while avoiding errors. No mention was made of any sequential structure embedded in the task. On each trial, the stimulus remained on the screen until the participant had responded or was timed out for not having pressed

a key within 4.25 s of the stimulus' appearance. RT was measured from the stimulus' appearance on screen until the computer detected a key press, and a 500-ms response–stimulus interval was used. If participants pressed an incorrect key or were timed out, then the trial terminated and the computer issued a short "beep." Similarly, if they anticipated a stimulus (i.e., they responded less than 100 ms after its onset), then the trial was aborted and a beep sounded. Following each block, participants experienced a 30-s break during which they were shown their average RT (in milliseconds) and their accuracy (as an error percentage) for the last block. They were also informed whether these scores were better or worse than those from the previous block.

Results

For the test phase data, we computed difference scores by subtracting the mean RTs and proportion of errors for trials that were consistent with each subsequence from the respective scores for trials that were inconsistent. To illustrate, consider the Group Same experimental condition subsequence XXX. Any X trial that is preceded by XX (which we label as an XXX trial) is consistent with the subsequence XXX, whereas any Y trial that is preceded by XX (i.e., XXY) is inconsistent with this subsequence. Therefore, the RT difference scores for subsequence XXX were calculated by subtracting the mean RT for XXX trials from the mean RT for XXY trials, and a similar subtraction gave us an equivalent difference for errors. This was done for all experimental subsequences, for each group and condition. In the control conditions, consistent/inconsistent was a dummy variable that was determined by the subsequences in the paired experimental condition.

Because the experimental conditions experienced different trial orders from their respective control conditions in training, any difference between the conditions here could be due to sequential effects (i.e., performance differences on different trial orders) instead of sequence learning (cf. Soetens, Boer, & Hueting, 1985). To minimize this confound, we used the method of calculating inconsistent minus consistent training difference scores described by Jones and McLaren (2009). Sequential effects of up to order $n - 3$ were controlled for by equal-weight averaging of the two versions of each subsequence that exist when the $n - 3$ trial is also considered (e.g., the score for XXX is the equal-weighted average of the scores for YXXX and XXXX). Insufficient data meant sequential effects of $n - 4$ and greater could not be controlled for in this way, but an inspection of our data suggests that $n - 4$ sequential effects are in the order of 1 ms (for a more detailed discussion, see Jones & McLaren, p. 543). This type of analysis sometimes leads to decreased degrees of freedom because some participants do not have sufficient data to contribute to all the analyses. To minimize this, we collapsed data from pairs of blocks; that is, in the analyses, data from Blocks 1 and 2 were treated as being from one block.

Our expectation was that, for both training and test, sequence learning should increase the difference scores of the experimental conditions compared with their respective controls because it should increase RT and errors on inconsistent trials and decrease RT and errors on consistent trials. Figure 1 shows these scores during training and test. We examine the training data first.

Considering the control conditions, the pattern here can be attributed to sequential effects. The Group Same control condition's scores were positive for both errors and RTs because the Group Same subsequences were trial orders that participants found relatively easy to respond to, but note that the Group Same experimental scores were higher still, suggesting that training on the Group Same subsequences led to sequence learning. The Group Different control condition scores were approximately the mirror image of those for the Group Same control condition, with any (small) deviation from this attributable to random variation. Thus, their inconsistent minus consistent scores are negative because their subsequences were those that participants did not find easy to perform. Once again, the experimental condition's scores were higher than their controls (much higher in the case of errors), suggesting that training on the Group Different subsequences resulted in sequence learning.

We can assess the main effect of condition (experimental vs. control) for Group Same in RTs, $F(1, 48) = 68.96$, $p < .001$, and errors, $F(1, 57) = 7.97$, $p < .008$; similarly for Group Different in RTs, $F(1, 48) = 72.93$, $p < .001$, and errors, $F(1, 54) = 44.59$, $p < .001$. In all cases, there was good evidence of superior performance in the experimental conditions, and the difference between experimental and control conditions increased over blocks in both groups, further supporting the conclusion that the participants learned at least some of the statistical structure of these sequences during the course of training. The interaction between condition and block that supports this assertion was significant in both RTs, $F(16, 768) = 5.31$, $p < .001$, and errors, $F(16, 864) = 3.82$, $p < .001$ in Group Different. The same Condition × Block interaction was also significant in Group Same for RTs, $F(16, 768) = 5.93$, $p < .001$, and errors, $F(16, 912) = 1.76$, $p < .04$. Hence, both groups exhibited reliable learning of the sequences during our experiment, but the difference between experimental and control conditions over blocks also differed for Group Different and Group Same, as the Group × Condition × Block interaction was significant in the RTs, $F(16, 1536) = 1.98$, $p = .012$, but not in the errors, $F(16, 1776) = 0.812$, $p = .67$. Inspection of the graphs in Figure 1 suggests that this reflects the somewhat faster learning in Group Different in Session 1 but not in Session 2.

The test data are based on performance on the pseudorandom blocks composing the last five blocks of the experiment and are shown on the right of Figure 1. Once again, there is evidence of learning in that RT differences for both Group Different and Group Same experimental participants were significantly higher than those for controls. Specifically, there was a main effect of condition (experimental vs. control) in Group Different's RTs, $F(1, 62) = 61.46$, $p < .001$, and errors, $F(1, 62) = 41.08$, $p < .001$, and a main effect of condition in Group Same's RTs, $F(1, 62) = 36.39$, $p < .001$, although the error data just fail to reach significance, $F(1, 62) = 3.97$, $p = .051$. Both RTs and errors show numerically better sequence learning expressed on test for Group Different than Group Same, with a significant interaction between condition and group in the errors, $F(1, 124) = 4.76$, $p = .031$, accompanied by a nonsignificant trend in the RTs, $F(1, 124) = 1.99$, $p = .16$. Using Brown's (1975) procedure for combining analyses that are not independent, we can generate an overall χ^2 for the RT and error measures of 7.65, with 2.3 *df*, which has an associated $p < .05$. Thus, we can conclude that the participants trained on the Group Different subsequences performed better on

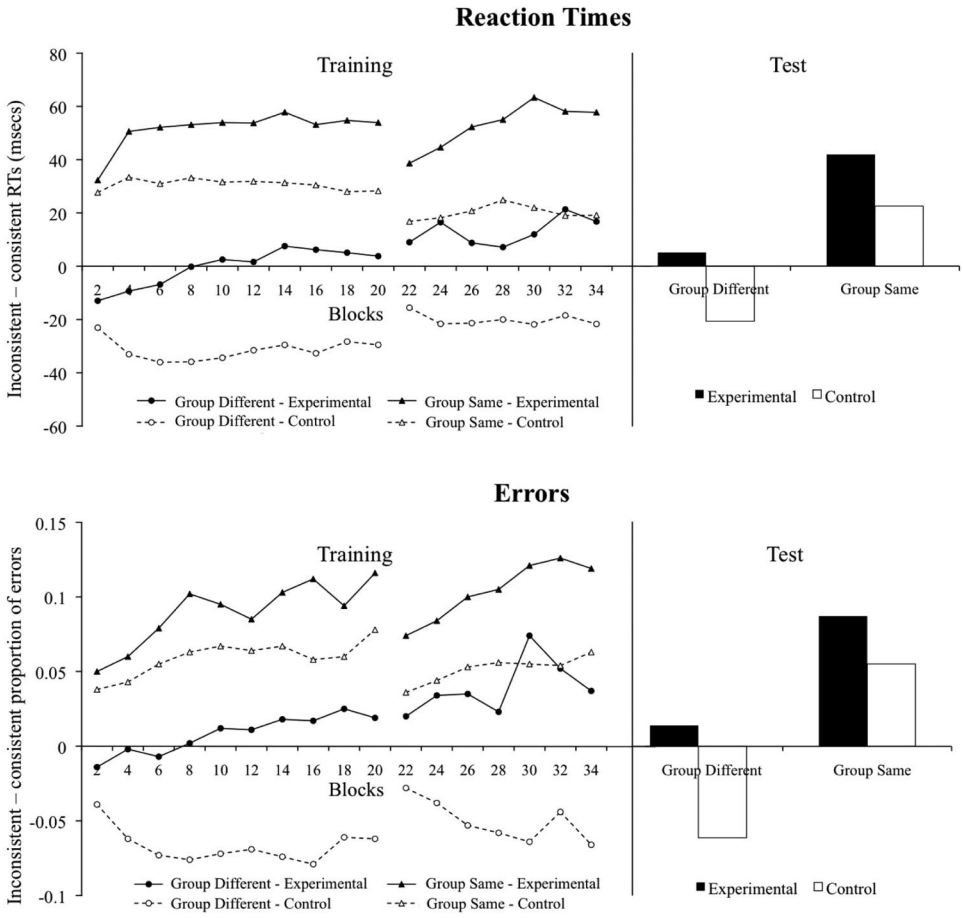

Figure 1. Reaction time difference scores (top half) and proportion of errors differences scores (bottom half) during training (left panels) and on test (right panels). Blocks are given in block pairs (i.e. 2 means the average of Blocks 1 and 2), and there was a break between Blocks 20 and 21 (two different sessions at least 2 hr apart). Only 14 blocks (seven block pairs) are shown in the second session as the last five blocks were used as the test phase.

test than those trained on the subsequences given to Group Same, with no hint of any speed–accuracy trade-off.

Discussion

Our findings are quite clear and straightforward. Under incidental conditions, participants trained on the subsequences experienced by Group Different, namely XXY, XYY, YXX, YYX, learned more than those trained on the subsequences experienced by Group Same, that is, XXX, XYX, YXY, YYY. This resulted in better performance in a final test phase that controls for any possible sequential effects by using the same pseudorandom sequences for all groups and conditions. The effect was not large, but it was entirely reliable and not compromised by any issues of speed versus accuracy.

We are also able to offer some reassurance that our results were indeed obtained under incidental conditions. By the end of the study, no participants were able to tell us what the subsequences in the experiment were. This fits well with the claims made by Jones and McLaren (2009) for this paradigm under the same conditions, and reassures us that our instructions and procedures placed our participants in this experiment on a footing similar to that in the earlier study. Given this, we can now inquire whether the model that fit the data in Jones and McLaren, the augmented SRN (Cleeremans & McClelland, 1991), is also able to fit these results.

We attempted to model the experiment using the augmented SRN with exactly the parameters given in Jones and McLaren (2009). In this model (shown in Figure 2; please disregard the two input units at the bottom labeled "Next trial" for this purpose), the two possible stimulus locations were each assigned an input unit (units labeled "Current trial") and were activated as appropriate on each trial, following the same sequential structure, number of trials, blocks, and sessions as used for participants. These fed-forward to a set of hidden units, which in turn activated two output units (labeled "Prediction of next trial"), with the activity of units in both layers determined by the logistic activation function (Rumelhart, Hinton, & Williams, 1986). At the end of each trial, the activations of the hidden units were copied via one-to-one feedback connections to a set of "context units" on the input layer (labeled as "Copy of last trial's hidden units"). This recurrence is the essence of Elman's (1990) SRN architecture. The output units corresponded to the two possible stimulus locations, and

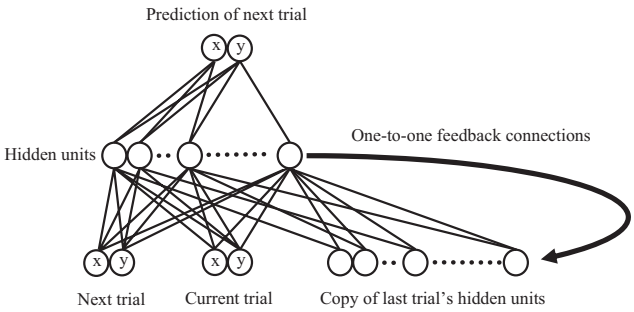

Figure 2. Model architecture for the augmented simple recurrent network (SRN) used in Jones and McLaren (2009, Figure 7), with the addition of two extra input units (corresponding to the two response locations) labeled "Next trial." See text for a description of the model and its operation.

their activation represented the model's prediction of the identity of the next trial.

The augmented SRN has a similar architecture to the SRN, but differs in its feed-forward connections. Where the SRN has modifiable network connections driven by a single learning parameter, the augmented SRN has two components to these connections, fast and slow (not shown separately in the figure), which have higher and lower learning rates, respectively. The fast components also decay by half their value at each time step, a feature adopted by Cleeremans and McClelland (1991) to help account for the robust short-term priming effect observed in their data. In addition, the augmented SRN includes a set of response units to capture the effect that responding to the previous trial has on the current trial, whereas the SRN lacks this component. Cleeremans and McClelland (1991), Cleeremans (1993), and Jiménez, Mendez, and Cleeremans (1996) have shown that this model can capture the detailed pattern of SRT data.

Using the stimulus location on the following trial as a training target, we modified the weights (both fast and slow in the augmented SRN) determining the strength of each connection between units according to the back-propagation algorithm (Rumelhart et al., 1986). As in the simulations conducted by Jones and McLaren (2009), the networks used had 20 hidden units and a slow learning rate parameter of 0.4, with the fast weights having a learning rate 1.33 times larger. Thirty-two networks were run in each of the four cells of the experimental design. The results are shown in the top panels of Figure 3. If we compare the simulations to the empirical results, then first impressions are that there is some correspondence, especially given that we have not fit the model to the data by varying parameters. The augmented SRN does quite well in predicting the basic pattern during training in that experimental and control groups are appropriately placed with respect to one another, and at least some of the trends observed in our data are also present in the simulations.

But the crucial point is that there are areas of significant disagreement between our data and the model predictions. Most important are those in the test data (top right panel of Figure 3). Both groups demonstrate a significant main effect of condition: For Group Different, $F(1, 62) = 290.65$, $p < .001$, and for Group Same, $F(1, 62) = 1383.87$, $p < .001$, but the significant interaction between condition and group, $F(1, 124) = 11.61$, $p < .002$, confirms reliably greater sequence learning in Group Same for the augmented SRN, which is the contrary pattern of results found in our empirical study. Thus, we have to reject the augmented SRN, or at least this version of it, as an adequate model for our data. Furthermore, with this architecture, we have, to date, been unable to find a set of parameters for the augmented SRN that will allow it to correctly predict the ordering of Groups Different and Same on test despite an extensive search over the parameter space for the model, suggesting that, as it stands, it cannot model our data. In any case, we can conclude that the version of the model that was successful in modeling the data in Jones and McLaren (2009) is demonstrably falsified by our results.

This outcome is surprising as the augmented SRN is our benchmark model of sequence learning and coped very well with the pattern of results that we obtained in Jones and McLaren (2009). Our initial response was to revert to the version of the SRN adopted in Spiegel and McLaren (2006). This proved capable of simulating more learning in Group Different than Group Same, but only at the cost of losing the ability to adequately simulate the overall pattern for our data set on test, and it does not provide as good a fit to the Jones and McLaren data. After considering various other modifications of the network, we finally realized that our simulation of this task using the augmented SRN was unrealistic in the following way. Recall that our architecture was feed-forward (and recurrent) with two input units set according to which location (i.e., left or right) was designated as a response on the trial just past and context units whose activation was set by the hidden unit activations on the previous trial as well. What we had failed to include in our model was anything that represented the stimulus—the filled circle—that always occurred just before a response was made. This was because this stimulus completely specified the response, and it would have seemed odd to include something so directly predictive when we were interested in the ability of the network to learn the sequential contingencies in play, not learn that when the left circle filled, it was to produce a left response! But our participants would have been exposed to just such a contingency, and so would have had the opportunity to learn about it. In some sense, this captures the idea of some automatization occurring in the course of experience that takes over from the instruction to press the corresponding key when one of the circles fills. Hence, we included these inputs in our new network (the input units labeled "Next trial" in Figure 2), but, because the circle only fills just before the response, we gave it a relatively low weighting in our model.[1] With this addition, we were able to successfully reparameterize the augmented SRN to produce our pattern of results in this experiment and still generate the pattern of results found by Jones and McLaren (2009). A typical set of simulation results for the current experiment is shown in the bottom panels of Figure 3. These simulations were run using the two additional input units, and activation of the unit corresponding to the response required on the to-be-predicted trial was set to 0.1. Input activations corresponding to the response units that had been acti-

[1] The low weighting was intended to reflect the fact that the timing was suboptimal for an associative network learning to predict the next response required, but this is a completely separate issue to the cue's predictiveness, which, of course, was 100%.

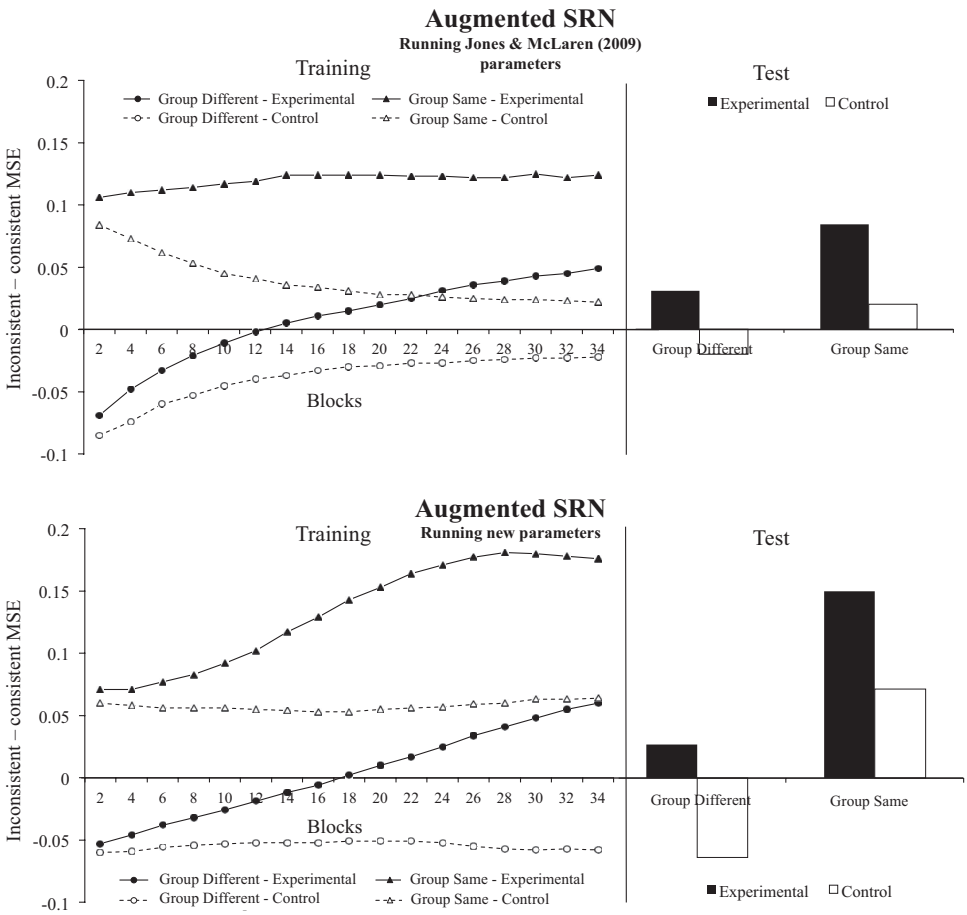

Figure 3. Mean square error (MSE) difference scores during training and test for the simple recurrent network (augmented SRN; top) and the revised augmented SRN (bottom). We did not attempt to simulate the delay between Blocks 20 and 21. Otherwise, it is laid out exactly as Figure 1.

vated on the previous trial were set to 0.75, and the context unit activations were set to 1.3 times the hidden unit activations from the previous trial. These parameters modulate the relative weightings of the contributions to learning from the different inputs and the context units, and were deliberately chosen to allow us to simulate our data. The learning rate parameters for the fast and slow weights were set to 0.5 and 0.2, respectively, and other parameters were the same as those used in Jones and McLaren.

The training data shown on the left of the bottom panel of Figure 3 again have the different conditions/groups in their appropriate relative positions. The control groups are once again approximate mirror images, and now produce a somewhat more stable pattern of performance over time. Learning proceeds at approximately the same rate in the two experimental groups (it is slightly faster overall for Group Different). The real data of interest are those at test, however, and here the pattern corresponds closely to that in our empirical data. There is a main effect of group, $F(1, 124) = 1682$, $p < .001$, a main effect of condition, $F(1, 124) = 717$, $p < .001$, and importantly, an interaction between group and condition, $F(1, 124) = 4.08$, $p < .05$, such that the difference between experimental and control conditions for Group Different is significantly greater than that for Group Same. In other words, Group Different sequences are better learned than Group Same sequences, although the effect is relatively small (roughly 10%) compared with the main effects. This is very much the pattern, and the power, that we observed in the empirical data we report here. We can also confirm that this model captures the pattern of subsequence learning observed in Jones and McLaren (2009), and also predicts that it is the difference in performance between XXX and the other sequences used in 2009, XYY, YXY, and YYX, that will be the easiest to detect. It would appear, then, that this revised model is a candidate to be our new benchmark for modeling sequence learning with this task under incidental conditions.

Why does this revised model succeed where the standard augmented SRN failed? The new version of the model differs from the old version in both including a more accurate representation of the stimulus conditions in the experiment (an unambiguously good thing) and in possessing more free parameters as a consequence of this modification. Is its success simply a consequence of greater

flexibility in fitting the data contingent on this increase in free parameters? We believe this is not the right explanation because when we simulated the Jones and McLaren (2009) data with the new model, we did not vary the parameters at all, implying that our success (the fit was actually better than in the Jones and McLaren paper) is unlikely to represent "overfitting" of the data. Instead, we believe that the inclusion of the units corresponding to the circle stimulus on the current trial is the critical feature making the difference, and if we take this out of the simulation, leaving everything else unchanged, then the model reverts to predicting that Group Same should perform better on test than Group Different. The change made is in some sense minor, but is also important. It does not represent a change in the algorithms used in the augmented SRN, or even in its basic architecture, but it is a departure from conventional simulation practice as far as the SRT task is concerned. As far as we are aware, no one else modeling sequence learning involving this type of task includes the current stimulus as an input in the model—but clearly it matters. Why has it typically been left out when simulating sequence learning in the SRT task? We think the reason is fairly straightforward: Researchers were (are) interested in sequence learning, and putting in this input contributes nothing to learning about the contingencies between sequences of events in this task, it just allows stimulus–response (S-R) learning. This type of information cannot actually assist in learning sequential structure, and would be expected to be the same in both experimental and control groups, and hence controlled for when assaying sequence learning. But, although this S-R learning cannot produce a difference between experimental and control groups in its own right, we now realize it can modulate our ability to learn about sequential structure, and so influence the size of that difference. It does this, we believe, via cue competition, which itself varies as a function of the local temporal sequence of events experienced. To see this, consider as an example the sequence LLL (three left responses required in a row). On the first two trials of this sequence, the augmented SRN will learn (transiently) that a left on trial $n - 1$ predicts a left on trial n, and also the association between the left stimulus and the required left response will have been incremented. The first effect makes learning the "LL is followed by L" structure difficult, because it partially blocks it (this is the explanation of why LLL is learned poorly under implicit conditions given in Jones & McLaren, 2009). But the second effect, incrementing the S-R learning, also contributes to blocking learning of the "LL is followed by L" contingency. Learning LLR is, relatively speaking, easier because the R is surprising in terms of the "L predicts L" transient learning, and, in this case, the R stimulus to R response association will not have just received two increments. Thus, the increments to the S-R associations are more of a problem for some subsequences (which turn out to be those in Group Same) than others, and so contribute to Group Same learning the sequential structure more slowly. For this reason, we cannot simply disregard this S-R information any more on the basis that it will be the same for both experimental and control groups. Clearly, we should not disregard this aspect of the task in any case if we are to hold to the view that these models are automatic in their operation and learn about all elements of the perceived stimulus array. If we are to believe that this is a real psychological model of associative learning, then, because the circles in the experiment signal which response to make, there must be something in the model that represents this, and the model will inevitably learn about this 100% reliable contingency. But now we can see that if we do neglect this aspect of the stimulus conditions, then our simulations do not match the empirical data, which is, in some sense, a rather encouraging outcome for this modeling approach.

Are there any discrepancies between our model's simulation and the data? One obvious discrepancy arises when comparing the training data from the model and our empirical results. The change from Session 1 to Session 2 is not captured by the model, but this is hardly surprising as we have no way of representing it in the model at present. Perhaps the worst aspect of the model as it stands is that it has Group Same learning faster than Group Different in the middle section of the graph, whereas the empirical data show the reverse, but even here it is difficult to know whether this is a reliable difference, and the analysis is compromised by both the effect of a change of session and a lack of power. We acknowledge that it is also possible to criticize the current model for being rather slow to learn. We speculate that a modification of backpropagation, APECS (Le Pelley & McLaren, 2001; McLaren, 1993, 1994, 2011), instantiated in a recurrent architecture (Jones, Le Pelley, & McLaren, 2002) might be the way forward here in terms of improving learning and incorporating aspects of memory that might permit the session effect to be captured. But, for now, on the basis of the data available in the literature on sequence learning in humans and contained in this article, the revised augmented SRN is our benchmark model of sequence learning.

References

Brown, M. B. (1975). A method for combining non-independent, one-sided tests of significance. *Biometrics, 31,* 987–992. doi:10.2307/2529826

Cleeremans, A. (1993). *Mechanisms of implicit learning*. Cambridge, MA: MIT Press.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General, 120,* 235–253. doi:10.1037/0096-3445.120.3.235

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14,* 179–211. doi:10.1207/s15516709cog1402_1

Jiménez, L., Mendez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 948–969. doi:10.1037/0278-7393.22.4.948

Jones, F. W., Le Pelley, M. E., & McLaren, I. P. L. (2002). The APECS-SRN: Towards a model of SRT sequence learning. *Proceedings of the 2002 international joint conference on neural networks* (pp. 692–696). Piscataway, NJ: IEEE.

Jones, F. W., & McLaren, I. P. L. (2009). Human sequence learning under incidental and intentional conditions. *Journal of Experimental Psychology: Animal Behavior Processes, 35,* 538–553. doi:10.1037/a0015661

Le Pelley, M. E., & McLaren, I. P. L. (2001). Retrospective revaluation in humans: Learning or memory? *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 54*(B), 311–352.

McLaren, I. P. L. (1993). APECS: A solution to the sequential learning problem. *Proceedings of the XVth Annual Convention of the Cognitive Science Society* (pp. 717–722). Hillsdale, NJ: Erlbaum.

McLaren, I. P. L. (1994). Representation development in associative systems. In J. A. Hogan & J. J. Bolhuis (Eds.), *Causal mechanisms of behavioural development* (pp. 377–402). Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511565120.018

McLaren, I. P. L. (2011). APECS: An adaptively parameterized model of associative learning and memory. In E. Alonso & E. Mondragón (Eds.), *Computational neuroscience for advancing artificial intelligence: Models, methods and applications* (pp. 145–164). Hershey, PA: IGI Global.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 318–362). Cambridge, MA: Bradford Books.

Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance, 11,* 598–616. doi:10.1037/0096-1523.11.5.598

Spiegel, R., & McLaren, I. P. L. (2006). Associative sequence learning in humans. *Journal of Experimental Psychology: Animal Behavior Processes, 32,* 150–163. doi:10.1037/0097-7403.32.2.150

Received December 10, 2011
Revision received December 21, 2012
Accepted January 9, 2013 ■