ELSEVIER

Cognitive Brain Research 9 (2000) 239–248

COGNITIVE BRAIN RESEARCH

www.elsevier.com/locate/bres

Research report

The effects of sequence structure and reward schedule on serial reaction time learning in the monkey

Emmanuel Procyk [a,c,*], Peter Ford Dominey [b], Céline Amiez [c], Jean-Paul Joseph [c]

[a] *Section of Neurobiology, Yale University School of Medicine, New Haven, CT 06510, USA*
[b] *Institut des Sciences Cognitives, 67 blv. Pinel, 69500 Bron, France*
[c] *Inserm U94, 16 av. du doyen Lépine, 69500 Bron, France*

Accepted 14 December 1999

Abstract

This research tests the hypothesis that sequence learning performance in non-human primates will be modulated both by the structure of the sequences to be learned and by the schedule of reward applied during learning. Sequence learning in humans has been extensively explored with serial reaction time (SRT) protocols where learning is revealed by reduced reaction times for stimuli presented in repeating sequences vs. stimuli presented in random series. The SRT protocol has been used to demonstrate that different types of sequential structure may be learned under different awareness conditions. Here, we consider surface and abstract structure of sensorimotor sequences such that sequences ABCBAC and DEFEDF (where A to F correspond to spatial locations on a touch sensitive screen) have different serial order or surface structure, but share the same abstract structure 123213, and are thus considered isomorphic. In four experiments, we manipulated the type of sequential structure to be learned, and the schedule of reward in spatial sequence learning tasks. Both of the two monkeys tested demonstrated significant SRT learning for serial order or surface structure, while they failed to learn and transfer abstract structure. Their learning performance was also modulated by the schedule of reward. These results are in support of our hypothesis and are discussed in the context of existing models of sensorimotor sequence learning.

Keywords: Serial reaction time; Monkey; Sequence; Learning; Reward

1. Introduction

Sequence learning in humans has been extensively studied with serial reaction time (SRT) tasks [12] in which response times (RTs) are reduced for stimuli presented in a repetitive sequence vs. stimuli presented in a random series. This improvement in RTs provides a quantitative measure of sequence learning. Variations both in the attentional state of the subjects, and in the structure of information present in the sequences have been used to dissociate different forms of sequence learning [2,6,8]. Human subjects in implicit conditions can learn the surface structure of sequences in an SRT task, but they fail to learn the abstract structure [8]. Surface and abstract structure are defined such that sequences ABCBAC and DEFEDF have different surface structures, but the same abstract structure 123213. These two sequences are thus defined to be ''isomorphic''. In these sequences, the elements corresponding to 213 are predictable by the abstract structure. Thus, when exposed to new isomorphic sequences, the reaction times for these predictable elements should be reduced if the abstract structure is transferred. In explicit learning conditions, where subjects are aware that such an abstract structure could exist, these subjects learn the abstract structure, and can transfer this knowledge to new isomorphic sequences, with reduced RTs for elements that are predictable by the abstract structure [8]. In contrast, subjects in implicit conditions learn the surface structure of the target sequences, but fail to learn and transfer abstract structure to the new isomorphic sequences. This suggests that surface and abstract structure learning processes are neurophysiologically dissociated [8].

In non-human primates, the serial organization of behavior has been investigated through learning procedures called 'forward' procedures, and by testing capacities in transitive inferences [5]. Such studies have provided evi-

* Corresponding author. Section of Neurobiology, Yale University School of Medicine, 333 Cedar Str., SHM B413, New Haven, CT 06510, USA. E-mail: procyk@lyon151.inserm.fr/procyk@kafka.med.yale.edu

0926-6410/00/$ - see front matter
PII: S0926-6410(00)00002-1

dence for complex serial representations constructed by monkeys [3,4,14]. On the other hand, using trial and error tasks, some investigators have analyzed the ability of monkeys to solve sequential problems, and designed tasks that are adaptable to neurophysiological studies [9,13]. In such experiments, the reward is, for monkeys, the primary motivation to perform, and a central event around which behavior is organized. The functional significance of reward, or stimuli that predict reward, is reflected in the discharge of dopamine producing cells of the substantia nigra pars compacta [10]. We have previously suggested that this release of dopamine in the striatum could form the basis for modification of cortico-striatal synapses required for sensorimotor sequence learning [7].

The current research explores sequence learning in the monkey using SRT protocols, in order to test the hypothesis that monkeys will demonstrate SRT learning modulated by sequence structure and reward schedule. We first investigate learning as a function of sequence length for simple sequences. We then test the ability to learn and transfer abstract structure between isomorphic sequences. In both cases, we test different reward schedules and their effect on SRT performance. The primary aim of the study is to explore the effects of sequence structure and the schedule of reward on sequence learning in the monkey. It is in addition designed to evaluate the feasibility of an animal model for SRT learning, such that data could be useful for future neurophysiological or neuropharmacological studies.

2. SRT protocol

The sequence learning tasks that we employed involved touching a spatial target that appeared in a sequence of different locations on a touch-sensitive screen. Sequence learning was quantified as the difference in response times (RTs) for stimuli presented in random vs. fixed sequence series. Two rhesus monkeys (P and E) were subjects in these experiments. The animal was seated in a primate chair in front of a tangent touch-screen (MicroTouch System) coupled to a TV monitor (30×40 cm), in a sound-attenuating chamber. The screen was located at arm's reach. In the front panel of the chair, an arm-projection window (10×10 cm) was opened, and allowed the monkey to touch the screen with one hand. A PC 486 DX 33 computer controlled the presentation of visual stimuli on the monitor, which served as light-targets (targets: 2×2 cm white squares). It also recorded and evaluated (for reward purposes) the correctness of each touch. The animal worked with nine targets arranged as illustrated in Fig. 1A.

We trained the animal to point to targets presented in isolation (one at a time) on the screen. If the monkey touched the target during the illumination, the target was extinguished immediately after the touch. The RT (delay between onset of the target and the touch) was recorded. During the interval between the response and the subsequent stimulus (RSI), the animal was free to keep its hand in a ready position. The stimuli were organized within sequences of fixed number of successive targets. The sequences were of three types, random, fixed, or isomorphic. In fixed sequences, a pre-selected set of either three or four targets was presented in fixed order. In random sequences, these targets were presented in a randomly chosen order. In isomorphic sequences, the targets were different from one sequence to another, but all sequences shared a common abstract structure, and were thus isomorphic. A block consisted of a fixed number of sequences, which were presented consecutively. The number of sequences within a block was chosen so that the number of target-presses was approximately 45. Thus, when the sequences had three or four targets, there were respectively 15 and 11 sequences within a block. Execution of a block lasted about 70 s. Execution of a block was followed by a resting period of 45 s. If the percentage of no-responses was above 25%, then the corresponding block was discarded. Thus, in Figs. 1 and 2, successive blocks (Random or fixed in abscissa) were not necessarily adjacent in time, but were performed at increasing delays on the time scale. The basis of the statistical analysis was the RTs. The role of different parameters on the amount of learning was assessed by variance analysis (ANOVA). Effects with $P < 0.05$ were considered to be significant. We ran successively four experiments that crossed sequence type (Fixed vs. Isomorphic) with reward schedule (after each touch vs. at the end of sequence).

3. Experiment 1: learning surface structure with reward after each response

This experiment tests the hypothesis that non-human primates can learn surface structure, or simple serial order, in an SRT protocol.

3.1. Experiment 1. Methods

A new set of three or four targets was selected each day. All targets in a sequence were different. Targets were presented one at a time, so there was no choice involved in the response. To encourage movement anticipation, the target illumination (i.e., the period during which a response could be made) was short (600 ms). Each correct response (touch before or during the 600 ms) was rewarded by a squirt of apple juice (Fig. 1A). The interval between a response and onset of the next stimulus (RSI) was 1000 ms. Thus, all sequence elements (targets) were equivalent regarding time delays and reward delivery. A training session always started with the performance of two random blocks. In random blocks, the targets were randomly chosen among the three or four targets of the sequence to learn. Then, the sequence to learn (one per day) was presented in successive sequence blocks. The daily number

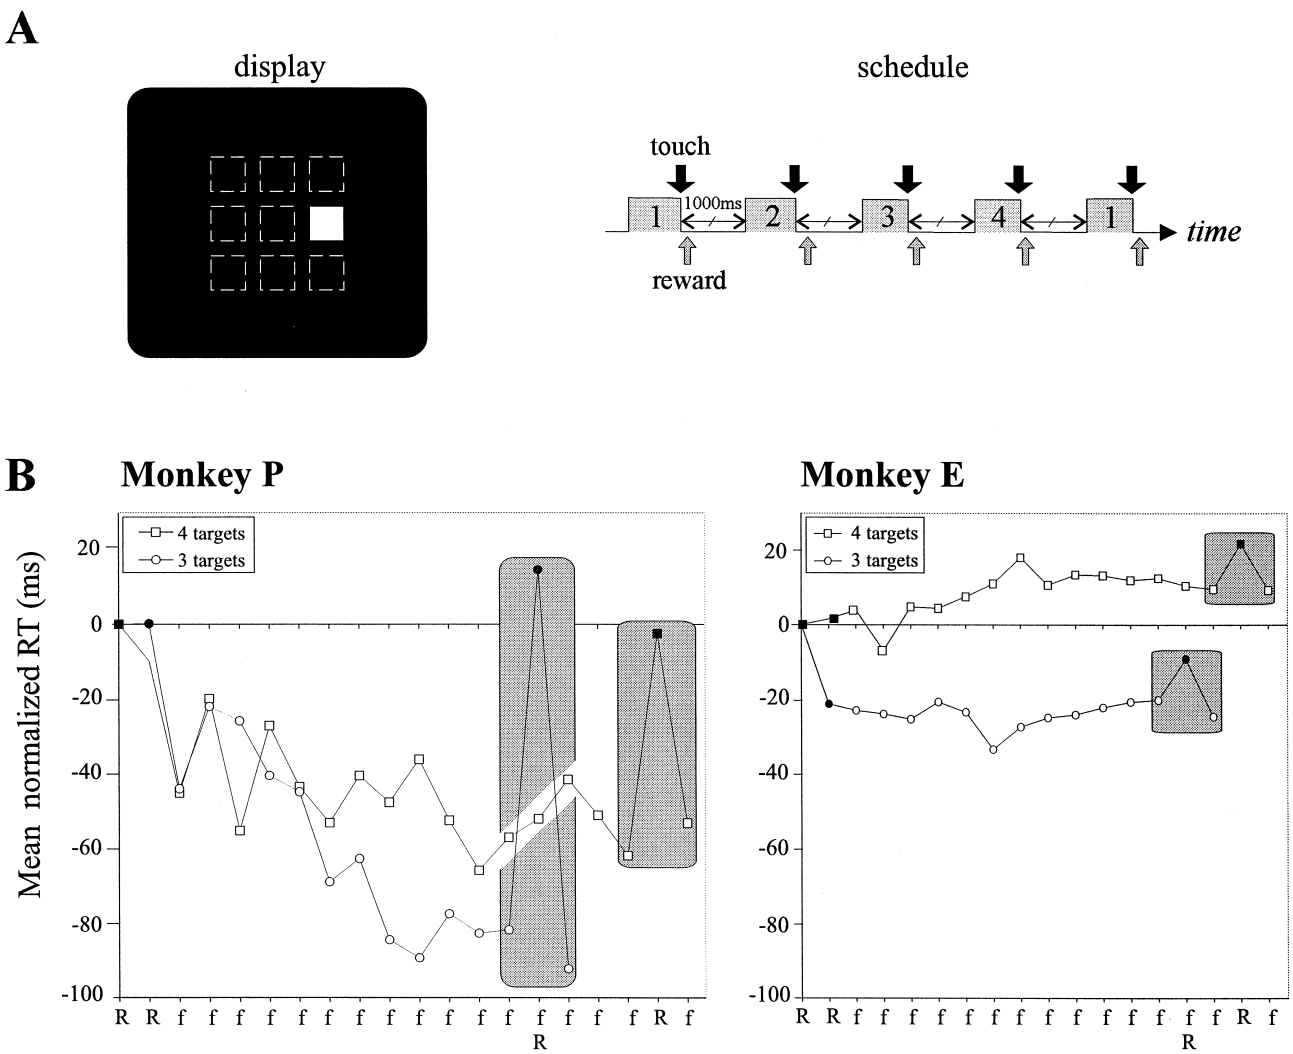

Fig. 1. Experiment 1: learning surface structure with reward after each response. (A) On the left, the target display is represented with one illuminated target. The location of the nine possible targets is indicated by empty squares. On the right, schematic representation of the task schedule for a four-target (1, 2, 3, and 4) sequence. Each gray square represents the period of corresponding target illumination. Black and gray arrows represent the occurrence of touches and rewards, respectively. The time delay between the touch of a target and the onset of the next is indicated. (B) The normalized average RTs plotted for each block of trials in fixed sequence learning for monkeys P and E. Each point corresponds to the mean over 4 days of the mean of RTs for one block of approximately 45 target presses. Means are normalized to the first random block (average RTs in the first random block: Monkey P, three targets: 488 ms, four targets: 472 ms; Monkey E, three targets: 442 ms, four targets: 362 ms). Blocks are labeled random (R) or fixed sequence (f) on the abscissa. Separate plots are shown for the learning of three- (circle) and four- (square) target sequences. Random blocks are illustrated by black symbols. The critical blocks for learning assessment are enclosed in the shaded areas. Both P and E display SRT learning (see text for details).

of successive sequence blocks was adjusted in such a way that the fixed sequence to learn was presented for a total of approximately 180 times. Then, the fixed sequence blocks were followed by a random test block, which was in turn followed by a fixed sequence block. This provides the critical comparison of RTs in sequential vs. random series for the quantification of sequence learning.

During 8 days of practice, four sequences of three targets, and then four sequences of four targets were successively learned. The mean RT within a block and the mean of means over the corresponding blocks in the 4 days were calculated. In random blocks, RTs corresponding to movements performed between two successive different targets, which exactly reproduce a succession of two movements in the sequence to learn were discarded from the statistics.

3.2. Experiment 1. Results

In Fig. 1B, mean normalized RTs are plotted for the successive blocks of sequence learning and as a function of the length of the sequence for the two monkeys. RTs are normalized to the first random block (the mean RT in the first random block is subtracted from the mean RT in each Block). For both monkeys in both conditions, we observe the signature of SRT learning in the last three blocks (shaded areas), as the RTs for targets presented in random

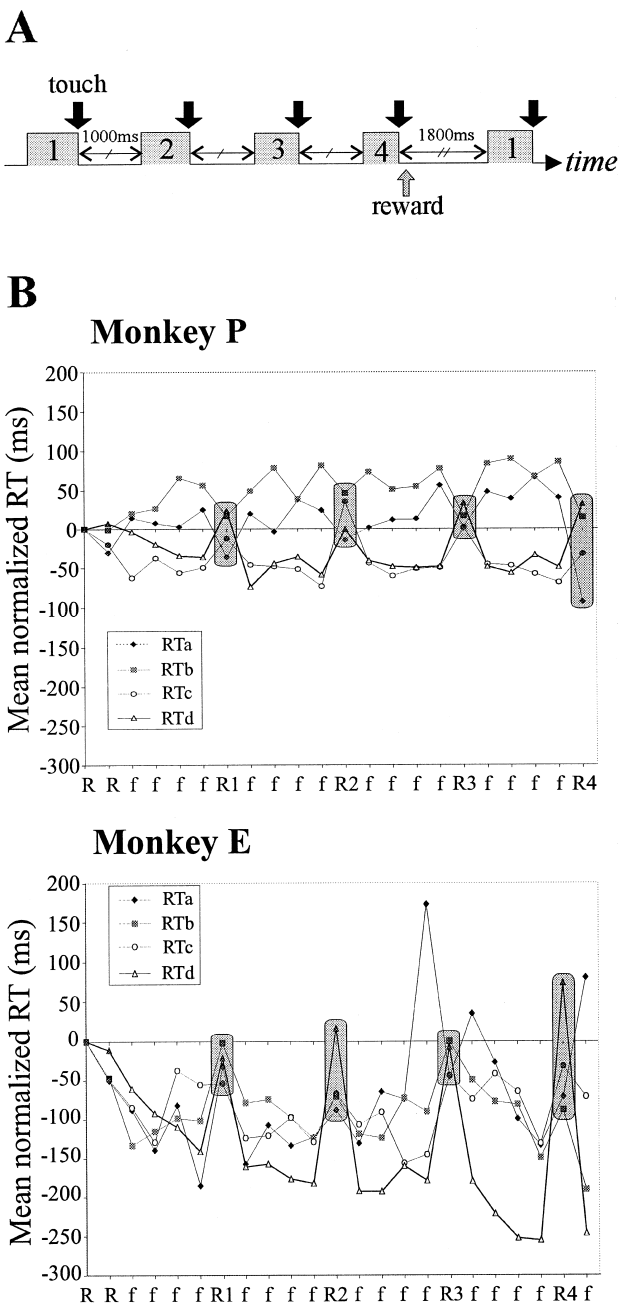

Fig. 2. Experiment 2: learning surface structure with reward at the end of the sequence (four targets). (A) Schematic representation of the task schedule. The time delay between the touch of a target and the onset of the next is indicated. Note that the delay between the fourth touch in a sequence and the onset of the first target in the next is longer. (B) The normalized average RTs plotted as a function of the target rank in the sequence (a–d) for each block of trials in four-target sequence learning. Means of 5 and 4 days for monkey P and E, respectively, normalized to the corresponding first random block (average RTs in the first random block: Monkey P: 594 ms, Monkey E: 656 ms). Blocks are labeled random (R) or fixed sequence (f) (abscissa). In every session, four random test blocks were performed (R1 to R4: data enclosed in the shaded areas). The analysis demonstrates that rank of the targets within a sequence induced significant learning variations. The decrease in RTs was present primarily for the last (rewarded) target (d) (see text for details).

order are augmented with respect to those presented in the fixed sequential order.

These observations in the four-target condition of reduced sequence vs. random RTs were confirmed by one-way ANOVAs, comparing the Random-test RTs to the mean RTs from the two surrounding Fixed blocks (one-way ANOVA, Block (Random vs. Fixed), monkey P: $F_{1,405} = 19.54$, $p < 0.01$; monkey E: $F_{1,428} = 7.91$, $p < 0.01$). In the three-target condition, a significant learning effect was found in monkey P only (one-way ANOVA, Block (Random vs. Fixed), monkey P: $F_{1,399} = 35.97$, $p < 0.01$; monkey E: $F_{1,527} = 2.99$, ns).

3.3. Experiment 1. Discussion

These results indicate that when each response is rewarded, non-human primates demonstrate SRT learning similar to that of humans. There are however, two points of discussion that must be addressed. First, in the four-target condition, monkey E displays RTs for sequence blocks that exceed those of the initial random series. However, the classic indicator of sequence learning in the form of elevated RTs in the random test block vs. the two surrounding sequence blocks was obtained. It is important to note that similar learning profiles have previously been observed in humans studies of SRT learning [2]. As the complexity of the sequence increases, the global RT reduction is diminished, but the random vs. sequence RT difference remains significant [2].

A related point is that this same monkey failed to demonstrate a statistically significant difference in sequence vs. random RTs for the three element sequence, despite the observation that they appear different in Fig. 1B. Modulation of responses within an individual in terms of attentional and motivational factors can contribute to this variability that can also be observed in human subjects (see Ref. [2]). In the four-target sequence however, this monkey displayed significant learning. In summary then, while subject to individual differences as seen in human subjects, these data permit us to conclude that non-human primates can demonstrate learning of serial order, or surface structure, in an SRT protocol in which each sequence element is rewarded.

4. Experiment 2: learning surface structure with reward at the end of the sequence

It is well known that reward and the expectation of reward during learning have significant neurophysiological impacts on the limbic motivation system, in part via the activity of dopamine producing neurons in the substantia nigra pars compacta of the basal ganglia [10]. Experiment

2, thus, tested the hypothesis that learning would be modified based on the reward schedule.

4.1. Experiment 2. Methods

The protocol was the same as in Experiment 1 with the following differences. In fixed and random blocks, the monkey was rewarded at the end of the sequence (i.e., after the response to the last element in the sequence) (Fig. 2A), and only if each of the targets of the sequence had been correctly pressed (i.e., within the 600-ms response period). A time delay (1.8 s) separated the end of a sequence and the beginning of the next. The succession of fixed and random blocks was also modified such that random test blocks, in which targets were randomly chosen among the nine possible targets, were presented at four separate points in each testing session (R1 to R4, see abscissa in Fig. 2B), rather than only at the end of the session as in Experiment 1.

4.2. Experiment 2. Results

Fig. 2B illustrates the evolution of RTs for fixed sequence and random elements as a function of their rank in the sequence (a–d), for the four-target sequence (normalized data: for each rank (a–d), the mean RT in the first random block is subtracted from the mean RTs in each Block). Learning is evaluated both in terms of Random vs. Fixed RTs, and in terms of the rank-dependent RT reduction, and can thus be observed in two different manners. First, comparing the fixed sequence vs. random blocks (f and R in Fig. 2B), we observe that the sequence block RTs are reduced with respect to those for the random blocks, in particular for the elements at the end of the sequence. Second, comparing rank (a–d) within the fixed sequence blocks, we observe that RTs are reduced for elements at the end of the sequence, particularly element d, with respect to those at the beginning of the sequence.

The observation of an SRT learning effect that depended on the rank of targets was confirmed by a two-way Rank (a,b,c,d) × Block (R,F) ANOVA. Monkey P: Rank × Block interaction: $F_{3,2320} = 26.72$, $p < 0.001$. Monkey E: Rank × Block interaction: $F_{3,1728} = 15.28$, $p < 0.001$), with the most significant learning for the target (d) closest to the reward (planned comparison $p < 0.001$ for both monkeys). The observation within the fixed sequence blocks that there was a significant effect for Rank was confirmed by a one-way Rank (a–d) ANOVA. Monkey P, Rank effect ($F_{3,3560} = 519$, $p < 0.0001$). Monkey E, Rank: $F_{3,2692} = 106$, $p < 0.0001$. Planned comparisons revealed that for both monkeys, RTs for the fourth target were most reduced ($p < 0.0001$). Likewise, for both monkeys ANOVAs performed for each target independently showed that only the fourth target (d) expressed a significant decrease of RTs across fixed blocks (Monkey P: $F_{19,890} = 3.36$, $p < 0.001$; Monkey E: $F_{16,673} = 9.18$, $p < 0.001$). These effects of learning over successive sequence blocks were not present in the RTs for random blocks, with no significant evolution across blocks found for any of the four targets in random blocks.

4.3. Experiment 2. Discussion

For both monkeys, sequence learning was revealed by two different measures. First, sequence learning was revealed by the decrease in RTs for sequence vs. random blocks, particularly for the last rewarded target. With respect to this random vs. fixed sequence comparison, one could argue that it is possible that part of the effect is due to the fact that all nine possible targets were used in the random blocks, vs. the three or four targets used in the fixed sequences. (Note that in Experiment 1, the random blocks used the same three or four targets as used in the fixed sequence blocks). The observation that the RTs in random block do not significantly evolve over time argues that they provide a stable reference against which we can measure sequence learning. However, the comparison of the nine-choice random condition with the four-choice sequence condition must be interpreted with caution.

More importantly, we note that sequence learning was also revealed by a second measure, completely independent of the sequence vs. random comparison. The RTs for elements in fixed sequence blocks decreased with learning differently as a function of their rank and proximity to the reward. Rewarded rank-d elements displayed a significant reduction in RT, with respect to the other elements, that increased in successive blocks. This indicates that over the successive blocks, the monkeys became increasingly proficient in predicting the rewarded element, which requires knowledge of its position in the sequence with respect to the other elements. These data argue that modification of reward schedule modulates learning, in this case by yielding a more significant learning effect (RT reduction) for elements that best predict the reward, with increase RTs for elements that least predict the reward. Indeed, elements early in the sequence that have a low reward prediction value actually display an increase in RT for monkey P. A similar modulation of RT as a function of temporal proximity to the reward has also been observed by Bowman et al. [1].

5. Experiment 3: learning abstract structure with reward after each response

The two preceding experiments demonstrate the learning of serial order or surface structure in non-human primates. In Experiment 3, we examine the capability for non-human primates to learn abstract structure in the same context. Recall that for a sequences like ABCBAC or DEFEDF, constructed from the abstract structure 123213, the elements in the second triple 213 are predictable based

on elements in the first triplet 123 that are themselves unpredictable. Thus, one method to assess the learning of abstract structure is to compare the RTs for elements that are predictable by the abstract structure vs. RTs for those that are not predictable. The following experiments focused on the generalization of sequential structure, and not on the generalization of spatial structure. Consider a sequence that went from upper left to diagonal down followed by diagonal up. If this sequence were learned and then moved down by one element throughout, one could test a form of spatial generalization. While interesting for future study, this type of generalization was not the focus of our current investigation.

5.1. Experiment 3. Methods

In Experiment 3, Monkey P was trained with isomorphic sequences such as GHG, CDC, EFE, etc., all of which share the abstract structure 121 or ABA (Fig. 3A). Based on the predictable nature of the abstract structure, RTs for movements from B to A (predictable from the abstract structure) should be reduced with respect to those for movements from A to B (unpredictable). Given the nine targets, there were 72 different isomorphic sequences based on this abstract structure. Within a given block, the time-delay between touch of a target and onset of the next target was fixed at a value between 500 to 1000 ms. The maximum target illumination was 800 ms. There were 15 sequences per block. The monkey was trained each day during 3 days. On day 1, 2, and 3, successively 540, 720, and 660 sequences were presented, hence a total of 1920 sequences ABA over 3 days.

The basis of statistical analysis was the difference in RTs for the second target (transition AB) and for the third target (transition BA). Both movements were of same amplitude but of opposite direction, and only transition BA was predictable from the abstract structure. In case of learning of the abstract structure ABA, positive values of the difference in RTs for the two transitions (AB–BA) were expected. This is because for a given isomorphic sequence, AB is unpredictable by the abstract structure ABA, while BA is predictable.

5.2. Experiment 3. Results

The plots of mean AB–BA values and a plot of mean RTs are illustrated in Fig. 3B as a function of day of practice. The magnitude of positive values of AB–BA indicates the degree to which the rule ABA has been applied to the sequences. We see in Fig. 3 (right) that during 3 days, the AB–BA difference remained negligible, though the overall reaction times reduced during the 3 days (Fig. 3, left). Indeed, AB–BA was actually negative in the two first sessions (AB vs. BA, *t*-test, $df = 1$, $p < 0.05$), and in the last session, AB and BA were statistically equal. The reduction in RTs for the two component responses AB and BA was confirmed (ANOVA, Main effect (Day), AB: $F_{2,1389} = 291.2$, $p < 0.0001$; BA: $F_{2,1389} = 287.3$, $p < 0.0001$). Thus, while there was an overall reduction in reaction times, it was not due to learning the ABA structure of the sequences.

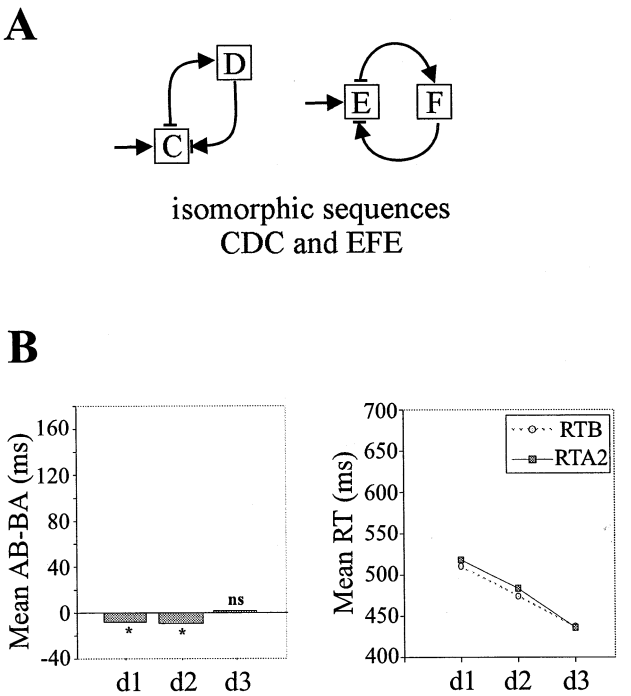

Fig. 3. Experiment 3: learning abstract structure with reward after each response. Monkey P. (A): schematic representation of two isomorphic sequences (CDC and EFE), which share the same abstract structure, i.e., 121. Letters C, D, E, and F symbolize four different spatial locations. Arrows indicate movement directions. The 72 isomorphic sequences were randomly presented in each session. (B) On the left: mean differences in RTs for B (movement AB) and for the second A (movement BA) expressed as (AB–BA) in isomorphic sequences of the type ABA. Values are plotted for successive training sessions. AB is unpredictable by the abstract structure, and BA is predictable, thus positive values of AB–BA reflect knowledge of the abstract structure. On the right: mean RTs are plotted for the successive sessions. Data show no learning of the abstract structure (AB vs. BA, one-way ANOVA, ns: not significant, $^{*}p < 0.05$).

5.3. Experiment 3. Discussion

This failure to extract the abstract structure under these conditions is not surprising. We have previously demonstrated [8] that if human subjects are not made explicitly aware of the possible existence of abstract structure in a sequence, they fail to express abstract structure learning, though their surface structure learning is intact. It was of interest, however, in this experiment to determine if this could change over the course of several days of exposure. The observation in this experiment of monkey P that previously learned surface structure but failed to learn abstract structure thus supports the hypothesis that surface and abstract structure are treated by dissociable neurophysiological processes [8].

6. Experiment 4: learning abstract structure with reward at the end of each sequence

In this experiment, we thus attempted to make the abstract structure more salient via two manipulations of the protocol, similar to those introduced in Experiment 2. First, we increased the temporal delay separating the end of one sequence from the beginning of another, in order to make the abstract structure more salient. Second, we placed the reward at the end of the sequence to provide a more salient indication of sequence boundaries. Concretely, it should be easier to extract the 121 regularity from ABA–CDC–EFE–GHG–IAI, than from ABACDCEFEGHGIAI.

In addition, we assessed abstract structure learning via a measure of the transfer of abstract structure knowledge to new isomorphic sequences. Recall that isomorphic sequences ABA and CDC have the same abstract structure. Learning of abstract structure with one of these sequences should allow transfer to the other. The measure of abstract structure learning will thus be to train subjects with one set of sequences (''training'' set) derived from a given abstract structure, and then to test the transfer of abstract structure knowledge to a ''transfer'' set of new isomorphic sequences constructed from the same abstract structure.

6.1. Experiment 4. Methods

The experiment, thus, tested the hypothesis that learning the abstract structure ABA would be more effective if each sequence was well isolated from the adjacent isomorphic sequences by the protocol changes described above. With respect to Experiment 3, two major differences were thus introduced. First, the reward was delivered at the end of the sequence and was followed by a time delay (1.8 s) before the beginning of the next sequence. Second, we explicitly examined the ability to transfer knowledge of the abstract structure to new isomorphic sequences. The 72 ABA sequences were thus divided into two sets (Training and Transfer) of, respectively, 16 and 56 sequences. The Training set was composed of all sequences ABA in which A was the upper-left or the middle target, and the Transfer set was composed of all other sequences. Thus, the repeated element of the sequences ABA (A) in the Training set was always different from the repeated elements in the Transfer set. Before the testing sessions began, the monkey had been familiarized with the task during a training session (170 trials) in which it worked with the Training set only. Then, each day, the experiment was proceeded in two phases. First, the monkey was re-exposed to the Training set (90 trials) to reinforce the learning of the abstract structure. Second, transfer was tested with sequences in the Transfer set (approximately 150 trials). Monkeys P and E were tested during three and four consecutive days, respectively.

6.2. Experiment 4. Results

Fig. 4 illustrates the monkeys' performance in the 16 sequences in the Training set and in the 56 sequences in the Transfer set. For both monkeys, at day 1 the difference in AB–BA for the Training set (labeled G1 in Fig. 4) is positive. For the Transfer test set (G2 in Fig. 4), however, the difference is much smaller for monkey P, and negative for monkey E. During the subsequent days, while the AB–BA difference for Transfer sequences becomes more positive, it never approaches the level observed for the Training sequences.

The observation of positive AB–BA difference in the Training set was confirmed for all sessions for both monkeys (one-way ANOVA, (AB vs. BA), $p < 0.01$ for each day). The same analysis for RTs in the Transfer set revealed, however, that the positive difference (AB–BA) became significant only in the second session (day) for

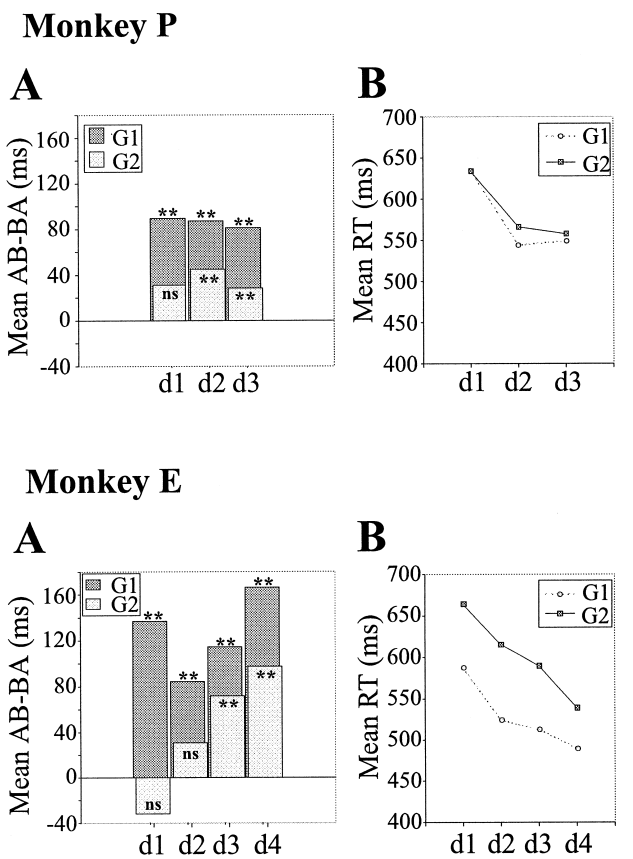

Fig. 4. Experiment 4: learning abstract structure with reward at the end of each correct sequence. (A) Mean differences in RTs for B and for the second A (AB–BA) in isomorphic sequences of the type 121, plotted for successive training sessions. (B) Mean RTs (average for all movements) are plotted for the successive sessions. Data for two separate sets of sequences, the 16 sequences in the Training set or group (G1), and the 56 sequences in the Transfer test set (G2) are illustrated separately. The figure shows sequence learning (AB vs. BA, one-way ANOVA, $^{*}p < 0.05$, $^{**}p < 0.01$), but a failure to transfer knowledge acquired in the Training set, to sequences in the Transfer set during the first day.

monkey P and in the third session for monkey E (at $p < 0.01$). Likewise, for both monkeys, in the Training and Transfer sequences, there was a significant reduction of AB and BA between the first and the last day (one-way ANOVA, $p < 0.05$). This is reflected in Fig. 4B as an overall decrease of mean RTs for training and transfer sequences. (ANOVA, day effect, $p < 0.001$).

6.3. Experiment 4. Discussion

The key observation in this experiment is that on day 1 of testing, both monkeys displayed a significant effect of learning (AB–BA) for the well-known sequences in the Training set, but they both failed to display a significant effect of transfer (by the same AB–BA measure) for the isomorphic sequences in the Transfer set. This difference only became significant in the subsequent days of testing. Note that monkey P appears to have some transfer in day 1, but recall that in Experiment 3, monkey P has already been exposed to all of the possible sequences (including the G2 sequences) during 3 days, and has thus likely acquired some knowledge of the G2 (transfer) sequences' surface structures. This is in stark contrast with the observation of transfer performance in human subjects, in which transfer is seen quite rapidly, even within the presentation of the first block of Ref. [8]. The failure to demonstrate transfer of abstract structure knowledge to the isomorphic sequences in the Transfer set on day 1 indicates a failure to learn the abstract structure.

Instead, the progressive improvement in the Transfer sequences over successive days indicates that the sequences were being learned verbatim, in terms of their surface structure. From this perspective, the fact that the Transfer set is larger (56 sequences) than the Training set (16 sequences) perhaps explains why the Transfer performance never approaches that seen for the Training sequences.

It is of interest to note that in this experiment, rewarding the movement BA will tend to make RTs for this movement diminish with respect to the AB RTs. This is based on the effects of reward on rank observed in Experiment 2 in which only the sequence-final elements were rewarded, and displayed the most significant RT reduction. From this perspective, the positive AB–BA values observed for the Training set would be due to sequence specific learning. This sequence specific effect does not transfer to the new isomorphic sequences in day 1, but only appears progressively, as a function of learning the individual sequences.

7. General discussion

The results of four SRT sequence learning experiments support our hypothesis that monkeys can learn in SRT tasks, and that the learning will be modulated both by the structure of the sequences to be learned and by the schedule of reward applied during learning.

7.1. Surface structure learning

The results of Experiment 1 demonstrate SRT learning in both monkeys, albeit with some inter-subject variability. For monkey P, the evolution of the RTs demonstrates the classical RT reduction with learning, and RT increase for random series, for both sequences. In monkey E, the signature of learning, in terms of this increase for random vs. sequential items is present in the four-element sequence, though the evolution of RTs does not demonstrate a progressive reduction. Experiment 2 demonstrates significant SRT learning in both monkeys. This learning in addition demonstrates sensitivity to the position of the reward, with the usual pattern of SRT learning observed primarily for the target most closely preceding the reward. This indicates that the animals have acquired knowledge of the sequential context since they know when and where the last target will be illuminated. The absence of this rank effect, or any learning, in random blocks indicates that we are not observing a non-specific counting effect, but rather a sequence-specific learning effect. One might argue that the rank effect is absent in random blocks because it is harder to ''count'' four random elements in a context of nine elements rather than four fixed. However, a pure counting strategy should function independently of the context. The fact that the rank effect appears only in the repeating sequences argues that sequence specific learning has occurred. These results from Experiments 1 and 2 are thus in agreement with the hypothesis that non-human primates will display SRT learning in which reward contingencies can produce rank-dependant learning effects.

In order to suggest a potential mechanism by which such rank selective RT reduction could occur, we can turn to our previous simulation studies of fronto-striatal function in sequence. We have developed a model of primate sensorimotor sequence learning in which learning takes place via the modification of cortico-striatal synapses [7] based on the reward-related release of dopamine in the striatum [10]. In the model, each sequence element is encoded in prefrontal cortex (PFC) neurons by a distributed pattern of activity that encodes the element itself and the previous context. The effect of reward is thus to link this cortical representation to the appropriate response-related neurons in the striatum by strengthening the appropriate cortico-striatal synapses. This leads to reduced RTs for predictable elements in SRT tasks [6]. If we reward only the sequence-final elements, the PFC will continue to encode the entire sequence, but RTs will be reduced only for the rewarded elements, as only these responses will benefit from increased synaptic efficacy due to learning.

7.2. Abstract structure learning

Given the successful demonstration of surface structure learning, we could then ask if monkeys can learn abstract structure under similar conditions. The choice of the abstract structure 121 in these experiments met two essential requirements. First, it qualifies as a legitimate abstract structure in that it can be used to generate an open class of isomorphic sequences (e.g., ABA CDC EFE, etc.). Second, it is short and thus should be relatively easy to learn. In Experiment 3, the monkey P was exposed at the outset to 72 isomorphic sequences of this form, and had no indication of the boundary between one sequence and the next. Neither the individual sequences nor their abstract structure were learned by monkey even after 3 days of training. This negative result is in fact in accordance with observations of failed abstract structure learning in human subjects who were unaware of the possible existence of an underlying abstract structure. Only subjects with such awareness displayed abstract structure learning [8]. An important difference is that while the human experiments took place in a single experimental session, the monkey was exposed to the training material for 3 days, during which we could have expected that knowledge of the abstract structure would be acquired and expressed. In Experiment 4, the reward was given at the end of the sequence, and an additional inter-sequence pause was introduced, both providing the animals with a clear sequence boundary that should facilitate extraction of the abstract structure. Additionally, initial training was provided with a Training set of 16 sequences, and then transfer was measured with the remaining Transfer set of 56 sequences. Given these changes, both animals learned the sequences in the training set, as revealed by reduced RTs for the BA component of the ABA sequences, but did they separately learn the 16 surface structures or the single abstract structure? We recall that the signature for abstract structure learning is transfer to the predictable elements of new isomorphic sequences in their first presentation [8]. The failure to transfer learning to the Transfer set of sequences in day 1 for both animals indicates then, that they learned the surface structure of the specific sequences, but not the abstract structure that was common to these sequences. Over the next days, performance improved for the sequences in the training set, but in a progressive fashion indicative of gradual surface structure learning, rather than the more discrete transfer of a well-learned abstract structure. We have additional evidence against abstract structure learning from Experiment 2. Here, a ''counting'' abstract structure of the form ''x x x x reward'', if learned, could have transferred to the random block. Again, the observation of no rank effect in the random block indicates that this generalization was not acquired.

These negative results should be interpreted with caution. One might argue that the repetition of A in ABA incurs some processing cost, as reflected, indeed by the negative learning effect for monkey P in Fig. 3, and for monkey E in day 1 of Fig. 4. Regarding the crucial measure of abstract structure learning as transfer to new isomorphic sequences, however, we note that for both monkeys in Experiment 4, despite the significant learning effect for the training set (G1), this is not seen in the testing set (G2). In other words, since any such cost has clearly been eliminated for G1, it should also be eliminated for G2. Thus, the poor performance for G2 indicates that G1 and G2 are treated as fundamentally different, with no recognition of their common abstract structure.

7.3. Methodological issues

There are important methodological differences between carrying out human vs. non-human primate behavioral experiments. In SRT experiments with humans, the experimenter can exploit verbal instructions to communicate to the subjects the objective to perform the task as quickly and accurately as possible. In working with non-human primates, the experimenter must communicate the objectives non-verbally. To induce the animal to respond as quickly as possible, duration of target illumination during which a response could be made was adjusted. It had to be small, but also it had to be large enough to allow the animal to press all targets in the prescribed time period. Likewise, the schedule of reward can be modified to modulate the behavioral salience of sequence elements near the reward.

From this perspective, we have observed that human subjects can learn abstract structure, but only when they have been made aware, by verbal instructions, of the possibility of an underlying abstract structure [8]. Such awareness could also potentially come from non-verbal means. Thus, before we can discount the possibility of abstract structure learning in non-human primates, we must consider alternate methods to instill awareness of the abstract structure, and alternate methods for expression of abstract structure learning.

In this context, it is of interest to note that a related experiment has recently been performed in 8-month-old human infants. The infants were exposed to 2 min of auditory sequences with abstract structure ABA. In an immediately following transfer test, their attentional orientation responses were tested for new sequences that either followed the ABA structure or a different structure ABB. Interestingly, the infants demonstrated a clear habituation response for the ABA sequences in the transfer test, and clear novelty responses for the ABB sequences, indicating their sensitivity to this abstract structure [11]. Marcus et al. [11] consider that this rule extraction capability may play an important role in extracting structural regularities from the speech signal during language acquisition. Again, before concluding that this behavior is absent in the non-human primate, we must recall that the attention orientation procedure is likely much more sensitive than the SRT

procedure we employed, and that future experiments should address this issue.

7.4. Conclusion

In conclusion, our results demonstrate that non-human primates are capable of learning sequential order at the level of surface structure as revealed by SRT protocols, and that this learning can be modified as a function of the reward schedule employed. This provides access to neurophysiological and pharmacological investigation of SRT learning in the non-human primate.

While we did not demonstrate a similar learning of abstract structure, it remains an open question as to whether this represents a true absence of capability in the non-human primates, or rather a problem of sensibility of our testing methods. Future research should further examine this question, as well as additional effects of manipulation of reward schedules, and the underlying electrophysiological correlates of the observed learning behavior.

Acknowledgements

We thank Patrick Monjaud and Marie-Line Loyalle for technical assistance, and Dr. William Clower for comments on a previous version of the manuscript. This work was supported by GIS Sciences de la Cognition, France.

References

[1] E.M. Bowman, T.G. Aigner, B.J. Richmond, Neural signals in the monkey ventral striatum related to motivation for juice and cocaine rewards, J. Neurophysiol. 75 (3) (1996) 1061–1073.

[2] T. Curran, S.W. Keele, Attentional and nonattentional forms of sequence learning, J. Exp. Psychol.: Learn. Mem. Cognit. 19 (1993) 189–202.

[3] M.R. D'Amato, M. Colombo, Representation of serial order in monkeys (*Cebus apella*), J. Exp. Psychol.: Anim. Behav. Process 14 (1988) 131–139.

[4] M.R. D'Amato, M. Colombo, Serial learning with wild card items by monkeys (*Cebus apella*): implications for knowledge of ordinal position, J. Comp. Psychol. 103 (1989) 252–261.

[5] C. De Lillo, The serial organisation of behavior by non-human primates; an evaluation of experimental paradigms, Behav. Brain Res. 81 (1996) 1–17.

[6] P.F. Dominey, Influences of Temporal Organization on Transfer in Sequence Learning: comments on Stadler (1995) and Curran and Keele (1993), J. Exp. Psychol.: Learn. Mem. Cognit. 24 (1) (1998) 234–248.

[7] P.F. Dominey, M.A. Arbib, J.P. Joseph, A model of cortico-striatal plasticity for learning oculomotor associations and sequences, J. Cognit. Neurosci. 7 (3) (1995) 311–336.

[8] P.F. Dominey, T. Lelekov, J. Ventre-Dominey, M. Jeannerod, Dissociable processes for learning the surface structure and abstract structure of sensorimotor sequences, J. Cognit. Neurosci. 10 (1998) 734–751.

[9] O. Hikosaka, M.K. Rand, S. Miyachi, K. Miyashita, Learning of sequential movements in the monkey: process of learning and retention, J. Neurophysiol. 74 (1995) 1652–1661.

[10] T. Ljungberg, P. Apicella, W. Schultz, Responses of monkey dopamine neurons during learning of behavioral reactions, J. Neurophysiol. 67 (1) (1992) 145–163.

[11] G.F. Marcus, S. Vijayan, S. Bandi Rao, P.M. Vishton, Rule learning by seven-month-old infants, Science 283 (5398) (1999) 77–80.

[12] M.J. Nissen, P. Bullemer, Attentional requirements of learning. Evidence from performance measures, Cognit. Psychol. 19 (1987) 1–32.

[13] E. Procyk, J.P. Joseph, Problem solving and logical reasoning in the macaque monkey, Behav. Brain Res. 82 (1996) 67–78.

[14] K.B. Swartz, S. Chen, H.S. Terrace, Serial learning by rhesus monkeys: I. Acquisition and retention of multiple four-item lists, J. Exp. Psychol.: Anim. Behav. Process 17 (1991) 396–410.