Trajectory Effects in a Novel Serial Reaction Time Task

George Kachergis[1], Floris Berends,

Roy de Kleijn, & Bernhard Hommel
[1]george.kachergis@gmail.com
Institute of Psychology / LIBC, Leiden University
Leiden, the Netherlands

Abstract

The serial reaction time (SRT) task, which measures how participants' keypress responses speed up as a repeating stimulus sequence is learned, is popular in implicit and motor learning research, and may help us understand the basic learning mechanisms underlying the acquisition of complex skills (e.g., riding a bike). However, complex action sequences are not simple stimulus-response chains, but rather require representing sequential context in order to learn. Moreover, human actions are continuous, temporally-extended movements that are not fully measured in the discrete button presses of the SRT task. Using a novel movement adaptation of the SRT task in which spatial locations are both stimuli and response options, participants were trained to move the mouse cursor to a continuous sequence of stimuli. We replicate the Nissen and Bullemer (1987) RT results with the trajectory SRT paradigm and show sequential context effects–predictive bends in response trajectories–that promise to reveal cognitive processes underlying sequential action learning.

Keywords: Sequence learning; serial reaction time task; sequential action; implicit learning; movement trajectory

Introduction

Sequence learning could be argued to be the most essential cognitive ability, for most of our daily behavior consists of sequential actions: from walking, cooking, and cleaning to spoken and written language. Consequently, sequence learning has been the topic of many studies, with widely varying approaches ranging from implicit sequence learning (Nissen & Bullemer, 1987; Cleeremans & McClelland, 1991; Stadler, 1992; Boyer, Destrebecqz, & Cleeremans, 2005) to language acquisition (Saffran, Newport, & Aslin, 1996) and typing (Fendrick, 1937; Gentner, LaRochelle, & Grudin, 1988). In the extensive area of implicit learning research, an important paradigm has been the serial reaction time (SRT) task, which asks subjects to press 4 buttons in a series, as indicated by corresponding lights in a sequence that repeats–unbeknownst to them–every 10 presses (Nissen & Bullemer, 1987). Subjects trained on this repeating sequence developed faster reaction times over the course of training, as compared to a control group responding to a random sequence of stimuli. The SRT paradigm has been cited as evidence for implicit learning, as subjects in the repeating condition, despite speeding up, are unaware of the sequence when later asked. However, performance does suffer somewhat when participants must simultaneously perform a second task (Nissen & Bullemer, 1987). The role of attention in the SRT task was further studied in Fu, Fu, and Dienes (2008), which demonstrated that reward motivation can improve the development of awareness of the sequence. Fu et al. (2008) reasoned that reward motivation

regulates the amount of attention paid towards the stimuli, which in turn facilitates sequence learning. The SRT task is neither wholly implicit or explicit: Willingham, Nissen, and Bullemer (1989) found that participants have varying degrees of declarative knowledge after a fixed training period in the SRT task, and that additional training resulted in greater explicit knowledge for many–but not all–subjects.

The dissociation of implicit and explicit processes facilitating sequence learning remains a topic of debate, yet learning remains robust under high degrees of noise and complex structure in the sequences. Cleeremans and McClelland (1991) used the SRT paradigm to train participants on a sequence with grammar-like structure and varying levels of noise. The presence of grammatical structure proved to facilitate sequence learning: participants responded faster on grammatical trials. Cleeremans and McClelland (1991) used a Simple Recurrent Network (SRN) to model these findings, which were confirmed by Boyer et al. (2005). The SRN demonstrates that associative mechanisms are able sensitive to statistical structure in sequences, suggesting that rule-like knowledge can be abstracted by networks that have been trained only on sequences of stimuli. Indeed, other modeling work shows that recurrent connectionist networks are able to learn hierarchically-structured, partially-ordered action sequences (e.g., coffee- and tea-making) merely from experiencing exemplars (Botvinick & Plaut, 2004; Kachergis, Wyatt, O'Reilly, Kleijn, & Hommel, submitted). Learning such sequences requires representing context, for as was noted by Lashley (1951), the majority of sequential actions cannot be explained by simply associating perceptual inputs with actions, since the prior action and the current percepts rarely sufficiently constrain action selection (for a review, see Rosenbaum, Cohen, Jax, Van Der Wel, & Weiss, 2007).

The SRT paradigm has proven to be a valuable tool for studying the basic mechanisms of sequence learning. However, to bridge the gap between button-pressing in the laboratory and real-world continuous actions such as language, we may need more nuanced measures to understand, for example, online error correction and dynamic effects of sequential context (Spivey & Dale, 2006). This paper introduces a mouse-tracking adaptation of the SRT task, in which participants react to stimulus changes by a corresponding move of the cursor. These movements are recorded, allowing us to analyze the trajectories for sequential context effects such as predictive motion (e.g., curving towards the location of the next stimulus). Such analyses promise to reveal more about

the underlying mechanisms of sequence learning than simply measuring response times, since we will be able to observe exactly when discrete action steps start to blend and merge into action phrases–much like neighboring phonemes in a word can be coarticulated. In this study, we show that our trajectory SRT task can replicate key findings from Nissen and Bullemer (1987), and offer additional measures of context effects.

Design

The purpose of the current study was to use the trajectory SRT paradigm to replicate earlier findings by Nissen and Bullemer (1987). This study used a recurring sequence of ten stimuli, horizontally displayed on a screen. Designating the stimulus positions from left to right as numbers, the original sequence read 4-2-3-1-3-2-4-3-2-1. To fit the trajectory paradigm the sequence was mapped to a square, left-to-right and top-to-bottom (i.e., 1=top left, 2=top right, etc.). Participants moved the mouse from one stimulus position to the next, corresponding to the sequence. We tested two groups of participants, one trained on the recurring sequence and the other trained on a random sequence. After ten blocks of training participants completed a generating task. This task consisted of the same basic test conditions, except participants were asked to predict the sequence instead of following it.

Nissen and Bullemer (1987) originally found participants showing improved performance within the first block of training. They continued to demonstrate performance suffered under dual-task conditions. The study also found that performance varied as a function of serial position in a pattern suggesting that learners were chunking the sequence into two pieces. In total, the study's results suggest that attention to the sequence is crucial for both implicit and explicit sequence learning, but that improved performance is not critically dependent on awareness of the sequence. For the purpose of this study only the initial experiment was replicated. We expect to replicate the basic improvement of performance, as well as the chunking pattern that was observed. Like Willingham et al. (1989), we included a final generation task, in which participants were asked to reproduce any action sequence they felt they had learned during training.

Methods

Participants Participants in this experiment were 23 Leiden University undergraduate students who participated in exchange for 3.5 euros or 1 course credit.

Procedure Participants were told to quickly and accurately move the mouse cursor to whichever square turned green. After arriving at the highlighted stimulus, another stimulus was highlighted after a 500 ms ISI. Participants completed 80 training trials, each of which contained a series of 10 locations. Participants were given a rest break every 20 training trials. Following the training phase, participants were asked to try to reproduce any sequence they had learned.

Participants were alternately assigned to one of the two

between-subjects conditions according to the order they signed up. In the NB87 sequence condition, participants were given a repeating sequence of 10 locations corresponding to the Nissen and Bullemer (1987) sequence (4-2-3-1-3-2-4-3-2-1). In the random condition, participants followed a random movement sequence in which no repetitions (i.e., staying at the same location) were allowed.

Each block contained 80 location stimuli (i.e., 10 repetitions of the NB87 sequence) which participants had to track with the cursor. The stimulus display consisted of four red squares (location 1 = upper left, 2 = upper right, 3 = lower left, 4 = lower right), displayed continuously. Monitors were 17", set to 1024 resolution, and each stimulus was 80 pixels on each side, separated by 440 pixels of white space. As a participant's cursor arrived at the green square, the square's color would change to red, like the other stimuli. The next stimulus in the sequence would change color after a 500 ms ISI. After training the participants performed a generating task which consisted of the same test conditions as training. They were asked to predict where they thought the stimulus would appear and move the mouse accordingly. A correct prediction would cause no color change while a mistake would cause the correct continuation of the sequence to show up green.

Results and Discussion

RT results Data were analyzed from the 22 participants (11 per condition) that completed the experiment. Median movement time to a target was 1,040 ms (sd: 1,776). Of 17,578 target arrival times, 84 were removed for being slower than 2,816 ms (median+sd). Each subjects' median RT for correct movements on each block was computed. Figure 2 shows the mean of median RTs by block for the two conditions. Participants in both conditions got faster over the course of the experiment, but participants in the NB87 sequence condition improved more than those in the random condition, replicating the Nissen and Bullemer (1987) speedup. There was a 25% reduction in reaction time over the course of training. These data were analyzed by a two-way analysis of variance, which indicated significant main effects of condition ($F(1,20)$ = 31.3, $p <$.001) and block ($F(7,168)$ = 6.3, $p <$.05), and a significant interaction effect ($F(7,210)$ = 14.7, $p <$.01) between the two.

The accuracy data is shown in Figure 3. Accuracy was high across training blocks although it dropped over time in the NB87 group, particularly after the first three blocks of training. A two-way analysis of variance confirmed a significant main effect of group ($F(1,20)$ = 36.7, $p <$.001) and a significant interaction effect ($F(9,210)$ = 14.1, $p <$.001).These results are evidence of sequence learning, replicating the Nissen and Bullemer (1987) keypress-based results. However, there was a speed-accuracy tradeoff in the NB87 condition: both accuracy and RT dropped over time. This was not present in the Nissen and Bullemer (1987) results, but can be explained through the difference in response execution. Key-presses are intermittent and can only be made in response to a stimulus (pre-stimulus responses were not recorded), while mouse

Figure 1: Characteristic movements in one trial from each condition.
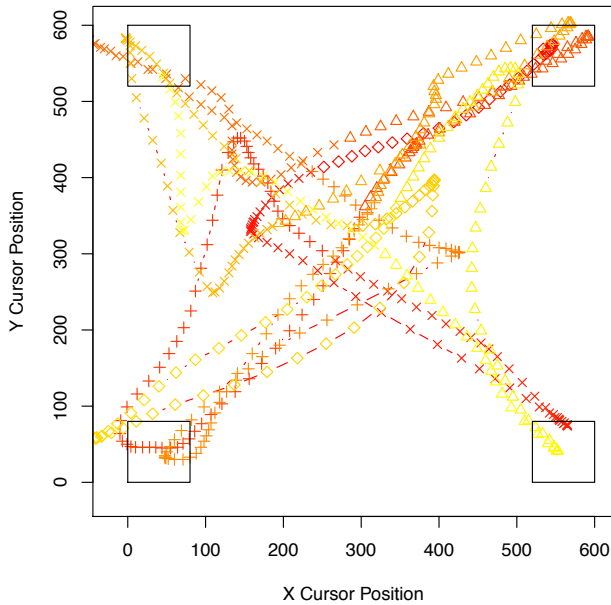
(a) A trial from the random condition, in which the next location was chosen at random, without repeats. All 11 random participants adopted a similar strategy of re-centering the cursor after each response. This is optimal in the sense that it was impossible to know which location will be highlighted next. (t_0 = red, t_{end} = yellow)
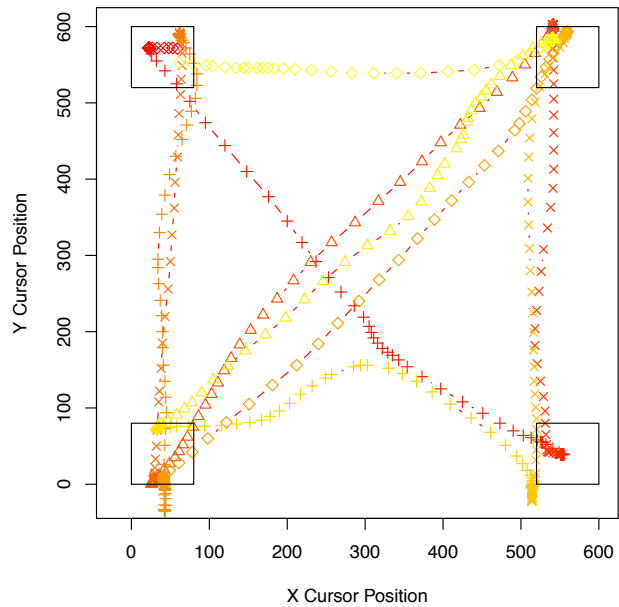
(b) A characteristic trial of a participant's movements during the NB87 sequence, beginning at location 4 (lower right) and ending at location 1 (upper left). These isomorphic trajectories can be compared for context effects. Only 4 NB87 participants showed centering movements in the last half of training.
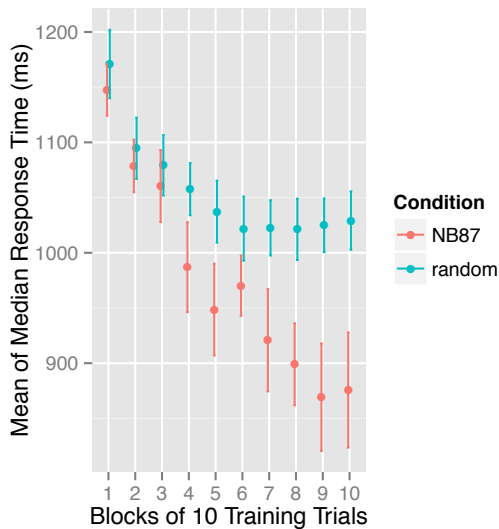

Figure 2: Mean of median RTs by block show that both conditions sped up over the course of Experiment 1, but that NB87 improved more. Error bars show +/-1SE.

movements are continuous and made constantly. Indeed, in the NB87 condition faster median hit RTs on a training block were significantly negatively correlated with the number of errors in that block (only for the 67 of 110 blocks containing errors; $r = -.56$, $t(65) = -5.48$, $p < .001$), evidenc-

ing a speed-accuracy tradeoff. The trajectory SRT paradigm encourages prediction, allowing participants to move freely while performing the experiment. Figure 4 shows that participants in the random condition move back to the center, while NB87 participants predictively move to another stimulus in the 500 ms before it is highlighted.

A two-way ANOVA with block as between- and serial position as within-subject factors, which showed significant main effects for block ($F(9,210) = 32.3$, $p < .001$ and serial position ($F(9,100) = 10.2$, $p < .01$). To determine whether participants became faster at the entire sequence or rather learned some chunks better than others, mean RT was plotted for each serial position, shown in Figure 5. Similar to the Nissen and Bullemer (1987) results, RTs on the second, fifth and eighth serial positions are slow, which may indicate that participants chunk the full sequence into two small, well-learned pieces.

Performance on the generating task was poor, as participants on average did not manage to reproduce the sequence without making many errors, as shown in Table 1. This indicates that, although training performance showed evidence of sequence learning, participants were not aware of the sequence. It is possible that participants would eventually be able to reproduce the sequence if training were extended, as in Willingham et al. (1989). Nissen and Bullemer (1987) originally found that participants were able to score around 80% correct on the generating task after two blocks of ten trials. Although the current study only required participants to
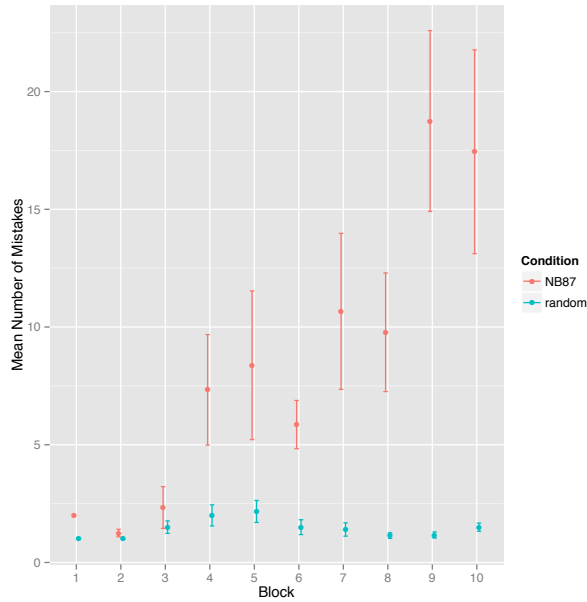
Figure 3: Mean number of mistakes by block show only the NB87 condition. Error bars show +/-1SE.
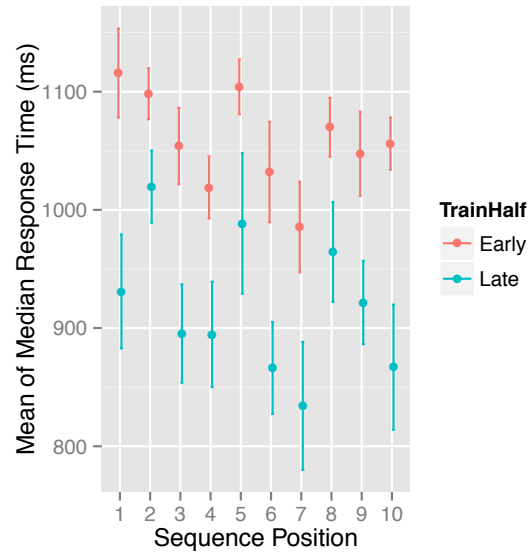

Figure 5: Mean of median RT by serial position during the early and late halves of training. Bars show +/-1SE.
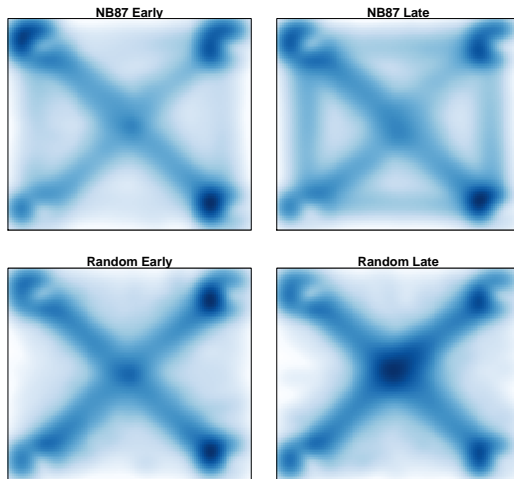

Figure 4: Smoothed scatterplots for each condition split by training half, showing only the movements made during the 500 ms just after arriving at the target before the next target is shown. Random condition participants re-center, whereas NB87 participants move towards other stimuli (vertical and horizontal lines, upper right).

Table 1: Generating task performance by NB87 participants.

Sequence Index	Hit RT	Average mistakes
1	1560.09	5.64
2	1576.82	5.70
3	1715.68	5.80
4	1541.14	6.70
5	1574.00	7.09
6	1448.09	5.70
7	1482.18	5.50
8	1427.00	5.10
9	1426.23	5.33
10	1334.05	5.18

complete one block of ten trials during the generating task, participants did not show any improvement during the task.

Trajectory results Figure 1 shows an example of mouse movements during a characteristic trial from each condition. Unexpectedly, some participants in the random condition (e.g., Figure 1 left) began re-centering the cursor after hitting a target, during the 500 ms ISI. This strategy is not unreasonable, as it minimizes the distance to potential targets, and the next target cannot be predicted in the random condition.

With learning, targets are predictable in the NB87 sequence condition, thus participants are expected to show faster reaction times (RTs) as training proceeds.

The NB87 sequence, 4-2-3-1-3-2-4-3-2-1, contains only one identical transition (3-2, a diagonal movement), although other movements are isomorphic (e.g., 4-2 and 3-1). We examined the development of sequential context effects–deflections in response trajectory caused by the prior or subsequent location–by plotting the average trajectories for the isomorphic movements: 4-2 vs 3-1. In the experiment, these movements are vertical, and we are interested in investigating the average deflections from the direct path from one stimulus center to another. We averaged position across subjects for these movements and plotted their deviation from the direct path (y-axis) over time (x-axis) in Figure 6, split by condition, and for each half of training. Early in training, some centering behavior is apparent in both conditions, most notably in the 4-

Figure 6: Averaged trajectories for vertical movements 4-2 and 3-1.
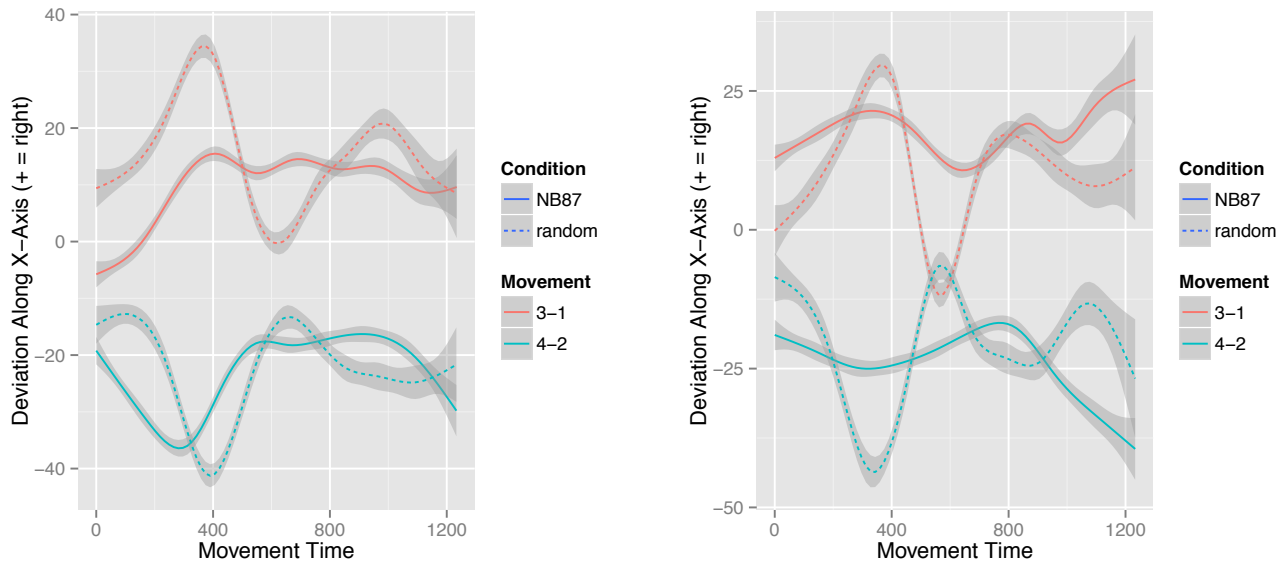
(a) Horizontal deviation during movement (i.e., over time) in early training. Trajectories in both conditions show some centering behavior, bending towards the middle (i.e., up for 3-1, down for 4-2). NB87 trajectories generally show less deviation.

(b) Horizontal deviation during movement in late training. The random condition shows more centering behavior, while the NB87 trajectories show little variation except at the end of the movements when they diverge, showing prediction of the subsequent stimulus.

2 movement. This movement also clearly shows the absence of centering behavior late in training for the NB87 condition. The 4-2 movement also shows participants tended to move towards the left after completing the movement. As the next target in the sequence is 3, which is situated to the bottom left of the current target, this indicates they were preparing to move towards the subsequent target.

General Discussion

This paper introduced the trajectory serial reaction time task and found that it replicates the results of Experiment 1 of Nissen and Bullemer (1987). Thus, while the trajectory SRT paradigm retains the essence of the original SRT, it also affords the opportunity to measure a variety of more detailed statistics about subjects' continuous motions. Response trajectories can reveal uncertainty, predictive movements, reversals in decision, and other phenomena that may reveal the dynamics of the learning mechanisms at work. The present study examined the average trajectories of two isomorphic vertical movements that appear in the NB87 sequence, as well as in the random condition. The two movements have different subsequent stimuli in the NB87 condition, and were thus expected to show a sequential context effect: as participants learn where the next stimulus will be, they may start to move towards this response even as they finish the previous movement–as a piano player may reach for the next key while the current one is being sustained (Soechting, Gordon, & Engel, 1996).

We found not only that the expected context effects had developed by late training, but also evidence of possibly strate-

gic adaptive behavior in the random condition. Many participants in the random condition developed a re-centering approach after each response, waiting for the next (unpredictable) stimulus to appear. In a way this behavior is optimal, since the center of the screen is as close as possible to all stimuli. Some participants in both conditions showed this behavior to a limited extent early in training, but those trained on the NB87 sequence lost this behavior over time as they learned to predict the location of the subsequent stimulus–hinted at by the decrease in reaction times in this condition, and confirmed by the deviation in average trajectory towards the subsequent stimuli.

Overall, the behavioral results show a striking similarity to the Nissen and Bullemer (1987) results. The pattern of reaction times over sequence position was almost entirely equivalent to the pattern observed in the original study, although the movement reaction times were higher throughout training and participants showed less overall improvement. This can be explained through the mechanics of the paradigm: mouse movements require more time to be executed than single keypresses, and require some fine motor control and error correction. The sensitivity of the mouse can be adjusted to achieve a balance between RT and error; we used a very low sensitivity to reduce overall noise. Participants in the NB87 sequence condition nonetheless showed increased number of mistakes during training, indicative of a speed-accuracy trade-off which was not present in the Nissen and Bullemer (1987) results. It is possible that extending the training would eventually lead to a reduction of mistakes, as participants would gradually become aware of the sequence.

The centering behavior we encountered was previously found by Duran and Dale (2009), who tracked arm movements in a sequence learning experiment. Employing a 12-position array of stimulus positions, they asked participants to respond to targets by selecting them with a Nintendo Wiimote. Similar to our results, participants tended to re-center after successfully completing an arm movement. In an effort to prevent participants from engaging in this behavior, Duran and Dale (2009) adapted the paradigm so that the array was visible during the entire experiment–as it was during our experiment. Although the centering strategy is somewhat intriguing in its own right–especially since it occurred rarely in the sequence condition, but subsequent research might benefit from discouraging it. Limiting the response-stimulus interval to somewhat less than the current 500 ms may limit the amount of time people can devote to re-centering.

Previous research has demonstrated that simple associative mechanisms are able to form these kinds of expectations when trained on a sequence (Cleeremans & McClelland, 1991; Stadler, 1992; Boyer et al., 2005). Although the SRN used by (Cleeremans & McClelland, 1991) is not suited for modeling the trajectory SRT paradigm, as the SRN expects discrete rather than continuous activation at the input level, other models such as the LEABRA neurocognitive architecture (O'Reilly, Hazy, & Herd, in press) might be used in subsequent research to model the paradigm. Training a cognitive model on continuous, structured sequences promises to reveal more about the underlying neural processes involved in online error-correction and associative learning.

Sequential action is complex, being hierarchically-structured and yet only partially-ordered, but it can be seen to be the foundation of nearly all human activity: from speaking and typing to cooking and sport. It is therefore important to study sequential action in detail, and we are confident that recording detailed action trajectories will grant additional insights into the dynamic, continuous mechanisms that allow us to learn sequential actions. In a second step towards this goal, we recently conducted a sister study extending the trajectory SRT paradigm to the design of a statistical word segmentation experiment (Saffran et al., 1996), using multiple recurring action subsequences instead of artificial words, but investigating the role of periodic rewards on learning multi-step actions (Kachergis, Berends, Kleijn, & Hommel, submitted). We hope that other researchers will join us in investigating action context effects in diverse sequential learning tasks.

Acknowledgments

References

Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: A recurrent connectionist approach to routine sequential action and its pathologies. *Psychological Review*, *111*, 395–429.

Boyer, M., Destrebecqz, A., & Cleeremans, A. (2005). Processing abstract sequence structure: Learning without knowing, or knowing without learning? *Psychological Research*, *69*, 383–398.

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

Duran, N. D., & Dale, R. (2009). Predictive arm placement in the statistical learning of position sequence. In *Proc. of the 31st annual conference of the cognitive science society.*

Fendrick, P. (1937). Hierarchical skills in typewriting. *Journal of Educational Psychology*, *28*(609–620).

Fu, Q., Fu, X., & Dienes, Z. (2008). Implicit sequence learning and conscious awareness. *Consciousness and Cognition*, *17*, 185–202.

Gentner, D. R., LaRochelle, S., & Grudin, J. (1988). Lexical, sublexical, and peripheral effects in skilled typewriting. *Cognitive Psychology*, *20*(4), 524–548.

Kachergis, G., Berends, F., Kleijn, R. de, & Hommel, B. (submitted). Reward effects on sequential action learning in a trajectory serial reaction time task. *ICDL-EpiRob*.

Kachergis, G., Wyatt, D., O'Reilly, R. C., Kleijn, R. de, & Hommel, B. (submitted). A continuous time neural model for sequential action.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–136). New York: Wiley.

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

O'Reilly, R. C., Hazy, T. E., & Herd, S. A. (in press). The leabra cognitive architecture: How to play 20 questions with nature and win! In S. Chipman (Ed.), *Oxford handbook of cognitive science.* Oxford: Oxford Uni. Press.

Rosenbaum, D. A., Cohen, R. G., Jax, S. A., Van Der Wel, R., & Weiss, D. J. (2007). The problem of serial order in behavior: Lashley's legacy. *Human Movement Science*, *26*, 525–554.

Saffran, J., Newport, E., & Aslin, R. (1996). Word Segmentation: The Role of Distributional Cues. *Journal of Memory and Language*.

Soechting, J. F., Gordon, A. M., & Engel, K. C. (1996). Sequential hand and finger movements: Typing and piano playing. In J. R. Bloedel, T. J. Ebner, & S. P. Wise (Eds.), *The acquisition of motor behavior in vertebrates.* MIT Press.

Spivey, M. J., & Dale, R. (2006). Continuous dynamics in real-time cognition. *Current Directions in Psychological Science*, *15*(5), 207–211.

Stadler, M. A. (1992). Statistical structure and implicit serial learning. *Journal of Experimental Psych.*, *18*, 318–327.

Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1047–1060.