

Modelling the Detailed Pattern of SRT Sequence Learning

F.W. Jones (psm1fj@surrey.ac.uk)¹

University of Cambridge, Department of Experimental Psychology,
Downing Street, Cambridge, CB2 3EB, UK.

I.P.L. McLaren (iplm2@cus.cam.ac.uk)

University of Cambridge, Department of Experimental Psychology,
Downing Street, Cambridge, CB2 3EB, UK.

Abstract

Any viable model of serial reaction time (SRT) sequence learning needs to be able to capture the relative extents to which participants learn different SRT sub-sequences. However, previous attempts to empirically establish a relative pattern of this sort have failed to fully control for 'sequential-effects', or have not satisfactorily ruled-out the influence of speed-accuracy trade-off. In this paper it is shown that, when sequential-effects are controlled for in a two-choice SRT task, participants learn those sub-sequences that end in an alternation better than those that end in a repetition. It is further demonstrated that, at least for the parameters investigated, a buffer network, a Jordan network, an SRN, an augmented SRN, and an AARN are unable to account for this pattern.

Introduction

Since its introduction (Lewicki, Czyzewska & Hoffman, 1987; Nissen & Bullemer, 1987) the serial reaction time (SRT) task has proved a popular way to assay human sequence learning. In the most common version of this task a stimulus can appear in one of several locations on a computer screen, each of which has a corresponding response key. Whenever a stimulus appears, participants simply have to press the appropriate key as quickly and accurately as possible. Crucially, usually unknown to the participants, the order of locations in which the stimulus is presented follows a sequence, at least for the majority of the time.

Typically, participants' reaction times (RTs) on trials that are consistent with this sequence become significantly faster, either than their RTs on occasional inconsistent probe trials, or than the RTs of a control group that is only exposed to a (pseudo-)random ordering (e.g. Anastasopoulou & Harvey, 1999). This is taken to imply that the participants have learnt at least part of the sequence, and that they are using this information to prepare for the stimulus or response on the next trial.

The widespread use of this task makes it a priority to develop an adequate model of human SRT performance. A particularly stringent test of any candidate model is to capture the relative extents to which participants learn different parts (sub-sequences) of an SRT sequence. Unfortunately, however, nearly all of the patterns of sub-

sequence learning that have been reported in the literature are potentially confounded by 'sequential-effects' (cf. Anastasopoulou & Harvey, 1999; Shanks & Johnstone, 1999).

'Sequential-effects' can be defined as the influence that the previous series of stimulus/response locations has on the participant's current response, in a choice-RT task in which the trial order is (pseudo-)random (e.g. Soetens, Boer & Hueting, 1985). To illustrate, in a two-choice task at a response-stimulus-interval (RSI) of 50ms, participants tend to respond faster when the stimulus appears in the same location as it did on the prior trial (Soetens et al., 1985).

Therefore, in order to more accurately assess SRT sequence learning, the performance of the sequence group and pseudo-random control group need to be compared on a 'test-phase' in which they are both exposed to the same order of trials; the assumption being that the two groups should manifest the same sequential-effects when responding to identical trial orders, allowing the difference between the groups to be used as an index of sequence learning.

However, the majority of previous SRT studies have failed to control for sequential effects in this manner (e.g. Nissen & Bullemer, 1987), calling into question the apparent patterns of sub-sequence learning they report.² Furthermore, the two previous studies that have adequately controlled for sequential effects, namely Anastasopoulou and Harvey (1999) and Shanks and Johnstone (1999), have not reported error data in sufficient detail to rule out the possibility that participants traded speed and accuracy between different sub-sequences.

Therefore it was necessary to carry out a new experiment in order to establish a relative pattern of SRT sub-sequence learning against which models of the SRT task could be tested.

Experiment

The experiment comprised a five session two-choice SRT task, with an Experimental Group that was trained on four sub-sequences and a Control Group that was exposed to a pseudo-random ordering for the same number of trials. To control for sequential effects, following training all

² Some authors have attempted to reduce the influence of sequential-effects by removing trials that they feel are particularly contaminated (e.g. Jimenez, Mendez & Cleeremans, 1996). However, it is not clear that this is sufficient.

¹Now at: University of Surrey, Department of Psychology (Clinical Psychology PsychD), Guildford, Surrey GU2 7XH, UK.

participants were 'tested' on pseudo-random orderings, with the difference between the two groups being taken as an index of sequence learning.

Method

There is not the space to give full methodological details, but these will be provided in Jones and McLaren (in prep.).³

Participants, Stimuli and Apparatus The participants were 24 Cambridge University students and members of the public, whose ages ranged between 18 and 47. 16 were randomly assigned to the Experimental Group and 8 to the Control. This difference in numbers did not cause any reliable differences in variance between the groups; therefore the subsequent ANOVAs are valid. The experiment was run on a Macintosh LCIII, and the stimulus comprised a circle 1.9 cm in diameter which could appear 2.2 cm to the right or left of the centre of the screen. The two response keys were spatially compatible with these locations.

Trial Order The trials were batched into blocks of 120, and are described here in terms of Xs and Ys. The assignment of X and Y to right and left was counterbalanced across participants. As illustrated in Figure 1, the building blocks of the 'sequence-blocks' were four sub-sequences, each three trials long; specifically XXX, XYY, YYX and YXY. These sub-sequences conform to the rule 'if the first and second trials are the same then the third is an X, but if they differ it is a Y'. Sequence-blocks were constructed by concatenating 10 of each of these sub-sequences in a random order. Thus every third trial in a sequence-block was predictable on the basis of the previous two. Participants were not informed about the special status of third trials, and the RSI following third trials was the same as that following other trials.⁴ It is also worth noting that when the properties of these sequence-blocks are considered on a trial-by-trial basis, then approximately two-thirds of the trials are consistent with the rule, and on average the four sub-sequences occur equally frequently. (This was confirmed by the Monte-Carlo method.)




Figure 1: Examples of portions of a sequence-block (top) and a pseudo-random-block (bottom), with some of the constituent triplets highlighted.

³ This paper will not include the simulation work presented here.

⁴ The sequence-blocks were constructed so that the third trials had a special status because the original purpose of the experiment was to compare the pattern of sub-sequence learning between this incidental condition and one in which participants were told about the third trials (see Jones & McLaren, in prep.).

'Pseudo-random-blocks' comprised 5 of each of the sub-sequences (XXX, XYY, YYX and YXY) and 5 of each of the complementary triplets with alternate endings (XXY, XYX, YYY and YXX), randomly concatenated.

Design Each session comprised 20 blocks. The first 10 blocks of session one and the last 10 blocks of session five constituted the pre- and post-training 'test-phases', and were formed of pseudo-random-blocks for both groups. The 80 block 'training-phase' in between these comprised sequence-blocks for the Experimental Group and pseudo-random-blocks for the Control Group.

Procedure The procedure followed that described for the standard SRT task in the introduction, with the addition that participants were paid a performance bonus, designed to encourage them to be as fast and accurate as possible. The RSI was 500ms.

Results and Discussion

Due to space constraints, only those analyses most relevant to the subsequent modelling are reported here, but see Jones and McLaren (in prep.) for full details. Furthermore, only the RT results have been presented because the same trends were observed in the errors, ruling out speed-accuracy trade-off.

To assay sub-sequence learning, the data from the third trials of the first five blocks⁵ of the post-training test-phase were divided up with respect to the identity of the previous two trials (i.e. XX, XY, YY and YX). They were then further divided on the basis of whether the third trial was consistent with the sub-sequence begun by the previous two trials (e.g. XXX) or inconsistent with it (e.g. XXY). For the Control Group consistency was a dummy variable, since they had never been trained on the sub-sequences.

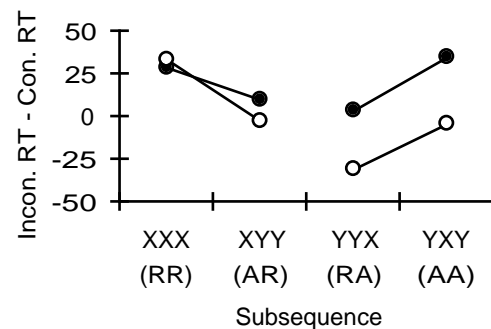


Figure 2: The mean difference scores from the third trials of the first half of the post-training test phase. Filled circles = Experimental Group. Open circles = Control Group.

A mean RT for each of the resulting eight trial types was calculated, on a per participant basis. To reduce variability, RTs were not taken from error trials or trials following an error. Then, for each sub-sequence, a difference score was constructed by subtracting the consistent RT from the inconsistent one. This measure was used because it subtracts

⁵ An analysis of the learning curves suggested that extinction had set in by blocks 6-10 (see Jones & McLaren, in prep.).

out individual variability in baseline RT. The mean difference score for each sub-sequence, per group, are shown in Figure 2.

Following the convention adopted in the two-choice sequential effects literature (e.g. Soetens et al., 1985), the four sub-sequences were coded in terms of whether each trial was a repetition (R) or an alternation (A) of the previous trial (i.e. XXX=RR, XYY=AR, YYX=RA and YXY=AA). The difference scores were then analysed using an ANOVA with the factors: group (experimental vs. control), first position in the alternation-repetition sub-sequence code (A or R), and second position (A or R).

This demonstrated that the Experimental Group's difference scores were significantly larger those of the Control Group ($F(1,22)=31.77, p<.01$). This suggests that the participants in the Experimental Group learnt at least some of the sequential contingencies, because sequence learning should produce slower RTs on inconsistent trials and faster RTs on consistent trials, and thus a larger inconsistent minus consistent score.

The group X first-position (reading from left to right) interaction was not significant ($F<1$), nor was the group X first-position X second-position interaction ($F<1$). However, the group X second-position interaction was reliable ($F(1,22)=5.14, p<.05$). According to a simple effects analysis, this interaction arose because the Experimental Group's difference scores were significantly greater than Control for those sub-sequences that ended in an alternation ($F(1,22)=18.12, p<.01$) but not for those that ended in a repetition ($F(1,22)=0.17, p>.1$).

Thus it would appear that differential learning of the sub-sequences occurred, with participants only learning those sub-sequences that ended in an alternation within the time of the experiment.

Furthermore, this pattern would appear not to be contaminated by a floor effect, because if just the results from the inconsistent trials are analysed then the same trend is observed. (The expression of learning cannot be masked by RT being at floor on inconsistent trials because learning should slow responding on such trials.)

Finally, when data from all trials was included in the analysis, rather than just from the third trials, a similar pattern of sub-sequence learning was observed. Thus there was no evidence to suggest that the participants in the Experimental Group had learned about the special status of third trials (i.e. that they were always consistent during training). Rather, participants appear to have learnt the 2/3 contingencies that are present on a trial by trial basis (see method section).

In summary, in a two choice SRT task in which sequential effects are controlled for, it appears that people learn those sub-sequences that end in an alternation better than those that end in a repetition. But can the current neural network models of SRT sequence learning capture this pattern?

Modelling

To address this question, a variety of different neural network models of SRT sequence learning were presented with an

exact analogue of the human task. These models all include some form of memory for previous trials. It is by associating a combination of the contents of this memory and a representation of the current trial with the identity of the next trial that they learn the sequential contingencies. As the models learn they become more accurate at predicting the next trial's identity, and consequently the mean squared error (MSE) of their prediction reduces.

Therefore, in the following simulations the MSE has been taken as an index of the models' RT to the stimulus that it is attempting to predict; the rationale being that a low MSE indicates that a model accurately expects the location of the next stimulus and so it should react more quickly to it when it occurs. While some authors transform the MSE using a decision rule or mechanism (cf. Cleeremans, 1993), this was avoided in order to prevent the properties of the models from potentially being obscured by the addition of an extra process.

In all the following simulations a localist code, with one unit for X and one for Y, was used to represent the input to the models and their predictions. The trial order was generated in exactly the same way as in the experiment, and for each type of model both an Experimental and Control Group were run. Finally, to assay sequence learning the models' MSEs from the first half of the post-training test-phase were analysed in the same way as the human data.

The different types of networks studied will now be considered in turn.




Figure 3: A buffer network.

The Buffer Network

The architecture of a buffer network is shown in Figure 3. In this model a memory is simply implemented by presenting as input to the network all the necessary elements of the sequence. Thus, the network's input comprises a representation of the current trial and a decayed (by half) representation of the previous trial. Trials prior to this are not included since the contingencies in the sequence-blocks are only depended upon the previous two trials. The architecture includes a layer of hidden units, because these enable it to learn non-linearly separable mappings (Rumelhart, Hinton & Williams, 1986), and the weights are updated using the backpropagation algorithm (Rumelhart, Hinton & Williams, 1986). For more details concerning buffer networks see Cleeremans (1993, pp. 141-143). (Note, a momentum term was not employed).

Thirty-two separate buffer networks, each with 4 hidden units, were run on an analogue of the experiment. Sixteen formed the Control Group and 16 the Experimental Group.

Each network 'subject' had a different set of randomly initialised weights. And, the learning rate was set to 0.8, since this value meant that the networks had learnt two of the four sub-sequences by the post-training test-phase, like the human participants had.




Figure 3: The Buffer networks' difference scores from the third trials of the first half of the post-training test phase. Filled circles = Experimental Group. Open circles = Control.

The mean post-training MSE difference scores are shown in Figure 3. The data were analysed using a group (2 levels) by sub-sequence (4 levels) ANOVA. This revealed that the Experimental Group's scores were significantly higher than Control ($F(1,30)=17.23, p<.01$), indicating that they had acquired some of the sequential contingencies. Moreover, there was evidence of differential sub-sequence learning (group X sub-sequence: epsilon corrected $F(2,57)=5.41, p<.05$); with a simple-effects analysis demonstrating that subjects had reliably learnt subsequences XYY/AR and YYX/RA (respectively $F(1,30)=22.05, p<.01$; $F(1,30)=5.84, p<.05$), marginally learnt XXX/RR ($F(1,30)=4.06, .1>p>.05$), and not learnt YXY/AA ($F(1,30)=0.24, p>.1$). The model made similar predictions with 20 hidden units.

However, these predictions differ from the pattern expressed by human participants (i.e. stronger learning of those sub-sequences that ended in an alternation). Therefore, the buffer network was unable to model the human data, at least with 4 or 20 hidden units.




Figure 4: A Jordan network.

The Jordan Network

The architecture of a Jordan network (Jordan, 1986), modified to model the SRT task (Cleeremans, 1993, pp. 139-141), is illustrated in Figure 4. In this model the identity of previous trials are not explicitly represented on separate pools of input units. Rather, a memory of the sequence is implemented by adding to each input unit a self-recurrent connection of fixed weight, which is 0.5 in this case. As with the buffer network the variable weights are updated using the backpropagation algorithm.




Figure 5: The Jordan networks' difference scores from the third trials of the first half of the post-training test phase.

To provide reliable results, 34 network subjects were run in both groups. Each network had 4 hidden units and a learning rate of 0.5. The results are shown in Figure 5. An ANOVA revealed reliable evidence of sequence learning ($F(1,66)=42.92, p<.01$) and of differential sub-sequence learning ($F(3,198)=17.28, p<.01$). A follow up simple-effects analysis demonstrated that the networks had learnt all the sub-sequences except XXX/RR (XXX/RR: $F(1,66)=1.47, p>.1$; XYY/AR: $F(1,66)=39.60, p<.01$; YYX/RA: $F(1,66)=20.23, p<.01$; YXY/AA: $F(1,66)=33.38, p<.01$). Therefore, with 4 hidden units, the Jordan network achieved a closer match to the human data than the buffer network, but it still was unable to capture the full detail of the pattern. When the number of hidden units was increased to 20, the networks' results were less similar to the human data.




Figure 6: An SRN.

The Simple Recurrent Network (SRN)

Figure 6 illustrates the architecture of the SRN, which was developed by Elman (1990). In the SRN a memory of previous items in the sequence is implemented by providing

as additional input to the network its own hidden unit activations from the previous trial. As with the other models, the weights are updated by the backpropagation algorithm.

With 4 or 20 hidden units, and a range of large learning rates (0.5, 0.8, and 1.0), the SRN was incapable of reliably learning any of the sequential contingencies by the post-training test-phase ($p > .1$). However, when sixteen 40 hidden unit SRNs were run in each group (learning rate 0.5), an ANOVA revealed that reliable sequence learning did occur ($F(1,30)=13.41, p < .01$). For the means see Figure 7.




Figure 7: The SRNs' difference scores from the third trials of the first half of the post-training test phase.

Moreover, there was reliable evidence of differential sub-sequence learning ($F(3,90)=8.91, p < .01$), and according to a simple-effects analysis the network subjects had learnt all the sub-sequences except YXY/AA (XXX/RR: $F(1,30)=4.97, p < .05$; XYY/AR: $F(1,30)=13.82, p < .01$; YYX/RA: $F(1,30)=9.37, p < .01$; YXY/AA: $F(1,30)=1.59, p > .1$). However, this pattern deviates from the advantage for those sub-sequences that ended in an alternation seen in the experiment.

The Augmented SRN

In order to allow the SRN to capture a higher proportion of the variance in their SRT data, Cleeremans and McClelland (1991) made two modifications to it, producing the Augmented SRN. First, to model the short term priming effect of previous learning episodes, weights and biases were divided into both a fast and slow component.




Figure 8: The Augmented SRNs' difference scores from the third trials of the first half of the post-training test phase.

Second, to capture response repetition effects the model's prediction of the next response was made dependent upon a decaying trace of previous responses as well as the SRN's output. For more details see Cleeremans and McClelland (1991).

To determine whether this model could capture the experimental results, 16 networks with 4 hidden units were run in each group. The learning rates were 1.3 and 1.0 for the fast and slow weights respectively. The other parameters were as described in Cleeremans and McClelland (1991). Given that the SRN's output activations are transformed in this model, the transformed activations were employed in the calculation of the MSE.

The results are shown Figure 8. While the trends may appear to be very small, an ANOVA did reveal that the networks had reliably learnt at least some of the contingencies ($F(1,30)=15.64, p < .01$) and that this learning varied across the sub-sequences ($F(3,90)=38.64, p < .01$). A subsequent simple-effects analysis demonstrated that XXX/RR and XYY/AR were learnt ($F(1,30)=28.17$ and $58.01, p < .01$), while the Experimental Group's scores were significantly lower than Control for YYX/RA and YXY/AA ($F(1,30)=51.41$ and $11.63, p < .01$). Thus the pattern of subsequence learning was the nearly opposite to that expressed by human subjects. Nor did increasing the number of hidden units to 20 substantially improve the situation.




Figure 9: An AARN.

The Autoassociative Recurrent Network (AARN)

The AARN, which was developed by Maskara and Noetzel (1993), is shown in Figure 9. It differs from a standard SRN in that the network is taught not only to predict the next trial, but also to predict the pattern of activity across the input layer. This is implemented by including an extra pool of output units, with each of these units corresponding to a particular unit in the input layer. The target for one of these new output units is the activity of its corresponding input unit on the same trial.

Fifty AARNs, with four hidden units and a learning rate of 0.8, were run in each group. This large number of subjects was required to produce significant results. It is also worth noting that only the activations of the two output units that comprised the network's prediction for the next trial were included in the calculation of the MSE.




Figure 10: The AARNs' difference scores from the third trials of the first half of the post-training test phase.

The networks' mean MSE difference scores are shown in Figure 10. An ANOVA revealed that the Experimental Group had learnt at least some of the sequential contingencies ($F(1,98)=47.29$, $p<.01$) and that the sub-sequences had been learnt differentially ($F(3,294)=9.77$, $p<.01$). According to a simple effects analysis, sub-sequences XYY/AR and YYX/RA were learnt ($F(1,98)=39.81$ and 11.59 , $p<.01$), while the remaining two were not ($F(1,98)=0.04$ and 1.37 , $p>.1$). This pattern contrasted with the stronger learning of those sub-sequences that ended in an alternation observed in the human data. Furthermore, changing the number of hidden units to 20 did not improve the situation.

Dominey's Network

The ability of Dominey's (1995) model to capture the human data has also been examined. However, thus far we have been unable to find a set of parameters that enables the model to display any evidence of sequence learning on the task.

General Discussion

To summarise, previous attempts to establish a relative pattern of SRT sub-sequence learning are flawed because either they fail to control for sequential effects or do not present error data in sufficient detail. In a two-choice SRT task designed to overcome these flaws, it was found that people learn those sub-sequences that end in an alternation better than those that end in a repetition. Surprisingly, however, a buffer network, the Jordan network, the SRN, the augmented SRN and the AARN all appear unable to capture this result, at least with the parameters investigated.

A critic could argue that our human data might be beyond the scope of these models because it may reflect 'explicit' rather than 'implicit' learning. However, other work in our laboratory suggests that people show a different pattern when learning explicitly, namely they find XXX/RR the most salient sub-sequence. Therefore, it would seem that current neural network models of SRT sequence learning will probably at least need to be modified in order to accommodate the data presented here.

Acknowledgements

This research was funded by an MRC research studentship awarded to F.W. Jones. The authors would thank Stephen Monsell, David Shanks, Nick Yeung and two anonymous reviewers for their helpful comments.

References

- Anastasopoulou, T., & Harvey, N. (1999). Assessing sequential knowledge through performance measures: the influence of short-term sequential effects. *Quarterly Journal of Experimental Psychology*, *52A*, 423-448.
- 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.
- Dominey, P.F. (1995). Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biological Cybernetics*, *73*, 265-274.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-211.
- Jimenez, 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.
- Jordan, M.I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: LEA.
- Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 523-530.
- Maskara, A., & Noetzel, A. (1993). Sequence recognition with recurrent neural networks. *Connection Science*, *5*, 139-152.
- Nissen, M.J., & Bullemer, P. (1987). Attentional requirements of learning: evidence from performance measures. *Cognitive Psychology*, *19*, 1-32.
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart and J.L. McClelland (Eds.), *Parallel Distributed Processing. Volume 1*. Cambridge, MA: Bradford Books.
- Shanks, D.R., & Johnstone, T. (1999). Evaluating the relationship between explicit and implicit knowledge in a sequential reaction time task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1435-1451.
- 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.