

Research report

A shared system for learning serial and temporal structure of sensori-motor sequences? Evidence from simulation and human experiments

Peter Ford Dominey *

Institut des Sciences Cognitives, UPR 9075, CNRS, 69008 Lyon, France

Accepted 22 July 1997

Abstract

This research investigates the influences of temporal structure on the representation of serial order. Experiments are performed in a neural network model of sequence learning and in human subjects. In the sequence learning model, a recurrent network of leaky integrator neurons encodes a succession of internal states that become associated, by reinforcement learning, with the correct sequential responses. First, the model is shown to learn a simple temporal discrimination task. The model is then exposed to two novel serial reaction time (SRT) experiments. In the standard SRT task (M.J. Nissen, P. Bullemer, Attentional requirements of learning: evidence from performance measures, *Cogn. Psychol.* 19 (1987) 1–32 [16]), reaction times for stimuli presented in a repeating sequence are reduced with respect to those for random stimuli, providing a measure of sequence learning. The novelty of the current experiments is that imbedded in the serial order of the sequences, there is a *temporal structure* of delays. The model is sensitive to both the serial structure and the temporal structure of the sequences. This observation is then confirmed in human subjects. These results demonstrate how a novel recurrent architecture encodes the interaction of temporal and serial structure and provide insight into related aspects of human sensori-motor sequence learning. © 1998 Elsevier Science B.V.

Keywords: Temporal sequence learning; Neural network; Fronto-striatal system; Recurrent network; Serial reaction time

1. Introduction

Virtually every aspect of cognitive function is imbedded in a sequential context. We see this in examples such as the perception and production of language and music, in game playing, problem solving and sensori-motor control. In each of these domains, the ability to correctly manipulate the order of sequential events is of primary importance. While order is a defining aspect of sequence, there is another dimension that is perhaps of equal importance, the temporal dimension, or organization *in time* of sequentially ordered events. What would music, speech or motor control be without a temporal structure imposed on the serial order of events?

Recent work in experimental psychology and related simulation studies have demonstrated that, indeed, the

disruption of the temporal organization of a learned sequence can have rather severe consequences on performance [5,11,21,22]. These studies are based on the serial reaction time (SRT) protocol of Nissen and Bullemer [16] in which sequence learning is demonstrated by reduced reaction times for stimuli that are presented in a repeating sequence vs. a random series. As part of an effort to investigate the effect of temporal structure in sequence learning, Stadler [22] tested subjects in an SRT task in several different conditions. One of the conditions closely corresponds to the standard SRT task. In this (no-pause) condition, the interval between each response and the successive stimulus (response-stimulus interval, RSI) was always 400 ms. In a second (pause) condition, the RSI was 2000 ms (instead of 400 ms) for approximately half of the stimuli, determined in a random fashion. In this pause condition, subjects were significantly impaired in sequence learning in comparison to subjects in the no-pause condition. Stadler interpreted this result as evidence that changes in the temporal organization of a sequence are in fact changes to the sequence itself that contribute to impair-

* Corresponding author. Vision et Motricité, INSERM Unité 94, 16 Avenue Doyen Lépine, Case 13, 69676 Bron, France. Fax: +33-472-91-34-01; E-mail: dominey@lyon151.inserm.fr

ments in learning. In this same line of thought, Stadler [21] also demonstrated that sequence learning was facilitated by the insertion of an RSI delay at the beginning of each repetition of a given sequence, and impaired if an RSI was introduced at a random location in each repetition.

We thus define *temporal structure* as the organization of RSI delays in an SRT task and observe, based on these results, that learning performance can be modified as a function of both the serial structure and the temporal structure of the sequence. An obvious question that arises is whether these two dimensions are treated by a single or separate mechanism. The working hypothesis of this research is that a single learning system should be capable of representing both serial and temporal structure of sequences. This hypothesis is tested by the examination of the interactions between representations of serial and temporal structures in a novel recurrent network. Recurrent networks [10,13,18–20] are inherently well suited for the treatment of sequence learning. In these systems, information about previous events, that is necessary to predict subsequent events, is maintained as an internal state via the recurrent connections. We recently described a model of sensori-motor sequence learning based on the primate frontostriatal system, in which the prefrontal cortex (PFC) is modeled as a recurrent network that encodes a sequence of internal states, and the striatum (caudate nucleus – CD) is modeled as an associative memory structure that binds internal states to their corresponding motor outputs [6–9]. While this model falls into the general category of recurrent networks, it has certain novel features that will shortly become apparent, that make it particularly well suited to address the hypothesis posed above concerning the co-representation of serial and temporal structures.

The following experiments test the prediction that the structure of temporal delays between sequence elements *and* the elements' serial order are encoded in an internal state that can subsequently drive the system's behavior. Expt. 1 demonstrates how the implicit representation of temporal delays can be exploited in learning a simple temporal discrimination task. Expt. 2 tests how changes in the RSI structure will influence the internal representation of the sequence and thus modify the behavior in terms of serial reaction times. Expt. 3 extends the results of Expt. 2 to temporal structure learning, and also provides specific predictions that are then tested in human subjects in Expt. 4.

The results of these experiments are discussed in the context of the proposed hypothesis, and the architecture of the model is compared with related recurrent architectures for sequence learning.

2. Model description

The model architecture that will be used in the rest of this paper, presented in Fig. 1, relies on a recurrent network to represent sequential state, and an associative mem-

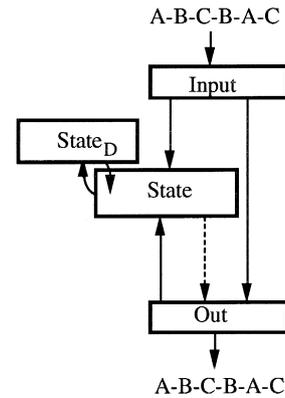


Fig. 1. Sensori-motor sequence learning model architecture. Each of the structures are 5×5 arrays of leaky integrator neurons. Sequences of spatial target locations are presented in the input array, Input. Responses are generated in the output layer, Out. Out units are influenced by Input, and also by modifiable, non-topographic projections from State. State encodes sequence state as a function of visual input from Input, response copy from Out, and self-input from State_D in order to generate a time varying sequence of internal states for each pre-response period in sequence reproduction and SRT tasks. These states become associated with Out activity for the correct response by a reinforcement learning mechanism. Note that due to the recurrent connections via State_D, if delays are introduced between input events, State will be modified, and thus time is coded in the same fashion as other external events.

ory that binds states to the appropriate responses. The general principle is that the neuron-like elements in the recurrent network generate a succession of distinct patterns of activation during the course of a sequence. Each of these patterns becomes associated, through learning, with each of the corresponding responses in the succession of sequence trails, yielding error-free sequence reproduction and/or reduced RTs as a function of the task. Recurrent sequence learning networks have previously been developed [1,10,13,18–20], and their sequence learning capabilities have been well documented [4]. Likewise, other effective sequence processing architectures have also been demonstrated (see [14] for a review). The recurrent network architecture described here is similar to previous recurrent models with three important differences. First, there is no learning in the recurrent connections (i.e., those that project from State_D to State), only between the State units and the Output units. Second, learning is based on a simple reinforcement mechanism rather than back-propagation of error, or related error-gradient calculation methods. Third, in the temporal domain, the computing elements are leaky integrators, and simulation time steps are not tightly coupled to input, output and learning processing. Indeed, the experimenter's capability to specify the time delays between external events is an integral part of this model originally developed to simulate primate behavioral electrophysiology experiments [3]. Following the model's production of a given response, the response is evaluated and the reinforcement learning rule (described below) is applied. The network can then run for an arbitrary number of time steps (i.e., the experimenter-specified

response-stimulus interval) before the next input is presented, making it ideal for studying the effects of RSI manipulation [5].

While the details of the model can be found in [5,6,8] we here summarize the most important interactions. Each of the layers displayed in Fig. 1 is modeled as a 5×5 array of leaky integrator neurons whose response latency is a function of their input intensity and their time constant. In the subsequent presentation, sequences will be referred to by letter strings, e.g. ABCBDC. Each of these letters corresponds to a single (x, y) location on the 5×5 input array. Thus, the 5×5 array is capable of encoding up to 25 distinct input symbols. Likewise, responses are generated in the 5×5 array, Output, with the same letter-location correspondence. The unit of time in the simulations is referred to as a simulation time step or STS and corresponds to a single update cycle of the simulation. In the equations that follow, the function $f(x)$ generates a non-linear (sigmoidal) output of x .

2.1. Recurrent state representation

The internal state of the system, i.e., the history of the input, output and previous state events that have so far occurred, is encoded in the recurrent network, State. The evolution of activity patterns in State over the course of a given sequence provides an ordered set of indices that become associated, via learning, with the ordered events in the sequence. For a given sequence, prior to each response, there will be a distinct pattern of activity in State. This pattern will become linked, by learning, to the response for that element, thus yielding error-free sequence reproduction after training.

Eq. (1.1) and Eq. (1.2) describe how State is influenced by external inputs from Input, responses from Out, and recurrent inputs from State_D. In Eq. (1.1) the leaky integrator, $s_i(t)$, corresponding to the membrane potential or internal activation of State, is described. In Eq. (1.2) the output activity level of State is generated as a sigmoid function, $f()$, of $s_i(t)$. The term t is the time, Δt is the simulation time step, τ is the leaky integrator time constant. As τ increases with respect to Δt , the charge and discharge times for the leaky integrator increase. In the simulations, Δt is 5 ms. For the 4 equations, the time constants are 10 ms, except for Eq. (2.1) which has 5 time constants that are 100, 600, 1100, 1600 and 2100 ms.

$$s_i(t + \Delta t) = \left(1 - \frac{\Delta t}{\tau}\right) s_i(t) + \frac{\Delta t}{\tau} \times \left(\sum_{j=1}^n w_{ij}^{IS} \text{Input}_j(t) + \sum_{j=1}^n w_{ij}^{SS} \text{StateD}_j(t) + \sum_{j=1}^n w_{ij}^{OS} \text{Out}_j(t) \right) \quad (1.1)$$

$$\text{State}(t) = f(s(t)) \quad (1.2)$$

The connections w^{IS} , w^{SS} and w^{OS} define the projections from units in Input, State_D, and Out to State. These connections are one-to-all, are mixed excitatory and inhibitory, and do not change with learning. This mix of excitatory and inhibitory connections ensures that the State network does not become saturated by excitatory inputs and also provides a source of diversity in coding the conjunctions and dysjunctions of input, output and previous state information. The n in the summation terms is 25, corresponding to the linearized size of the 5×5 layers.

Recurrent input to State originates from the layer State_D and Out. State_D (Eq. (2.1) and Eq. (2.2)) receives input from State, and its 25 leaky integrator neurons have a distribution of time constants from 20 to 420 simulation time steps (100–2100 ms), while State units have time constants of 2 simulation time steps (10 ms). This distribution of time constants in State_D yields a range of temporal sensitivity similar to that provided by using a distribution of temporal delays [14].

This coding of the internal state of sequence execution in the State recurrent network is a central point of the model. The fact that recurrent connections allow events that have occurred several elements in the past to influence the current context provides a powerful tool for resolving ambiguities in learning complex sequences [6]. This coding of sequence context in recurrent networks has been demonstrated in a number of related models [1,10,13,18–20].

$$sd_i(t + \Delta t) = \left(1 - \frac{\Delta t}{\tau}\right) sd_i(t) + \frac{\Delta t}{\tau} (\text{State}_i(t)) \quad (2.1)$$

$$\text{State}_D = f(sd(t)) \quad (2.2)$$

2.2. Associative memory

During learning, for each correct response generated in Out, the pattern of activity in State at the time of the response becomes linked, via reinforcement learning in a simple associative memory, to the responding element in Out. The required associative memory is implemented in a set of modifiable connections (w^{SO}) between State and Out. Eq. (3) describes how connections between State and Out are modified during learning. The threshold function in $f()$ is sigmoidal. When the response in any unit of Out reaches a certain level, the simulator considers that a response has been generated, and the response is evaluated. Thus, Eq. (3) is executed only once for each response. When a response is evaluated, the connections between units encoding the current state in State, and units encoding the current response in Out are strengthened as a function of their rate of activation and learning rate R . R is positive for correct responses and negative for incorrect responses, which are evaluated purely in terms of whether or not the right element in Out is activated, independent of its timing. Weights are normalized to preserve the total synaptic output weight of each State unit.

$$w_{ij}^{SO}(t + 1) = w_{ij}^{SO}(t) + R * \text{State}_i * \text{Out}_j \quad (3)$$

The network output is thus directly influenced by the Input, and also by State, via learning in the w^{SO} synapses, as described in Eq. (4.1) and Eq. (4.2). Learning thus yields reduced reaction times by selectively augmenting state-driven activity in Out.

$$o_i(t + \Delta t) = \left(1 - \frac{\Delta t}{\tau}\right) o_i(t) + \frac{\Delta t}{\tau} \left(\text{Input}_i(t) + \sum_{j=1}^n w_{ij}^{SO} \text{State}_j(t) \right) \quad (4.1)$$

$$\text{Out} = f(o(t)) \quad (4.2)$$

3. Experiments

The following experiments have been designed to test the hypothesis that the model described above can learn both the order (serial structure) and the temporal structure of input sequences. In this context, *serial structure* refers to the order of events, such as A–B–C–A–D–C, and temporal structure refers to the organization of RSI delays that separate the ordered events in time. Before directly testing this hypothesis we first provide a simple test of the model's sensitivity to temporal structure.

3.1. Experiment 1 (simulation): explicit encoding of time in a recurrent network

If neural activity in the recurrent network State is reliably sensitive to the passage of time as suggested above, then the model should be capable of discriminating arbitrary time intervals. That is, State's dynamics should provide a "clock" function. In order to test this prediction, a simple protocol was used in which a single stimulus (F) was presented, and then after different delays, a choice had to be made between three elements (A, B, C) presented together as the "go signal". The correct choice was determined as a function of the delay between offset of the stimulus F, and onset of the choice go signal. A 0 delay corresponded to choice A, a 10 ms delay (20 simulation time steps) corresponded to choice B, and a 20 ms delay (40 simulation time steps) corresponded to choice C. The model learned this task after less than 500 trials.

Fig. 2 displays the time-dependent changes in activity for the 25 State units over the time course of the three choice tasks, with the three different conditions (delays) superimposed on the same figure. At each of three choice selections, the state vectors are sufficiently distinct to be discriminated by the associative memory, in order to link to each state its appropriate response.

The recurrent network translates the passage of time into a trajectory of spatial patterns of activation in the State layer which can then be associated with the desired output. This result provides a simple verification that the

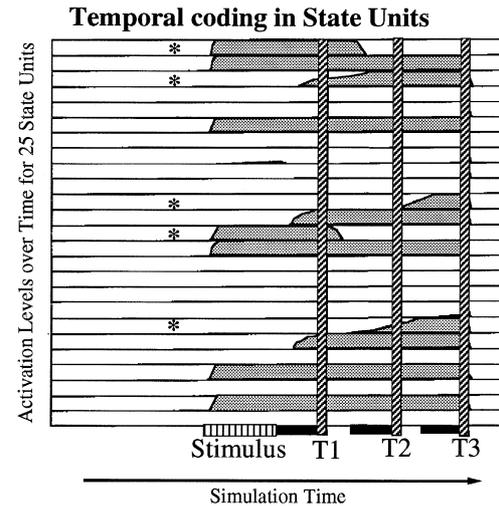


Fig. 2. Temporal coding in State units during a temporal discrimination task. Modification of State activity is seen to explicitly encode the passage of time after the Stimulus. Each row represents the activity of one of the 25 State units over time. Simulation time runs from 0 to 0.6 time units (120 time steps). After Stimulus presentation (indicated by vertically hatched bar on the time scale), the "go signal" (simultaneous onset of three targets A, B and C indicated by a dark bar on the time scale) can occur either with a delay of either 0 (T1), 0.10 (T2) or 0.20 (T3), corresponding to 0, 20 or 40 time steps. After the onset of the go signal, the State activity biases the response via the learned State-Out connections, and the correct choice is made. Three separate trials with the three delays are here superimposed. Vertical bars indicate the State activation at the time that each of the three choices are made (i.e., three separate trials overlaid on the same trace) after the model is trained to error-free performance. State units marked with "*" differentiate between two or three of the time-dependent conditions.

modifications in the internal state induced by different temporal delays are sufficiently distinct so as to be discriminated and associated with the desired output by the associative memory.

3.2. Experiment 2 (simulation): influence of temporal structure on sequence learning

Expt. 1 demonstrated the model's sensitivity to simple temporal structure. In order to study the interaction between serial and temporal structure, an SRT task is used that allows the comparison between RTs for sequential vs. random series (the standard measure of learning), using the same serial structure (i.e. A–B–C–A–D–C), but different temporal structures. Three temporal structures, T1, T2 and T3, were studied, each consisting of a repeating series of 20 RSI values of either 0.1 or 0.5 simulation time units (20 or 100 time steps, respectively). T1 = 0.1, 0.5, (0.1, 0.5)*, T2 = 0.1, 0.5, 0.5, (0.1, 0.5)*, T3 = 0.1, 0.5, 0.5, (0.1, 0.1, 0.5, 0.5)*, where "*" indicates repetition. By training the model with one temporal structure and then testing with another (same or different), we can assess the model's sensitivity to temporal structure.

A set of 10 instances of the model were generated by using different seed values for the random number genera-

tor to initialize the weights in w^{IS} , w^{SS} , w^{OS} and w^{SO} . The 10 instances were replicated to yield three equivalent groups of 10, corresponding to the three temporal structures. Each group was trained on 6 blocks of 80 trials using the repeating 6-trial sequence A–B–C–A–D–C using one of the three temporal structures. The models were then tested with one block of 80 trials using the trained sequence (SEQ) and another using a random (RAND) series in three different conditions. The first condition (Same) used the same temporal structure as that used for training, and the remaining two conditions (Different) used the two temporal structures different from that used for training. In the SRT test, there is only one possible response at each go signal, thus in each of these three conditions the subjects' performance was evaluated in terms of the RTs for the random and sequence series, where reduction in RT for sequence vs. random blocks is evidence that the sequence has in fact been learned.

3.2.1. Performance effects

An example of the model's performance is displayed in Fig. 3. The RTs for responses in sequence blocks 1–6 reduce progressively and then increase when a random series of targets is introduced in block 7. This SEQ (block 6) vs. RAND (block 7) difference is significant (1-factor ANOVA, $F_{1,162} = 96.29$, $P < 0.00001$), indicating significant learning of the serial structure of the sequence.

The performance effects of the temporal structure used for training and testing were quantified by a 2-way ANOVA in which the two factors were the temporal structures used respectively in Training (T1, T2, T3) and Testing (Same or Different from that used in training). The dependent variable was the learning as measured by the difference in reaction times for sequence and random blocks (see Fig. 3). The significant main effect for Training ($F_{2,84} = 19.61$,

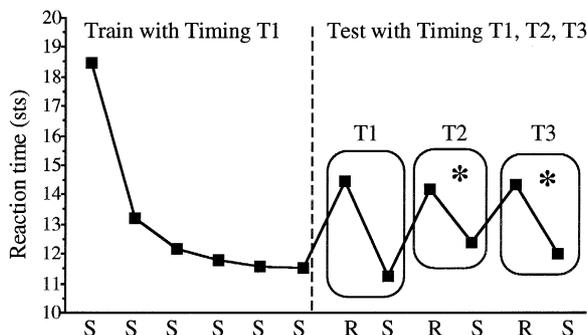


Fig. 3. Simulation performance in an SRT task. Training with Temporal Structure T1: 6 successive blocks of 80 trials using the sequence A–B–C–A–D–C, and temporal structure T1. Test with Timing T1, T2, T3. Random (R) and Sequence (S) blocks using three temporal structures T1, T2, T3. RTs are expressed in STS units as means and S.D.s for the 80 trials in each block for the 10 models trained on T1, then tested on T1, T2 and T3. Learning is measured as the difference in RAND-SEQ RTs. Learning in conditions T2 and T3 is present, but significantly reduced from that in T1 ($t_9 = 3.98$, $P < 0.005$; $t_9 = 4.89$, $P < 0.001$) indicated by “*”. Note that the RT is measured as the delay between stimulus onset and response and does not include the RSI delay.

$P < 0.0001$) indicates that the model is sensitive to temporal structure used in training, with RAND-SEQ learning measures of 2.7, 1.6 and 1.5 STS for T1, T2 and T3 respectively. This advantage for T1 derives from its regular alternating structure which ensures that a given element of the sequence will always have the same RSI. That is, T1 is coherent with A–B–C–A–D–C, which is not the case for T2 and T3.

The significant main effect for Testing ($F_{1,84} = 10.94$, $P < 0.005$) indicates that testing with the Same temporal structure as that used in training yields significantly better performance (2.3 STS) than that obtained when the training and testing temporal structures are Different (1.7 STS). The lack of significant interaction ($F_{2,84} = 2.7$, $P > 0.05$) indicates that the performance advantage for testing and training with the same temporal structure does not significantly differ between T1, T2 and T3.

3.2.2. Influence of internal state on performance

According to our hypothesis, these performance deficits following changes in temporal structure are the result of modifications that occur in the internal state representations for sequential events due to the change in the temporal structure surrounding these events. Since the model learns via an associative memory that links internal states to their corresponding responses, changes in the internal state representation for a given element will yield an impaired recall, and thus an increased reaction time. In order to test this hypothesis, one should compare the internal states during execution of the same sequence in conditions in which the temporal structure used during testing and training are the same (e.g. trained and then tested with T1, indicated as T1-T1), and in conditions in which they are different (e.g. T1-T3), in order to determine if the changes observed in RTs are accompanied by changes in the internal states. To make this comparison, the internal states are described as 25-element vectors, defined by the level of activation of the 25 State elements at the time a response is made. Note that it is thus this State vector that is associated with the appropriate response via the w^{SO} synapses. Vector pairs corresponding to the internal states for the same sequence element occurring in different temporal structures are then compared in terms of their relative amplitudes and the cosine of the angle between them.

3.2.2.1. State vector magnitudes. Fig. 4 displays the ratios of the state vector magnitudes for 20 successive stimuli obtained after training with T1 and then testing with T1, T2 and T3 (condition T1-T1, T1-T2 and T1-T3, respectively). In the T1-T1 condition there is little variation in the successive magnitude ratios, since the training and testing conditions are functionally identical. There is an increase in the variability in magnitude ratios for conditions T1-T2 and T1-T3, but the mean magnitude rests near unity (100), and there is no significant difference between the magnitude ratios for the 20 vector pairs in T1-T1 vs.

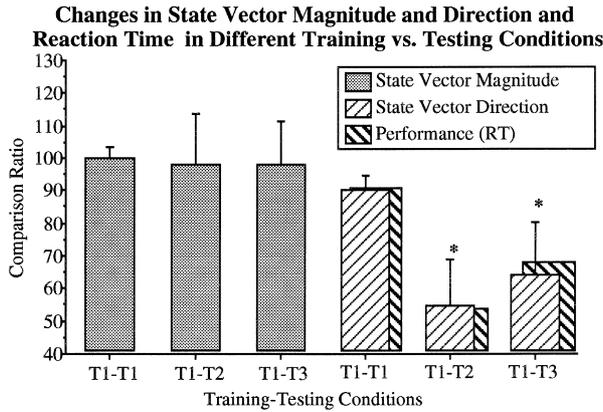


Fig. 4. Effect of temporal structure used in Training vs. Testing on State vector magnitude and direction, and their relation with Performance, for 20 successive responses. Magnitude is computed as the sum of the squares of the vector components. Direction computed as \cos^2 . Mean and S.D. for Effects of Timing used in Training vs. Testing on State vector magnitude and direction, and comparison of relative changes in RTs obtained in testing with T1, T2 or T3 after training with T1. RAND-SEQ reaction time values for models trained with T1 and tested with T1, T2 and T3 (from Fig. 3) are presented on a normalized scale for comparison with the State vector direction ratios. 100% indicates no difference between the two conditions in question. The relative learning performance $T1 > T3 > T2$ in RAND-SEQ has the same ordering as the state vector direction ratios, indicating that the performance impairments are correlated with the divergence in State vector directions.

T1-T2 or T1-T3 as measured by a set of paired t -tests: ($t_{19} = 0.70$, $P > 0.4$; $t_{19} = 0.74$, $P > 0.4$, respectively). This is of interest because it allows us to conclude that the RT increase in the conditions in which the temporal structures used for training and testing are different is not due simply to a reduction of activity in State.

3.2.2.2. State vector directions. Fig. 4 displays the mean of the \cos^2 of the angle between the respective state vectors for 20 successive stimuli. In condition T1-T1, this value remains near unity throughout the 20 sequential responses indicating little difference in the direction of the state vector pairs. Condition T1-T2 shows a divergence in the state vector pair directions, with condition T1-T3 showing a level of divergence intermediate between those for T1-T1 and T1-T2. The resulting change in direction of the state vectors obtained by comparing the T1-T1 values with those for T1-T2 and T1-T3 in a paired t -test is significant ($t_{19} = 12.5$, $P < 0.0001$; $t_{19} = 7.8$, $P < 0.0001$). This is of interest because it allows us to conclude that in the conditions in which the temporal structures used for training and testing are different, the RT increase is accompanied by a change in State coding activity, a change that violates the learned state-response associations. Indeed, the degree of change in state vector direction corresponds well with the change in performance (RAND-SEQ RTs) as displayed in Fig. 4. That is, in terms of both performance impairment and divergence of state vector directions, the relative ordering is the same: T1-T1 followed by T1-T3 and then T1-T2.

3.3. Experiment 3 (simulation): influence of coherent temporal structure on sequence learning

Expt. 2 provided evidence that changes in temporal structure can disrupt previously acquired learning. Expt. 3 extends these results to address two additional questions. First, recall that the greatest perturbation in Expt. 2 was seen in the case where a coherent temporal structure T1 (that is, one in which a given sequence element was always preceded by the same RSI) was used in training, and an incoherent one (T2 or T3) used in testing. A potential criticism is that our temporal structure effect is simply the result of a change from coherent to non-coherent temporal structure, rather than a specific learning effect. Thus, in Expt. 3, the initial training occurs with a coherent temporal structure, and then a different, but coherent temporal structure is used in the transfer test. In addition, the sequence used in Expt. 2 was of low complexity. In Expt. 3, a more complex sequence, B-C-B-D-C-A-D-A-C-D is used.

As in Expt. 2, an SRT task is used that allows the comparison between RTs for sequential vs. random series (the standard measure of learning), using the same serial structure but different temporal structures. Two coherent temporal structures T3 and T4 were studied, each consisting of a repeating series of 10 RSI values of either 0.1 or 0.5 simulation time units (20 or 100 time steps, respectively). In T3, elements A and D are always preceded by RSIs of 0.1, while elements B and C by are preceded by RSIs of 0.5, and the opposite in T4.

A set of five instances of the model were generated by using different seed values for the random number generator to initialize the weights in w^{IS} , w^{SS} , w^{OS} and w^{SO} . The five models were replicated to yield two equivalent groups of five, according to the temporal structure used for their training (T4 or T5). The two groups were trained on six blocks of 80 trials with the repeating 10-trial sequence B-C-B-D-C-A-D-A-C-D using temporal structures T4 and T5, respectively. The trained models were then tested with the same serial structure while using the same temporal structure as that used for training (blocks 6 and 8), and with the same serial but different temporal structure from the one used for training (block 7), and finally with random serial and the same temporal structure (block 9).

3.3.1. Performance effects

The model's performance in the SRT task is displayed in Fig. 5, where the RTs for responses after training in sequence blocks 6 and 8 are reduced with respect to those in the transfer block 7 that uses the different temporal structure, and also with respect to the random block 9.

3.3.1.1. Learning serial structure. The observation that overall learning occurred was examined in a multifactor ANOVA in which Training Group (T4 or T5) and Block (sequence or random) were the independent variables and

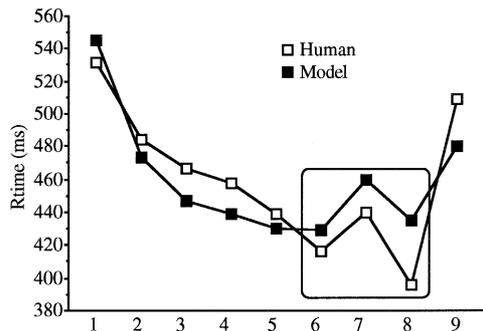


Fig. 5. Simulation and human performance in a modified SRT task. Six successive blocks of 80 trials using the sequence B–C–B–D–C–A–D–A–C–D, and one of the two temporal structures. Transfer to the different temporal structure occurs in block 7. Block 8 is the same as blocks 1–6, and block 9 is random. RTs for the simulation (Expt. 3) and human (Expt. 4) results are expressed as means for the 80 RTs in each block in ms, with the simulation equivalent to ms as calculated by the linear regression (see text). Learning is measured as the difference in RAND-Seq RTs. Performance in the transfer condition in block 7 is significantly reduced from that in blocks 1–6 and 8 (see text).

RT was the dependent variable. RTs were significantly reduced for sequence blocks (mean of blocks 6 and 8) vs. random (block 9) blocks ($F_{1,16} = 8.95$, $P < 0.01$), indicating significant learning of the serial structure of the sequence. The lack of effect for Group ($F_{1,16} = 0.111$, $P > 0.5$), and Group \times Block interaction ($F_{1,16} = 0.089$, $P > 0.5$) indicates that this learning is not influenced by which of the structures is used for training, that is, that both are equally effective.

3.3.1.2. Learning temporal structure. The observation that performance was disturbed by the change of temporal structure was examined via two measures. First, in a multifactor ANOVA in which Training Group (T4 or T5) and Block (transfer block 7 and random block 9) were the independent variables and RT was the dependent variable. There was no significant difference between RTs for the transfer condition (block 7) vs. a series with random serial structure (block 9) ($F_{1,16} = 1.721$, $P > 0.2$), indicating that the performance deficit from a change in the temporal structure is similar to that from a change in the serial structure. Neither the Group nor Group \times Block interactions were significant. These transfer results indicate that changing the temporal structure yields a significant performance deficit. This difference in performance for same vs. different temporal structure was confirmed in a paired t -test that revealed a significant difference in RTs for sequence blocks with the initial serial and temporal structures (mean of blocks 6 and 8) versus the transfer condition with the same serial but different temporal structure (block 7), $t_9 = 2.131$, $P < 0.05$.

3.4. Experiment 4 (human subjects): influence of coherent temporal structure on sequence learning

Expt. 3 provided evidence that changes in coherent temporal structure can disrupt previously acquired learning

and provides a specific prediction that can be tested in human subjects. That is, when the initial training occurs with one coherent temporal structure, and then a different, coherent temporal structure is used in the transfer test, a performance impairment should be observed in the transfer phase. To test this prediction, Expt. 3 is replicated using human subjects.

As in Expt. 3, an SRT task is used that allows the comparison between RTs for sequential vs. random series, using the same serial but different temporal structures. The same sequence and coherent temporal structures, T4 and T5 from Expt. 3, were also used in Expt. 4. Thus, elements A and D are always preceded by RSIs of 200 ms and B and C by RSIs of 1000 ms in temporal structure T4, and the opposite in T5.

Five human subjects were trained on 6 blocks of 80 trials with the repeating 10-trial sequence B–C–B–D–C–A–D–A–C–D using temporal structures T4 ($n = 2$) and T5 ($n = 3$), respectively. The subjects were then tested with the same serial structure while using the same temporal structure as that used for training (blocks 6 and 8), and with the same serial but different temporal structure from the one used for training (block 7), and finally with random serial and temporal structures (block 9). As Expt. 3 indicated that equivalent learning occurred with both of the temporal structures, we focused our analysis for the present experiment on the effects of changing the temporal structure, which was predicted by Expt. 3 to have significant effects on performance.

3.4.1. Performance effects

The human subject's performance in the SRT task is displayed in Fig. 5, where, as in the simulation, the RTs for responses after training in sequence blocks 6 and 8 are reduced with respect to those in the transfer block 7 that uses the different temporal structure, and also with respect to the random block 9.

3.4.1.1. Learning serial structure. The observation that overall learning occurred was examined in a 1-factor ANOVA in which block (sequence or random) was the independent variable and RT was the dependent variable. RTs are significantly reduced for sequence blocks (mean of blocks 6 and 8) vs. the random (block 9) block ($F_{1,9} = 5.81$, $P < 0.05$), indicating significant learning.

3.4.1.2. Learning temporal structure. The observation that performance was disturbed by the change of temporal structure was examined via two measures. First, in a 1-factor ANOVA in which Block (transfer block 7 and random block 9) was the independent variable and RT was the dependent variable. There was no significant difference between these RTs ($F_{1,9} = 2.17$, $P > 0.1$), indicating that the performance deficit from a change in the temporal structure is similar to that from a change in the serial structure. These transfer results indicate that changing the

temporal structure yields a significant performance deficit. This performance deficit linked to the change in temporal structure was confirmed in a paired *t*-test that revealed a significant difference in RTs for sequence blocks with the initial serial and temporal structures (mean of blocks 6 and 8) versus the transfer condition with the same serial but different temporal structure (block 7), $t_4 = 3.184$, $P < 0.05$.

These results in human subjects support the simulation results and the initial hypothesis: the modification of temporal structure induces perturbations in previously learned serial and temporal structures. This agrees with the possibility that serial and temporal structure are represented in a common system based on the coding properties of recurrent networks.

3.4.2. Comparison of simulated and human subject results

As displayed in Fig. 5, the simulation results of Expt. 3 are quite comparable to those of the human subjects from Expt. 4. The performance in sequence blocks 6 and 8 is superior to that of the random block 9, indicating that sequence learning has occurred. Most important for the current investigation, this performance is also significantly superior to that of the transfer block 7 in which the same sequence is presented with a different temporal structure. This demonstrates that sequence learning and performance are dependent not only on the serial structure, but on the temporal structure as well.

One difference between the simulations and the human studies is related to the temporal structure used in block 9. In the simulations the temporal structure for block 9 was the same as that used in blocks 1–6 and 8. In the human experiments block 9 had random serial and temporal structures. The effects of changed temporal and serial structure are thus combined in the comparison between sequence (blocks 6 and 8) vs. random (block 9). This has no influence, of course, on the observation that changing the serial structure (block 7) disrupts performance with respect to the initial serial structure (blocks 6 and 8).

The observation of a similarity between the simulation and human performances was tested in a linear regression on the simulated and human RTs. This analysis yielded a significant linear relation between human and model performance, with $y = 35.404x - 11.65$, $r^2 = 0.73$, $P < 0.005$, where y is simulated RT and x is RT in ms.

4. Discussion

4.1. Summary of results

It is well known that recurrent architectures are capable of demonstrating human-like performance in SRT tasks. Cleeremans and his colleagues [4] have extensively studied and documented the behavior of Elman's [10] "simple recurrent network" in different SRT task variations. We have recently studied the current model in related SRT tasks using RSI manipulations to simulate dual task performance deficits [5] as observed by Stadler [22]. The current

experiments extend such results in the temporal domain and make predictions that are tested in human subjects, again using a model initially developed to study sequence learning in the primate frontostriatal system. First, Expt. 1 demonstrated that the model is capable of discriminating temporal delays. Expt. 2 demonstrated that in addition to learning the serial order of a repeating sequence, the model also learned the temporal structure of the sequence and was sensitive to changes in this temporal structure. Expt. 3 extended these results demonstrating that even when the timing is changed in a manner that is coherent with the serial order, the model is still sensitive to this change. These results then served as a prediction for human performance. Expt. 4 tested and confirmed these predictions in human subjects, demonstrating that these subjects, like the model, after learning the serial order with one temporal structure were impaired when that temporal structure was changed. This perturbation occurred despite the fact that the serial order was identical. This result is in agreement with the earlier observations of Stadler [21,22] and support his suggestion that temporal structure is an integral part of the sequence that when changed changes the sequence itself. Indeed, while Stadler showed that learning is impaired for incoherent temporal structure, the current results go even further, demonstrating that learning is similarly impaired in the transfer to a new but coherent temporal structure. That is, the system (i.e., the model and the human subjects) is not only sensitive to the coherent vs. incoherent distinction, but can also discriminate between different coherent temporal structures.

It is of interest to recall that the model in question was originally developed to simulate the frontostriatal system in sequence learning and reproduction [8], and that it is now demonstrated to be sensitive to temporal structure in sequence processing. In this context, it has been shown that dysfunction of the frontostriatal system in Parkinson's disease yields specific impairments in the ability to discriminate and reproduce temporal structure [2,17]. This provides further evidence for a functional relation between the current model and the primate frontostriatal system [9], as both appear to exploit recurrent networks to encode serial and temporal structure. Related studies in the generation and maintenance of fixed movement frequencies have demonstrated the important role of the cerebellum in this timing-related behavior. Lateral cerebellar lesions appear to yield deficits in a central timing process, while medial lesions leave timing intact but yield impaired response generation [12]. Further studies should be made to determine the cerebellar contribution to temporal structure learning, notably in cases where the execution timing is provided internally rather than by external go signals.

4.2. Comparison with related models

Recurrent networks have been used to predict complex sequential structure in a number of domains including natural English letter and word sequences [10] and to

analyze human performance in SRT tasks as noted above [4]. In these studies, each time step involved presentation of the next element and application of the learning algorithm. That is, time was strictly linked to input/output events. Other studies have examined the capacity of neural networks to generalize over changes in the sampling rate of sequences [23,24], while retaining the property that inputs and outputs are processed at each time step.

Part of the motivation to maintain a strict correspondence between input/output processing and network time steps in some of these models stems from technical issues in the learning-related modification of recurrent synapses. After an input is presented, and several network cycles occur, an output is generated and the error is evaluated and corrected. A given connection weight in the recurrent network has contributed to the error, but in a different way on each successive cycle of information passing through this weight in the network. How is this weight's contribution to the error over these sequential time steps to be unraveled in order to implement the error-reducing learning? One way is to simply require that only one time step can pass between input/output events and weight updates, so that a weight's contribution is only made in one pass through the network. This is the case in the SRN [10], in which the learning algorithm and its application to the recurrent connections requires that at each time step a new input is provided, an output is generated and the learning algorithm is applied. Thus, time steps cannot pass independent of input, output and learning processing.

Otherwise, in order to modify connections in a way that takes into account the weight's contribution to error over a number of time steps, one must calculate the effect of that weight change for each of these time steps. One way to do this, referred to as back-propagation through time, is to "unfold" the recurrent network into an n -layered network where each layer is a copy of the original network [18,19]. One update cycle for this n -layered network is equivalent to n time steps in the original recurrent net. To change RSIs in such a network implies using different cascade depths for the different RSIs.

In recurrent back-propagation [2,20] the dynamic state of the recurrent network is allowed to settle into a fixed point after each input. The error is then calculated, and the learning algorithm applied. Again, because the connections in the recurrent network are modified by the learning algorithm, computational resources are necessary to determine each recurrent connections contribution to the final error as described above.

It is worth noting that these methods of resolving the problem of learning in recurrent networks over multiple time steps are considered to be biologically implausible, because they are not consistent with forward running time, and/or because they have excessive computational and memory storage requirements [25].

In the current model there is no learning in the recurrent part of the network, only in the feedforward connections

between the State units and the Out units, and only at the time that a response is evaluated. Thus, the above described complexity of recurrent learning is no longer an issue. Simulation time steps are mapped to real time or experiment time and are not tightly coupled to input, output and learning processing. Following the model's production of a given response, the response is evaluated and the reinforcement learning rule is applied. The network can then run for an arbitrary number of time steps (i.e. the experimenter-specified RSI) before the next input is presented. During this time the State activity is modified due to its recurrent connections, providing an explicit representation of the effects of time. Thus, while the model in question is certainly not unique in the capacity to simulate the passage of time, it is distinguished in the computational simplicity by which this is achieved [5–8].

4.3. Theoretical implications

In the current study, dissociable temporal and serial learning phenomena are examined by manipulation of temporal structure of RSIs during training and testing in simulation and human experiments. These issues can be addressed in simulation because (a) the model's state system can be updated by its recurrent connections during the arbitrary delays between a response and the subsequent stimulus (RSI), and (b) network time steps are not tightly coupled to I/O and learning processing. Thus, during the RSI, there is no input nor output processing, only an evolution of the state trajectory that thus encodes the effects of the RSI delay. This trajectory is an analog of the passage of time, and individual states in that trajectory can become associated, via a simple reinforcement learning procedure, with output events, thus providing a temporal discrimination capacity. The dynamics of this recursive net provide a next state function that takes into account not only the serial organization of sequences but their order-independent temporal aspects as well, thus time and order become inseparably coded in the recurrent network. These results are consistent with the observations of Stadler [21,22] that under conditions of RSI modification, humans can benefit from repetitive RSI structure, and do not generalize to modified or random RSI structure.

In Expts. 2 and 3, this performance dependence on temporal and serial structure derives from a single learning system. The observation that human performance in Expt. 4 is in agreement with that of the model suggests that a single mechanism might also be responsible for human sensitivity to temporal and serial structure. Though the current results are not sufficient to verify this conclusion exclusively, they clearly support it. Future research will attempt to resolve this question. In a recent, related study, Mayr [15] demonstrated that different spatial and object sequences could be simultaneously learned, which might suggest the simultaneous operation of distinct mechanisms. In his study, blocks of 72 trials were used with the imbedded spatial sequence of length 8 and the imbedded

object sequence of length 9. Thus, they were incoherent with respect to each other. This simultaneous learning of incoherent sequences provided stronger evidence that they were being learned by distinct systems. In the current experiment the temporal and serial structures were coherent with respect to each other. Future simulation and human experiments will address the simultaneous learning of temporal and serial sequences that are not coherent with respect to each other, in an effort to further examine the possibility of a common or distinct learning mechanisms.

5. Conclusion

One defining feature of a sensori-motor sequence is the order or *serial structure* of the events in the sequence. A perhaps equally important dimension of such a sequence is the set of timing relations between these elements, or the *temporal structure* of the sequence. The current research addresses the issue of how both the serial structure and the temporal structure of a sensori-motor sequence may be coded by a single system. A recurrent neural network, based on the primate frontostriatal system, is demonstrated to be such a system. The recurrent connections encode sequential state, thus providing a basis for encoding serial structure, and they also provide a system dynamics that is sensitive to the passage of time, providing a basis for encoding temporal structure. The model is exposed to several experimental conditions, including a modified version of the serial reaction time task that requires learning both serial and temporal structure. The model's performance in this task provides the basis for predictions of human performance in the same task. Indeed, the human performance confirms the model's predictions. These results suggest that recurrent networks in the primate frontostriatal system can provide a shared system for encoding serial and temporal sequential structure, though the possible intervention of the cerebellum in these timing-related functions is not excluded.

Acknowledgements

This work was supported by the Fyssen Foundation (Paris), and by INSERM Unit 94, Marc Jeannerod, Director. I thank M. Stadler for useful discussions of his work and initial simulation results, Mohammad Arzi and Alessandro Farne for useful discussion and comments, and an anonymous reviewer for useful comments on a previous version of the manuscript.

References

- [1] L.B. Almeida, A learning rule for asynchronous perceptions with feedback in a combinatorial environment, in: Proc. IEEE 1st Int. Conf. Neural Networks, San Diego, CA, June 21–24, 1987, p. 609–618.
- [2] J. Artiedia, M.A. Pastor, F. Lacruz, J.A. Obeso, Temporal discrimination is abnormal in Parkinson's disease, *Brain* 115 (1992) 199–210.
- [3] P. Barone, J.-P. Joseph, Prefrontal cortex and spatial sequencing in macaque monkey, *Exp. Brain Res.* 78 (1989) 447–464.
- [4] A. Cleeremans, J.L. McClelland, Learning the structure of event sequences, *J. Exp. Psychol.: Gen.* 120 (1991) 235–253.
- [5] P.F. Dominey, Influences of temporal organization on sequence learning and transfer: comments on Stadler (1995) and Curran and Keele (1993), *J. Exper. Psychol.: Learn., Mem. Cogn.* 1998, in press.
- [6] P.F. Dominey, Complex sensory-motor sequence learning based on recurrent state-representation and reinforcement learning, *Biol. Cybern.* 73 (1995) 265–274.
- [7] P.F. Dominey, An anatomically structured sensory-motor sequence learning system displays some general linguistic capacities, *Brain Lang.* 59 (1997) 50–75.
- [8] P.F. Dominey, M.A. Arbib, J.P. Joseph, A model of cortico-striatal plasticity for learning oculomotor associations and sequences, *J. Cogn. Neurosci.* 7 (1995) 311–336.
- [9] P.F. Dominey, J. Decety, E. Broussolle, G. Chazot, M. Jeannerod, Motor imagery of a lateralized sequential task is asymmetrically slowed in hemi-Parkinson's patients, *Neuropsychologia* 33 (1995) 727–741.
- [10] J.L. Elman, Finding structure in time, *Cogn. Sci.* 14 (1990) 179–211.
- [11] P.A. Frensch, C.S. Miner, Effects of presentation rate and individual differences in short-term memory capacity on an indirect measure of serial learning, *Mem. Cogn.* 22 (1994) 95–110.
- [12] R.B. Ivry, S.W. Keele, H.C. Diener, Dissociation of the lateral and medial cerebellum in movement timing and movement execution, *Exp. Brain Res.* 73 (1988) 167–180.
- [13] M.I. Jordan, Learning to articulate: sequential networks and distal constraints, in: M. Jeannerod (Ed.), *Attention and Performance XIII*, Lawrence Erlbaum, Hillsdale, NJ, 1990.
- [14] R. Kühn, J.L. Van Hemmen, Temporal association, in: E. Domanay, J.L. Van Hemmen, K. Schulten (Eds.), *Physics of Neural Networks*, Springer-Verlag, Berlin, 1992, pp. 213–280.
- [15] U. Mayr, Spatial attention and implicit sequence learning: evidence for independent learning of spatial and nonspatial sequences, *J. Exp. Psychol.: Learn. Mem. Cogn.* 22 (1996) 350–364.
- [16] M.J. Nissen, P. Bullemer, Attentional requirements of learning: evidence from performance measures, *Cogn. Psychol.* 19 (1987) 1–32.
- [17] M.A. Pastor, J. Artiedia, M. Jahanshahi, J.A. Obeso, Time estimation and reproduction is abnormal in Parkinson's disease, *Brain* 115 (1992) 211–225.
- [18] B.A. Pearlmutter, Learning state space trajectories in recurrent neural networks, *Neural Comput.* 1 (1989) 263–269.
- [19] B.A. Pearlmutter, Gradient calculation for dynamic recurrent neural networks: a survey, *IEEE Trans. Neural Netw.* 6 (1995) 1212–1228.
- [20] F.J. Pineda, Recurrent backpropagation and the dynamical approach to adaptive neural computation, *Neural Comput.* 1 (1989) 161–172.
- [21] M.A. Stadler, Implicit serial learning: questions inspired by Hebb (1961), *Mem. Cogn.* 21 (1993) 819–827.
- [22] M.A. Stadler, The role of attention in implicit learning, *J. Exp. Psychol.: Learn. Mem. Cogn.* 21 (1995) 674–685.
- [23] D.L. Wang, M.A. Arbib, Timing and chunking in processing temporal order, *IEEE Trans. Syst. Man. Cybern.* 23 (1993) 993–1009.
- [24] D.L. Wang, X. Liu, S.C. Ahalt, On temporal generalization in simple recurrent networks, *Neural Netw.* 9 (1996) 1099–1118.
- [25] P.J. Werbos, Backpropagation: basics and new developments, in: M.A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge, MA, 1995, pp. 134–139.