Consciousness and Cognition 21 (2012) 413–421

Contents lists available at SciVerse ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

Short Communication

Out of control: An associative account of congruency effects in sequence learning

Tom Beesley [a,*], Fergal W. Jones [b,c], David R. Shanks [a]

[a] Division of Psychology and Language Sciences, University College London, London, UK
[b] Department of Applied Psychology, Canterbury Christ Church University, Kent, UK
[c] Sussex Partnership NHS Foundation Trust, Sussex, UK

ARTICLE INFO

Article history:
Received 16 May 2011
Available online 16 November 2011

Keywords:
Control
Sequential congruency
Associative learning
Connectionism

ABSTRACT

The demonstration of a sequential congruency effect in sequence learning has been offered as evidence for control processes that act to inhibit automatic response tendencies (Jiménez, Lupiáñez, & Vaquero, 2009) via unconscious conflict monitoring. Here we propose an alternative interpretation of this effect based on the associative learning of chains of sequenced contingencies. This account is supported by simulations with a Simple Recurrent Network, an associative (connectionist) model of sequence learning. We argue that the control- and associative-based accounts differ in their predictions concerning the magnitude of the sequential congruency effect across training. These predictions are tested by reanalysing data from a study by Shanks, Wilkinson, and Channon (2003). The results support the associative learning account which explains the sequential congruency effect without appealing to control processes (either conscious or unconscious).

1. Introduction

Events within the world are often contingent on one another and, at times, these events will occur in a rapid sequence that requires a complex series of actions to be performed. For example, when driving a car, to manoeuvre at a junction it is necessary to learn a structured sequence of movements involving the brake and clutch pedals, the gear change, the steering wheel, and the indicator levers. A change in the state of one of these elements (e.g., a failure to engage a gear) has implications for the execution of subsequent events (e.g., the clutch should not be released; the accelerator should not be engaged). Thus, learning about the occurrence of individual elements is insufficient: in order to respond effectively in these situations one must learn a sequence of temporally spaced contingent events. In this article, we consider evidence that purports to show that control processes, operating via unconscious conflict monitoring, act to inhibit automatic response tendencies. We present an alternative account of this evidence that makes no appeal to such control processes. This is an important issue concerning the nature of top-down control in the execution of behaviour.

In order to study sequence learning in the laboratory, researchers have employed (amongst other things) the serial reaction time (SRT) task. The SRT task is a computerised choice-response task in which a target stimulus appears in one of a set number of locations on the display (often four). Participants are given response keys that correspond directly to these positions and are asked to respond to the location of the target as quickly and as accurately as they can. Response times are typically around 500 ms, and the next target appears shortly after the response to the last target. The movement of the target stimulus is determined by an underlying sequence, although the instructions given to participants make no reference to this

* Corresponding author. Address: Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London WC1H 0AP, UK.
E-mail address: t.beesley@ucl.ac.uk (T. Beesley).

1053-8100/$ - see front matter
doi:10.1016/j.concog.2011.09.020

feature of the task. Commonly a probabilistic sequence is used, such that "training" transitions (T) occur frequently, whilst "control" transitions (C) occur less frequently. The difference in reaction times (RTs) on these two types of trial therefore provides a measure of the extent to which the sequence has been learnt. This task captures the probabilistic and unsupervised way in which sequence learning is likely to occur in many real-world settings.

In one of the most influential papers in the field of sequence learning, Cleeremans and McClelland (1991; see also Cleeremans, 1993) provided a demonstration of how a recurrent neural network (connectionist) model employing the back-propagation algorithm, the Simple Recurrent Network (hereafter SRN; Elman, 1990), can offer considerable insight into the performance of participants in sequence learning situations. The success of the SRN is shown in its continual application to current data throughout the implicit learning literature (e.g., Dienes, 1993; Dienes, Altmann, & Gao, 1999; Jiménez, Méndez, & Cleeremans, 1996; Jones & McLaren, 2009; Spiegel & McLaren, 2006) and, whilst additional mechanisms are required to explain certain patterns of data (e.g., Beesley & Le Pelley, 2010), it continues to provide a compelling description of human sequence learning in purely associative terms.

The SRN (see Fig. 1) is a connectionist learning model comprising four layers of processing units, interconnected with associative weights, which act to modulate the strength of the signal passing between connected units. The *input units* reflect the current location of the target stimulus. Based on this input the model makes a prediction as to the location of the stimulus on the next trial of the sequence, by passing activation to the *output units*. Input and output units are connected indirectly via a set of *hidden units*. These units reflect the model's internal representation of the input pattern and play a crucial role in its ability to learn sequenced patterns. After each trial, the model copies the activation on the hidden units to a set of context units. By doing this, it is able to use its own internal state of activation as a further source of input on subsequent trials, which provides a rich context that reflects the previous temporal contingencies within the sequence. A formal description of the SRN is provided in Appendix A. For the present purposes it is sufficient to appreciate that the model is able to learn *higher-order contingencies*. For example, in the sequence 121432413423, every possible first-order transition (e.g., 12,13,14) occurs equally often (i.e., once) and so learning of higher-order contingencies within the sequence is necessary for the next element of the sequence to be predicted accurately (e.g., 14 is followed by 3 but 24 is followed by 1). The SRN is able to use the context layer to distinguish between otherwise ambiguous states (e.g., 14-3 vs. 24-1) and so can simulate learning of these sequences. Learning in the SRN is driven by prediction error, which is used to adjust weights within the network, such that when the model experiences similar sequenced transitions in the future, its predictions regarding the next target element are more accurate. Hence, knowledge of the sequence is represented by the evolving weight structure within the model.

To summarise, the SRN is a purely associative model of learning. The expression of sequence knowledge within the model is reflected only in its output activation and, importantly, these responses are produced automatically by a chain of activation in the model; there is no sense in which the model can facilitate or inhibit responses by "top-down" control processes.

Jiménez, Lupiáñez, and Vaquero (2009) recently examined whether control processes act to modulate the expression of sequence knowledge in the SRT task and concluded that control processes do indeed play a role in performance. In their task participants were trained on a second-order sequence presented in a probabilistic fashion, such that sequenced transitions were chosen from the trained sequence with a high probability (.9) or from an alternative control sequence with a low probability (.1). Two sequences were used: 121432413423 and 323412431421. If a participant was trained on the first sequence, after experiencing transition 12, position 1 occurred with a high probability (a training transition, T), whilst 4 (selected because it is the successor to 12 in the second sequence) appeared with a low probability (a control transition, C). Thus, a faster reaction time after a T transition (1), and a slower reaction time after a C transition (4), indicated that participants possessed knowledge about the second- and perhaps higher-order relationships in the training sequence.

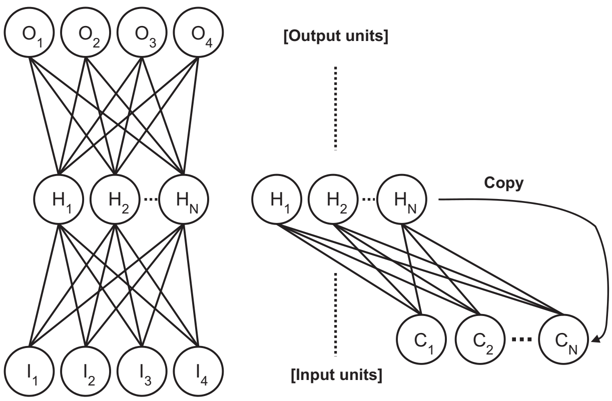

Fig. 1. Simple Recurrent Network (SRN). Left panel: a back-propagation network with input units (I_1–I_4), hidden units (H_1–H_N), and output units (O_1–O_4). *Right panel*: the context loop of the SRN, which allows the network to make a copy of its internal activation at the end of each trial to a set of context units (C_1–C_N), allowing the internal activation to be used as input to the network on subsequent trials. The number of hidden units can be varied.

By chaining T and C sequences, Jiménez et al. were able to examine whether the experience of a C transition led participants to "...modulate the expression of sequence learning..." (p. 693) on subsequent trials. Indeed, Jiménez et al. (2009) found that the difference in the latency of responding on T and C transitions was smaller following a previous C transition, compared to the difference following a previous T transition (i.e., TC–TT was greater than CC–CT). In two additional experiments, originally reported in Jiménez, Vaquero, and Lupiáñez (2006), participants were trained on second-order conditional sequences, like those described above, with control sequences substituted in full with a probability of .2. That is, at the end of each presentation of the entire training sequence, the training sequence would be presented again with a probability of .8 or the control sequence would be presented with a probability of .2. In reanalysing the data from these studies, Jiménez et al. (2009) found that the results again suggested a sequential congruency effect in that the facilitation in responding on TT sequence chains compared to TC chains was larger than the facilitation for CT over CC. Although these data were consistent with the sequential congruency hypothesis, Jiménez et al. (2009) conceded that the transitions of entire sequences could lead to strategic processes operating once participants had noticed a long sequence of unfamiliar transitions had occurred.

Jiménez et al. (2009) interpreted these sequential congruency effects as arising from the operation of control processes within a 'conflict monitoring model' such as that proposed by Botvinick, Braver, Barch, Carter, and Cohen (2001). According to this model, the experience of an incongruent C transition would give rise to an increase in "conflict", triggering a control process to limit the activation of future incorrect response tendencies. Thus the normal facilitation of responding (i.e., fast responses) on T transitions compared to C ones will be attenuated after a preceding C transition as a result of the latter inducing suppression of activation via top-down control. Since the sequential congruency effect did not differ between intentional and incidental learning conditions, and since it occurred on sequenced transitions for which participants showed no conscious awareness, Jiménez et al. (2009) argued that the conflict monitoring mechanism operates at an implicit level, rather than as an overt decision process.

The control interpretation of the sequential congruency effect challenges the associative account offered by the SRN, in that this model has no mechanism for controlling the spread of activation; the SRN has no top-down control process that acts to modulate responses when an unexpected event is experienced. In this article, we will examine an alternative account of the sequential congruency effect based purely upon the learning of higher-order associations between elements of a sequence. Specifically, the experience of a control transition necessarily results in a decrease in the amount of sequenced information that can be used to make predictions for subsequent trials. In fact, Jiménez et al. (2009) made reference to the possibility of this alternative account (p. 699) of the sequential congruency effect. In the current article, we examine this account in detail by simulating the experimental design with the SRN and assessing whether the model predicts a sequential congruency effect. We then present an analysis of the sequential congruency effect in data originally reported by Shanks et al. (2003), in order to test a specific prediction from the SRN regarding the emergence of the sequential congruency effect during training. This work therefore addresses the important question of whether implicit control processes play a role in the expression of learnt action sequences.

2. An SRN simulation of the sequential congruency effect

The model was trained with sequences constructed in the same manner as those reported in Jiménez et al. (2009), presented in 12 blocks of 120 trials. Two second-order sequences were used as training and control sequences (121432413423 and 323412431421). To start the sequence a first-order transition was selected at random. On all subsequent trials the next location in the sequence was selected on the basis of the previous two elements of the sequence. A transition from the training sequence was selected with a probability of .9 and from the control sequence with a probability of .1. For example, after randomly selecting 34 as the first two elements, 2 would be selected with a probability of .9, whilst 1 would be selected with a probability of .1. The next transition would then be selected in the same manner on the basis of either 42 or 41, and so on. As in the design employed by Jiménez et al. (2009), in blocks 10 and 11 the probabilities of generating training and control transitions were set at .2 and .8, respectively, before returning to the previous probabilities in a final block (12). During this "transfer" period, Jiménez et al. found that participants were still sensitive to the training sequence (albeit showing slightly weakened effects). Furthermore, Jiménez (personal communication, August 2011) noted that the sequential congruency effect was not present over these transfer blocks. We sought to examine whether this same pattern of results would be revealed in the behaviour of the SRN.

One hundred independent simulations were run, initialised with randomly determined starting weights within the range −.5 to +.5. Bias units were used on both the hidden and output units and were initialised within the same range. On each trial of the simulation a Luce Ratio was calculated – the activation of the target output unit divided by the total activation of all output units – where high ratios reflect accurate performance (Luce, 1959). Model performance is presented as 1 minus the Luce Ratio in order to mimic the decreases in reaction times commonly seen in the sequence learning performance of participants. All model data were calculated by averaging the results of the 100 simulations.

We explored a number of combinations of values for the learning rate and the number of hidden units used. A summary of the results from these parameter combinations is shown in Table 1. Four descriptive statistics are presented for each parameter combination. "Learning" provides a score of the basic learning effect, calculated as the difference in responding on T and C transitions during the two blocks immediately prior to the transfer phase (blocks 8 and 9), whilst "Learning (t)" provides the same data for the transfer phase. Thus the impact of the transfer block on the expression of learning would be seen in a

Table 1

Learning rate						
	.1		.3		.5	
Number of hidden units						
50	Learning	.01	Learning	.08	Learning	.17
	Learning (*t*)	.03	Learning (*t*)	.10	Learning (*t*)	.16
	Seq. Con.	−.01	Seq. Con.	−.03	Seq. Con.	−.04
	Seq. Con. (*t*)	−.01	Seq. Con. (*t*)	−.02	Seq. Con. (*t*)	−.02
100	Learning	.04	Learning	.20	Learning	.30
	Learning (*t*)	.06	Learning (*t*)	.17	Learning (*t*)	.20
	Seq. Con.	−.02	Seq. Con.	.02	Seq. Con.	.07
	Seq. Con. (*t*)	−.02	Seq. Con. (*t*)	.00	Seq. Con. (*t*)	.04
150	Learning	.08	Learning	.27	Learning	.36
	Learning (*t*)	.09	Learning (*t*)	.19	Learning (*t*)	.22
	Seq. Con.	−.04	Seq. Con.	.14	Seq. Con.	.17
	Seq. Con. (*t*)	−.01	Seq. Con. (*t*)	.02	Seq. Con. (*t*)	.02

Note: Results of the parameter exploration for the simulations with the Simple Recurrent Network. 'Learning' is the average difference in performance on control and training trials across blocks 8 and 9. 'Learning (*t*)' is the average difference in performance on control and training trials across the transfer blocks 10 and 11. 'Seq. Con.' is the average sequential congruency effect (see text) across blocks 8 and 9. 'Seq. Con. (*t*)' is the average sequential congruency effect across the transfer blocks 10 and 11. All values reflect differences in Luce Ratio scores, calculated as the activation of the target output unit divided by the total activation of all output units.

smaller numerical score for Learning vs. Learning (*t*). "Seq. Con." provides the average sequential congruency effect over blocks 8 and 9 (i.e., [TC–TT] − [CC–CT]), whilst "Seq. Con. (*t*)" provides the equivalent sequential congruency calculation over the transfer blocks. Positive values reflect sequential congruency effects for that particular parameter set. The absence of an effect for Seq. Con. (*t*) would be shown by a zero value for this statistic (i.e., [TT–TC] = [CT–CC]).

It is clear from Table 1 that whilst a learning rate of .1 was able to establish a learning effect with sufficient hidden units, there was no evidence for any impairment of performance on the transfer blocks. Likewise, when only 50 hidden units were used, similar results were found even with higher learning rates. The remaining four parameter combinations (learning rates of .3 and .5, and 100 and 150 hidden units) showed strong learning effects and evidence of impaired, but positive, transfer performance. Furthermore, all four combinations showed a positive sequential congruency effect over training blocks 8 and 9 (Seq. Con.) and a reduced sequential congruency effect over the transfer blocks. In choosing the parameter combination to explore in more detail, we selected the set which produced the largest sequential congruency effect in proportion to the basic training effect (a learning rate of .3 and 150 hidden units). This parameter set also showed a substantial reduction in the sequential congruency effect over the transfer blocks (from .14 to .02) in line with the observation by Jiménez (personal communication, August 2011).

Fig. 2A presents the SRN simulation of the basic learning effect (the difference in performance on control and training transitions), which develops across the course of the experiment.[1] During the transfer blocks 10 and 11, the model demonstrates an impairment in the learning effect, although there is still a difference between C and T over these blocks.

Fig. 2B shows the data averaged across blocks 1–9 and as a function of both the status of the current trial and the previous trial. It is clear from Fig. 2B that the SRN predicts an interaction effect: the difference in responding on training and control transitions is smaller if that transition was preceded by a control transition, compared to when it was preceded by a training transition. That is, the model demonstrates a sequential congruency effect.

Fig. 2C shows the data from Jiménez et al. (2009), collapsed across all training blocks and also across both the incidental and intentional learning conditions ($N = 39$). The sequential congruency effect is shown by the significant interaction between the factors of current and previous transition.

3. Reanalysis of Shanks et al. (2003), Experiment 3

Whilst the preceding analysis shows that the SRN is able to demonstrate an interaction effect which resembles the sequential congruency effect, it is clear that the pattern of data is not identical to that observed by Jiménez et al. (2009). Specifically, Jiménez et al. (2009) failed to find a benefit for current T transitions over current C transitions following a previous C transition (i.e., CT = CC). The SRN is unable to simulate this absence of an effect on trials occurring after a previous C transition (except of course for circumstances in which the model fails to learn at all). We believe the model cannot show a difference between TT and TC but not between CT and CC, simply because the former effect is dependent on learning of second- *and* third-order transitions in the sequence, whilst the latter is only dependent on learning of the second-order

[1] Note that all figures contain error bars representing the standard error of the mean. However, with the exception of Fig. 3H, the error bars for the model simulations are so small as to be imperceptible.

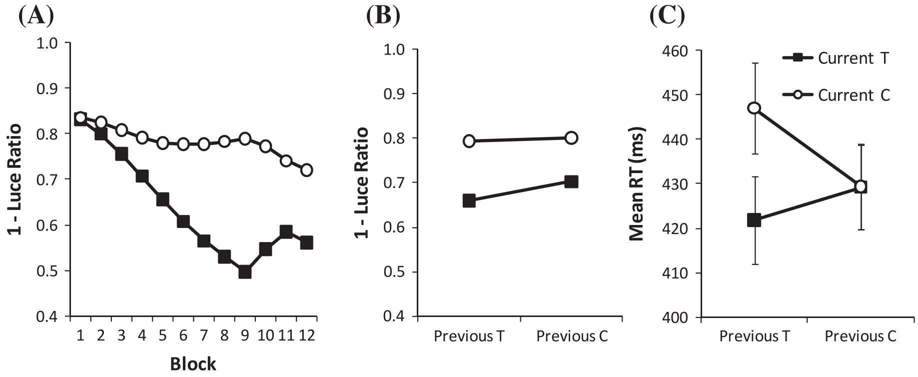

Fig. 2. The performance of the SRN (A and B) and the empirical data (C) from the experiment reported in Jiménez et al. (2009). *Current T*: responses to training transitions; *Current C*: responses to control transitions; *Previous T*: responses following a previous training transition; *Previous C*: responses following a previous control transition. The Luce Ratio is calculated as the activation of the target output unit divided by the total activation of all output units. This figure was reproduced using data sent personally by Luis Jimenez, which originally appeared in Jiménez et al. (2009).

transitions in the sequence. It is likely that the SRN will learn more slowly about third-order transitions in the sequence because these transitions will occur less frequently, since they rely on two successive high-probability choices in the sequence generation procedure, whilst second-order transitions rely on just one. Thus – despite its ability to predict an interaction – the absence of any difference in response latencies for T and C transitions following C transitions (Fig. 2C) is a challenge to the SRN model and it is important to ask whether this is a replicable pattern. We consider this issue below.

This description of learning in the SRN reveals a key prediction of the model with respect to the sequential congruency effect, namely that the effect should emerge as training progresses, being absent early on in training and becoming stronger as training progresses. In contrast, the control account does not predict an absence of the effect early in training. By this account, if sequence knowledge can be expressed, then an experience of conflict in the sequence will cause a suppression of knowledge on the subsequent transition. Therefore, according to the current control account, as soon as sequence knowledge can be expressed in the task a sequential congruency effect should be observed.[2]

To formally test this prediction against that of the SRN we reanalysed data originally reported in Experiment 3 of Shanks et al. (2003). This experiment used a very similar training procedure to the study reported by Jiménez et al. (2009), which therefore allows for an identical analysis to be conducted to assess the presence of a sequential congruency effect. Moreover, the dataset contains data from 100 participants, which means it is particularly useful for two reasons. Firstly, the increased sample size allows us to conduct a more detailed analysis of how the sequential congruency effect develops over time. Secondly, this larger sample size will help to re-assess the pattern of data following a previous C transition (i.e., whether there is any difference between response latencies to CC and CT transitions).

The training procedure was similar to that reported in Jiménez et al. (2009), in which second-order conditional sequences were presented in a probabilistic fashion. The two sequences were structurally identical to those reported in Jiménez et al. (2009) and participants were randomly assigned one of these sequences as the training sequence and the other as the control sequence. Sequence substitution was conducted in the same manner as in Jiménez et al. (2009) except that the probability of a T transition appearing was .85 rather than .9. Participants were trained with 12 blocks of 100 trials (this study did not contain a transfer phase). As before, we assessed the mean response latencies for four trial-types: TT, TC, CT, and CC.

The reanalysed data from Shanks et al.'s (2003) Experiment 3 are shown in the top row of Fig. 3. The corresponding data from an SRN simulation (100 simulated subjects, learning rate = .3, hidden units = 150) are shown in the bottom row. Fig. 3A shows the data averaged across the whole of the experiment, as was presented for the study by Jiménez et al. (2009) shown in Fig. 2C. These data were subjected to an analysis of variance (ANOVA) with within-subject factors of current and previous transition. This analysis revealed a main effect of the current transition, $F(1,99) = 59.99$, $p < .001$, indicating that responses to current T transitions were faster than to current C transitions. There was also a main effect of the previous transition, $F(1,99) = 6.03$, $p < .05$, indicating that participants' responses were slower overall if they had experienced a C transition immediately prior to the current transition. The interaction between these two factors was also significant, $F(1,99) = 12.60$, $p < .01$, suggesting that the performance benefit for a current T transition over a current C transition was greater when the previous transition was a T transition; the data replicate the sequential congruency effect. This interaction

[2] It is possible that the control account could be extended in some way to predict the emergence of the sequential congruency effect over time. For example it could be extended to include a threshold which determines the influence inconsistent events have on learnt behaviour. That is, until the learning of the sequence has reached some predetermined level, control processes would fail to influence the expression of learning and so a congruency effect would not be observed in the early stages of training.

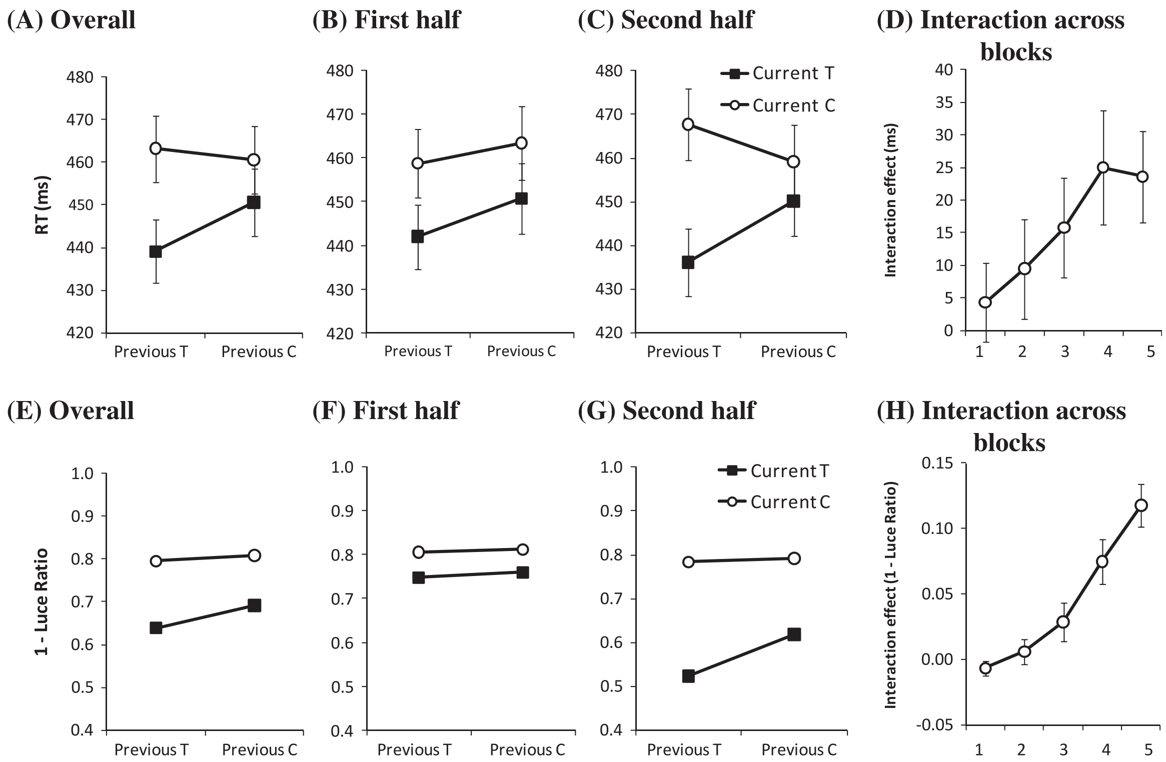

Fig. 3. (A) The empirical data from Shanks et al. (2003) analysed in terms of the sequential congruency effect. (B) The data from the first half of the training. (C) The data from the second half of training. (D): A plot of the interaction effect (i.e., the extent to which the difference in responding for training and control trials was greater following a previous training trial) as a function of training. (E) The performance of the SRN on the same experimental design, across the whole (E), the first half (F), and the second half (G) of training. (H) The interaction effect in the SRN data, plotted as a function of training. *Current T*: responses to training transitions; *Current C*: responses to control transitions; *Previous T*: responses following a previous training transition; *Previous C*: responses following a previous control transition.

effect was driven largely by a difference in RTs on current T transitions (TT < CT), $t(99) = 8.04$, $p < .001$, whilst the difference across the current C transitions was not significant (TC = CC), $t < 1$.

As described above, the SRN predicts that the sequential congruency effect should emerge as training progresses in the task. To assess this, we conducted the above analysis for the first 600 trials and the last 600 trials of the task. For the first half, shown in Fig. 3B, an ANOVA revealed a main effect of the current transition, $F(1,99) = 27.79$, $p < .001$, and of the previous transition, $F(1,99) = 7.89$, $p < .01$. Importantly, there was no interaction between these effects, $F < 1$; there was no evidence for a sequential congruency effect in the first half of training, despite clear evidence that the sequence had been learnt. For the second half of the training session, shown in Fig. 3C, an ANOVA revealed a main effect of the current trial, $F(1,99) = 60.47$, $p < .001$, but there was no main effect of the previous trial, $F < 1$. Importantly, for this phase of training the interaction between these two factors was highly significant, $F(1,99) = 16.74$, $p < .001$; there was strong evidence for a sequential congruency effect in the second half of training. As before, this interaction was driven largely by a decrease in RTs on current T transitions (TT < CT), $t(99) = 6.99$, $p < .001$, whilst the difference in RTs for current C transitions was not significant (TC = CC), $t(99) = 1.63$, $p = .11$. A combined analysis which includes the factor of training period (first half vs. second half) revealed a significant three-way interaction, $F(1,99) = 7.83$, $p < .01$, demonstrating that the sequential congruency effect was greater in the second-half of the experiment than in the first-half.

Panel D shows the size of the congruency effect as a function of training duration. For this analysis the data from the 1200 trials were split into five 240 trial blocks. Within each block the size of the congruency effect was calculated by subtracting the difference between CC and CT from the difference between TC and TT. Positive scores therefore reflect how much greater the expression of sequenced knowledge was after a previous T transition than after a previous C transition. The data show that the congruency effect increases as a function of training.

The SRN predictions for this experimental design are presented in the bottom row of Fig. 3. It is clear that the SRN reproduces the sequential congruency effect in the data averaged across the whole of training (Fig. 3E), as was demonstrated in the simulations of the Jiménez et al. (2009) procedure. More importantly, early on in training (Fig. 3F) the model predicts a main effect of the current transition, but no interaction of this effect with the previous transition. That is, the model has learnt about the sequence by this point in training, but its failure to acquire knowledge about the higher-order contingencies by this stage means that there isn't a sequential congruency effect. In contrast, in the second half of training (Fig. 3G), the model is clearly sensitive to contingencies of greater length, and there is an interaction between the states of the previous and the current transitions on responding. The development of the congruency effect in the model data is also shown in Fig. 3H, where the congruency effect is plotted as a function of training block (calculated in the same way as for the data in 3D). These results demonstrate how the sequential congruency effect emerges over the course of training and mirrors the participants' data, suggesting that the effect is based in the learning of longer chains of sequence contingencies, the representations for which require longer periods of training to develop, with no necessity to impute control processes (implicit or otherwise).

We noted above that there was one pattern in the empirical data that the SRN appeared unable to capture. Specifically, in the data presented by Jiménez et al. (2009) shown in Fig. 2C, it was observed that the difference in responding between T and C transitions was only apparent following a previous T transition and not after a previous C transition. We noted that the SRN cannot predict this effect. However, in the analysis of Shanks et al.'s (2003) data, there was an effect following a previous C transition, $t(99) = 2.78$, $p < .01$, with 62 participants showing an effect in this direction, $Z = 2.3$, $p < .05$. Given that Shanks et al.'s experiment contained more than double the number of participants, one possibility is that the study by Jiménez et al. (2009) lacked sufficient power to detect this effect. A post hoc power analysis using the effect size of .28 from Shanks et al.'s experiment found that the power to detect an effect after a previous C transition (across all between-subject conditions) in Jiménez et al.'s (2009) study ($N = 39$) was only .40 ($\alpha = .05$). Interestingly, for the much larger effect following a previous T transition, which was shown to be significant in both the Jiménez et al. (2009) and Shanks et al. analyses, taking the effect size of 1.12 from the Shanks et al. study, we observe that Jiménez et al. (2009) had power of .99 to detect this effect in their experiment. Thus, one conclusion we may draw from this analysis is that Jiménez et al. (2009) had sufficient power in their analysis to detect the larger effect present after a previous T transition, but low power to detect the weaker effect following a previous C transition. Overall, then, there is little in the data that the SRN cannot accommodate.

4. Discussion

In this article, we have considered two alternative accounts of the mechanisms responsible for learning and responding to sequenced events. According to the control-based, or conflict monitoring account, put forward by Jiménez et al. (2009), the experience of an unexpected event in a learnt sequence automatically triggers a more cautious response to subsequent events. A more cautious response here refers to the inhibition of expression of the participant's learnt tendencies in the task, such that the benefit of trained over non-trained transitions is weakened or even eliminated. An alternative account is that the sequential congruency effect emerges simply from the process of learning more complex contingencies within the sequence. Specifically, we suggested that the SRN – which can use an internal context loop to encode a representation that reflects the rich context of temporal contingencies within a sequence – can predict a sequential congruency effect. According to this model, the experience of a control transition in a sequence results in a weakening of the model's ability to predict the next element of the sequence; the experienced sequence being a degraded version of the learnt sequence.

Using the same sequences as were used to show a sequential congruency effect by Jiménez et al. (2009), we were able to demonstrate a sequential congruency effect in the SRN. Furthermore, by reanalysing the data from a very similar experiment conducted by Shanks et al. (2003), we were also able to show how the sequential congruency effect develops over the course of training in humans, a result which the SRN also accurately predicts. Thus the SRN provides a precise account of the sequential congruency effect without the need to appeal to an additional conflict monitoring mechanism. Whilst we have not addressed the issue directly here, the nature of the distributed memory representation within the SRN means that this account of the sequential congruency effect is at least consistent with an automatic learning process.

In conclusion, we have shown here that the sequential congruency effect can be explained by the gradual acquisition of more complex sequence knowledge by a purely associative system, and that there is insufficient evidence to necessitate additional control processes (whether conscious or unconscious) to explain the expression of sequence knowledge. On this last point, it should be noted that the control account offered by Jiménez et al. is not in itself an account of sequence learning, but rather a theoretical account of a particular aspect of sequenced behaviour or the expression of sequence knowledge. Thus, even if one was to accept the control account of the sequential congruency effect, it would still be necessary to adopt the additional representational framework of a model such as the SRN in order to provide a full account of sequence learning. We would argue therefore that not only does the SRN offer a better fit to the observed data than the control-based account (i.e., with respect to the emergence of the sequential congruency effect over time) but it also continues to offer the most parsimonious account of sequence learning and behaviour.

Acknowledgment

The authors would like to thank Luis Jiménez for his helpful comments on an earlier version of this article.

Appendix A

The SRN was based on the network described by Elman (1990): a backpropagation multi-layer network with a recurrent loop. On each trial, the activation of one input unit was set to 1 and all other input units were set to 0. The activation was then fed forward to the hidden units by calculating the sum of all products of input unit activations and their respective connection strengths with each hidden unit. Thus, for hidden unit h:

$$in_h = B_h + \sum_{i=1}^{I} w_{hi} \cdot a_i \tag{A.1}$$

where in_h is the input for hidden unit h; B_h is the bias associated with hidden unit h; w_{hi} is the weight of the connection between hidden unit h and input unit i; a_i is the activation of input unit i; and I is the total number of input units. This input is then transformed into an activation value for hidden unit h, by the activation function given in Rumelhart and McClelland (1986):

$$a_h = \frac{1}{1 + e^{-in_h}} \tag{A.2}$$

The input to, and activation of, the output units is calculated in much the same way as that of the hidden units, such that for output unit o:

$$in_o = B_o + \sum_{h=1}^{H} w_{oh} \cdot a_h \tag{A.3}$$

$$a_o = \frac{1}{1 + e^{-in_o}} \tag{A.4}$$

where in_o is the input for output unit o; B_o is the bias associated with output unit o; w_{oh} is the weighted connection between output unit o and hidden unit h; a_h is the activation of hidden unit h; and H is the total number of hidden units in the network.

When applied to sequence learning, the target output on each trial is the next element in the sequence. The accuracy of the model in selecting the next element was calculated as the activation of the target output unit divided by the total activation of all output units, commonly referred to as the Luce Choice Ratio (hereafter LCR; Luce, 1959).

Target values for 'active' and 'inactive' stimulus positions were set at .9 and .1, respectively. Values of .9 and .1 are used as targets, rather than 1 and 0, as the latter cannot be reached without infinitely large weights, and so effectively cannot be achieved (Rumelhart and McClelland, 1986).

Following each response made by the network, the error of each output unit is back-propagated through the network to update the weights between each layer of units. Error terms for output and hidden units were calculated as follows:

$$\delta_o = (t_o - a_o) \cdot (1 - a_o) \cdot a_o \tag{A.5}$$

$$\delta_h = \left(\sum_{o=1}^{o=O} \delta_o \cdot w_{oh} \right) \cdot (1 - a_h) \cdot a_h \tag{A.6}$$

where δ_o and δ_h refer to the error on output unit o and hidden unit h, respectively, and t_o is the activation target for output unit o, while O is the total number of output units. These errors were then used to update the weights and biases in the network using the generalised delta rule (Rumelhart and McClelland, 1986):

$$\Delta w_{oh} = LR \cdot \delta_o \cdot a_h \tag{A.7}$$

$$\Delta w_{hi} = LR \cdot \delta_h \cdot a_i \tag{A.8}$$

$$\Delta B_o = LR \cdot \delta_o \tag{A.9}$$

$$\Delta B_h = LR \cdot \delta_h \tag{A.10}$$

where LR denotes the learning rate parameter and Δ means 'change in'. At the end of each trial, the activation on the hidden units is copied to a set of 'context' units. Context units function as an additional set of input units and are treated identically to input units in equations A1 and A8.

References

Beesley, T., & Le Pelley, M. E. (2010). The effect of predictive history on the learning of sub-sequence contingencies. *Quarterly Journal of Experimental Psychology, 63*, 108–135.

Botvinick, M., Braver, T., Barch, D., Carter, C., & Cohen, J. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108*, 624–652.

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

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

Dienes, Z. (1993). Computational models of implicit learning. In D. C. Berry, & Z. Dienes (Eds.), *Implicit learning: Theoretical and empirical issues*. Hove: Erlbaum.

Dienes, Z., Altmann, G. T. M., & Gao, S.-ji. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science, 23*, 53–82.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14*, 179–211.

Jiménez, L., Lupiáñez, J., & Vaquero, M. (2009). Sequential congruency effects in implicit sequence learning. *Consciousness and Cognition, 18*, 690–700.

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.

Jiménez, L., Vaquero, J. M. M., & Lupiáñez, J. (2006). Qualitative differences between implicit and explicit sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 32*, 475–490.

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.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1): Foundations*. Cambridge, MA: MIT Press.

Shanks, D. R., Wilkinson, L., & Channon, S. (2003). Relationship between priming and recognition in deterministic and probabilistic sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 248–261.

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