

# Attention and Awareness in Sequence Learning

Axel Cleeremans

Laboratoire de Psychologie Industrielle et Commerciale  
Université Libre de Bruxelles  
Avenue F.-D. Roosevelt, 50 — CP 122  
1050 Bruxelles — BELGIUM  
axcleer@ulb.ac.be

## Abstract

How does implicit learning interact with the availability of explicit information? In a recent series of experiments, Curran & Keele (1992) demonstrated that sequence learning in a choice reaction setting involves at least two different processes, that result in differing availability of the acquired knowledge to conscious inspection, and that are differentially affected by the availability of attentional resources. In this paper, I propose a new information-processing model of sequence learning and explore how well it can account for these data. The model is based on the Simple Recurrent Network (Elman, 1990; Cleeremans & McClelland, 1991; Cleeremans, 1993), which it extends by allowing additional information to modulate processing. The model implements the notion that awareness of sequence structure changes the task from one of anticipating the next event based on temporal context to one of retrieving the next event from short-term memory. This latter process is sensitive to the availability of attentional resources. When the latter are available, performance is enhanced. However, reliance on representations that depend on attentional resources also results in serious performance degradation when these representations become less reliable, as when a secondary task is performed concurrently with the sequence learning task.

## Introduction

In recent years, sequence learning in choice reaction settings has elicited considerable interest as a vehicle to study implicit information processing (e.g., Cleeremans & McClelland, 1991; Lewicki, Hill, & Bizot, 1988; Nissen and Bullemer, 1987; Perruchet, & Amorim, 1992). In such tasks, subjects are presented with a visuo-spatial choice reaction task, but, unknown to them, the sequence of successive stimuli is structured, so that the uncertainty about the next event may be reduced based on the constraints set by previous events. Typically, subjects exhibit detailed sensitivity about these sequential constraints, yet their explicit knowledge of the sequence remains very limited. This kind of outcome, where detectable performance

improvements are not accompanied by correlated improvements in explicit, reportable knowledge, is referred to as implicit learning (Reber, 1989). Implicit learning contrasts with explicit learning (exhibited for instance by subjects engaged in problem-solving behavior), in which processing is usually goal-directed and fully available to conscious inspection. This notion of two “modes of learning” has led many to formulate dichotomous theories of cognition in which implicit and explicit processing are generally thought to be complementary (in the sense of one mode being most efficient in the exact conditions where the other is least efficient) and independent (see Hayes and Broadbent, 1988, Reber, 1989, for examples).

However, it seems reasonable to assume that learning in general is never purely implicit or purely explicit. On the contrary, it is likely that most tasks that have been dubbed “implicit” do in fact involve—to various degrees—explicit strategies and knowledge. Goal-directed, intentional processing cannot simply be “turned off”. Many recent studies (Curran & Keele, 1993; Perruchet & Amorim, 1992, Howard, Mutter, & Howard, 1992) have begun to explore the effects of various factors relevant to the implicit/explicit distinction on performance in implicit learning tasks. These factors are maybe best described as characteristics of explicit learning, that is, (1) awareness of the material, (2) intentionality, and (3) sensitivity to the availability of attentional resources. The picture that emerges from these studies is far too complex to be discussed in detail here, but in a nutshell, all three factors may facilitate or interfere with performance in implicit learning tasks, depending on other factors such as stimulus salience or material complexity. Thus, there seems to be reasonable empirical grounds for distinguishing between learning processes that are differentially affected by the variables listed above.

Taking such an implicit/explicit dichotomy for granted, if only in a purely functional sense, one may have different theories about the nature of the representations and mechanisms that produce this dichotomy. Three positions have been expressed in the implicit learning literature. First, some authors (e.g., Perruchet & Amorim, 1992) argue that performance in implicit learning tasks does not necessarily reflect the operation of an independent implicit learning system. Rather, performance would be mostly based on explicit

---

The author is a Senior Research Assistant of the National Fund for Scientific Research (Belgium).

processing, but the resulting knowledge is fragmented enough that verbal reports probing for general information are unlikely to reveal the extent of subjects' knowledge. Other authors (e.g., Knowlton, Ramus & Squire, 1992) assume that implicit and explicit learning are supported by different memory systems, and that these systems are completely independent from each other. Implicit and explicit learning would thus proceed in parallel, but without interacting. They produce different kinds of knowledge, and are most likely to operate efficiently in contrasted settings. Finally, there may be an intermediate position where one assumes that implicit and explicit processing indeed rely on distinct memory systems, but in which some interactions between the two systems are allowed, and in which some processing resources are shared.

In this paper, I would like to explore how one may start thinking about these issues by proposing a new information-processing model of learning of sequential material in choice reaction settings. The model is based on the simple recurrent network (SRN) connectionist architecture first proposed by Elman (1990), and subsequently applied to implicit learning phenomena by Cleeremans and McClelland (1991). By contrast with the SRN and other models of sequence processing, this model uses different sources of knowledge to produce its responses. Thus, it instantiates the third theoretical position described above. To start exploring how well this kind of model is able to account for relevant sequence-learning data, I compared its performance with that of human subjects in three experiments conducted by Curran and Keele (1993). In the next section, I describe these experiments and provide an empirical context for the simulation work described in the rest of this article.

### The Curran and Keele Studies

Curran and Keele (1993) conducted four experiments that explore how implicit and explicit learning interact in a sequence-learning task. For lack of space, and because Experiment 4 is somewhat different from the others, I will not discuss it in this paper. In the first three experiments, subjects were exposed to a four-choice reaction time task divided in blocks of 120 trials each. Curran and Keele manipulated three factors. First, the material could either be random or sequential. When sequential, the target's movement followed a repeating sequence of length six (e.g., 1-2-3-2-4-3). Positive differences between RTs elicited by random blocks and RTs elicited by sequential blocks would indicate that subjects are learning about the sequence. Second, an attention-demanding secondary task could either be present or absent. When present, either a low-pitched or a high-pitched tone appeared between any two RT trials. Subjects were to count the number of high-pitched tones and report their count at the end of the block. Third, subjects could either receive typical

implicit learning instructions ("incidental subjects"), or could be told that the material would sometimes follow a sequence, and that knowing the sequence would be helpful in carrying out the main RT task ("intentional subjects"). These latter subjects were also given a minute to study the actual sequence.

All three experiments started with 2 blocks of practice on random material in dual-task conditions. In Experiment 1 (see Figure 2, top panel), a group of intentional subjects and a group of incidental subjects were first exposed to 4 single-task, sequential blocks. Next, they received one block of random material followed by another sequential block, again in single-task conditions. Learning was assessed by averaging performance on the last two sequential blocks and by subtracting this average from performance on the intermediate random block. In a second, dual-task, phase of the experiment, subjects were exposed to 2 blocks of random material, followed by one block of sequential material and a final block of random material. Learning in this second phase was again assessed by computing the RT difference between the two random blocks and the intermediate sequential block. Based on awareness reports obtained after the first, single-task phase of the experiment, incidental subjects were classified as "More aware" or "Less aware" according to how much knowledge of the sequence they were able to report. The results showed that all three groups of subjects differed in their performance on the first, single-task, phase. Intentional subjects exhibited the largest facilitation, followed by "More aware" and "Less aware" subjects. However, these differences disappeared in the second, dual-task phase, with all three groups exhibiting comparable and small facilitation effects on the sequence material. Thus, large between-groups performance differences that can be attributed to intentionality and awareness disappear when attentional resources are no longer available. Lack of attentional resources blocks the expression of explicit knowledge.

In Experiment 2 (see Figure 3, top panel), a group of intentional subjects and a group of incidental subjects were first exposed to 8 sequential blocks. Further, incidental subjects were learning under dual-task conditions, whereas intentional subjects were learning under single-task conditions. All subjects then transferred to a second phase identical with that used in Experiment 1. Both groups showed a small sequence learning effect in this second phase, but failed to differ, despite the fact that intentional subjects exhibited considerably more sequence learning than incidental subjects in the first phase of the experiment. Here, then, the lack of attentional resources is shown to block the acquisition of explicit knowledge.

Experiment 3 (see Figure 3, bottom panel) assessed whether the small sequence learning effects observed in the dual-task phases of Experiments 1 and 2 may be an artifact resulting from reaction times being near the ceiling. A single group of incidental subjects was first exposed, under dual-task conditions, to 8 blocks of

sequential material followed by one random and one sequential block. Next, subjects transferred to a second, single-task phase consisting of two random blocks, one sequence block, and a final random block. If the acquisition of explicit knowledge is blocked by the lack of attentional resources in the first phase of the experiment, one would expect equivalent, implicit, sequence learning in both phases, at least until explicit knowledge has had time to develop in the single-task phase. The results confirmed the hypothesis.

What kind of mechanism may account for these effects? I address this question in the next section.

### The DSRN model

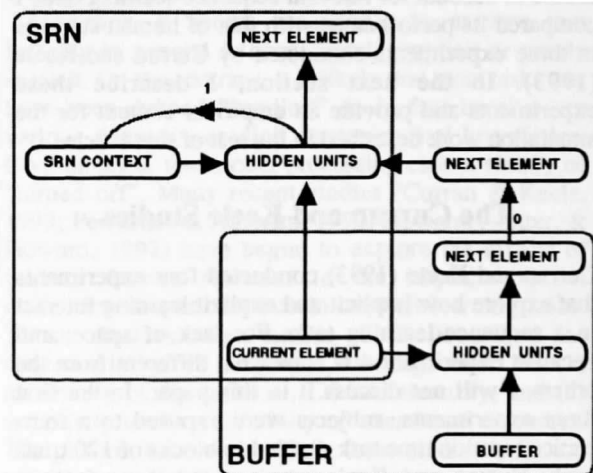
The Dual SRN (DSRN, for lack of a better name) model is based on the simple recurrent network (SRN) introduced by Elman (1990). The SRN (see Figure 1, boxed area) is able to predict successive elements of sequences presented one element at a time. Thus, on each time step, the network receives element  $t$  of a sequence as input, and is trained (using the back-propagation algorithm) to produce element  $t+1$  on its output units. To perform this prediction task, the SRN has a three-layers architecture. Fixed recurrent connections from the hidden units to a pool of context units enable the network to develop a representation of its own past activity, and to become gradually sensitive to the temporal constraints set by previous elements of a sequence in predicting the next one (see Servan-Schreiber, Cleeremans & McClelland, 1991; for a detailed analysis of learning.). Cleeremans and McClelland (1991) showed how the SRN could be used as a model of human implicit learning of sequential material in choice reaction settings. We considered the activations of the SRN's output units to represent response strength, and assumed that human subjects prepare implicitly for the next element. With these assumptions in place, the SRN model is able to account in substantial detail for the results of several sequence-learning experiments, such as those reported by Cleeremans (1993), by Lewicki, Hill, and Bizot (1988), and by Cohen, Ivry, and Keele (1990). The model implements a series of principles central to implicit learning performance, such as elementary, gradual, associative learning, processing that is local and results in fragmentary knowledge, and sensitivity to context information.

As it stands, however, the SRN model is incapable of accounting for the effects reported by Curran and Keele (1993), essentially because its architecture does not allow different sources of knowledge about sequence structure to influence processing. In the following sections, I describe how the SRN was extended to incorporate mechanisms relevant to implementation of three assumptions about (1) the effects of explicit knowledge, (2) the effects of intentionality and awareness, and (3) the effects of attentional demands.

### Effects of explicit knowledge

The central assumption upon which the DSRN model is based is that subjects who are aware of the sequence (either because they have memorized it before the task, or because they have managed to memorize it while performing the task) use their explicit memory of the sequence to tell themselves what the next stimulus will be. Thus, awareness of sequence structure changes the task from one of (implicitly) anticipating the next event based on temporal context to one of (explicitly) retrieving the next event from short-term memory. In the model, this is implemented by augmenting the basic SRN with an additional pool of input units that hold a representation of the next element of the sequence. The SRN has therefore two different ways of producing the next element. It can either develop its own representation of the temporal context, by learning how to represent successive elements as an activation vector over its context units, or it can act essentially as an encoder on the information in the pool of input units representing the next element.

How are these representations of the next element learned? Essentially, some other part of the system has to produce these representations. In the case of human subjects, this information is presumably stored and retrieved from short-term memory, using mechanisms that are outside the scope of this model. For modeling purposes, I assumed that these representations are produced by a buffer network (see Figure 1).



**Figure 1.** The DSRN model. An SRN and a buffer network are both assigned the task of predicting the next element of a sequence presented one element at a time. The SRN may use information coming from its own mechanisms for maintaining the temporal context (direct pathway) to produce the next element, or it may base its performance on information produced by the buffer network (indirect pathway). Gray-colored connections represent copy operations with the indicated time delay, and are not subject to modification by back-propagation. See the text for additional details.

The task of this network is identical with that of the SRN: To predict what the next element of a sequence will be. To do this, the buffer network has several pools of input units, each corresponding to the sequence element that occurred on a particular time step. Information shifts down the buffer as each new element is presented. The buffer network is trained concurrently with the SRN, but independently so: Error information is not back-propagated from the SRN to the buffer network.

In short, the model consists of two main processing pathways: A direct pathway, that involves connections from the "current element" and "context" pools of input units to the SRN's hidden units, and an indirect pathway involving the entire buffer network as well as connections from the "next element" pool of input units to the SRN's hidden units. Processing will tend to be distributed among these two pathways as a function of the reliability of the information flowing through each of them. For instance, if the information coming through the indirect pathway is unreliable, the model will tend to rely more on information coming from the direct pathway.

### Effects of awareness and intentionality

"Intentional" subjects in Curran and Keele's experiments were exposed to the sequence before starting the experiment. Subjects were told that the movement of the target would often follow a pattern, and were then given one minute to study the pattern. I assumed that the effects of this exposure to the stimulus material are (1) to store the sequence in memory, and (2) to bias subjects into using this knowledge during the task. To implement these assumptions in the model, the indirect pathway was trained before the actual task started. The buffer network was trained as if performing the actual experiment. In the SRN, the only trainable connections were (1) the connections from the "next event" pool of input units to the hidden units, and (2) the connections from the hidden units to the output units. The effects of this pre-training are to bias the SRN to use the information produced by the buffer network, at the expense of information coming from the direct pathway. It may seem inconsistent to assume that mere exposure to the sequence results in learning in some connections of the SRN, but this assumption is backed by recent data by Howard, Mutter, and Howard (1992), who showed that observation results in as good subsequent choice RT performance than actually performing the task.

Subjects may also become spontaneously aware of the sequence to varying degrees, depending presumably on factors such as individual differences in the allocation of attentional resources and short-term memory capacity. In the model, I again assumed that such differences could be represented by the degree to which the SRN's hidden units are sensitive to

information coming from the indirect pathway. The model was thus pre-trained in the manner described above, but on random material. This has the effect of strengthening the indirect pathway without giving the network information about the actual sequence.

### Effects of attentional demands

Cleeremans and McClelland (1991) showed how attentional effects in sequence learning could be simulated by interfering with processing, by means of adding normally distributed random noise to a network's hidden units. Because of its auditory nature, the tone-counting task used in these experiments is likely to interfere more with the storage and retrieval of traces in short-term memory than with response execution or other processes. To represent this asymmetry in the simulations, noise was added to the net input of both networks' hidden units, but to a larger extent in the buffer network than in the SRN.

## Simulations

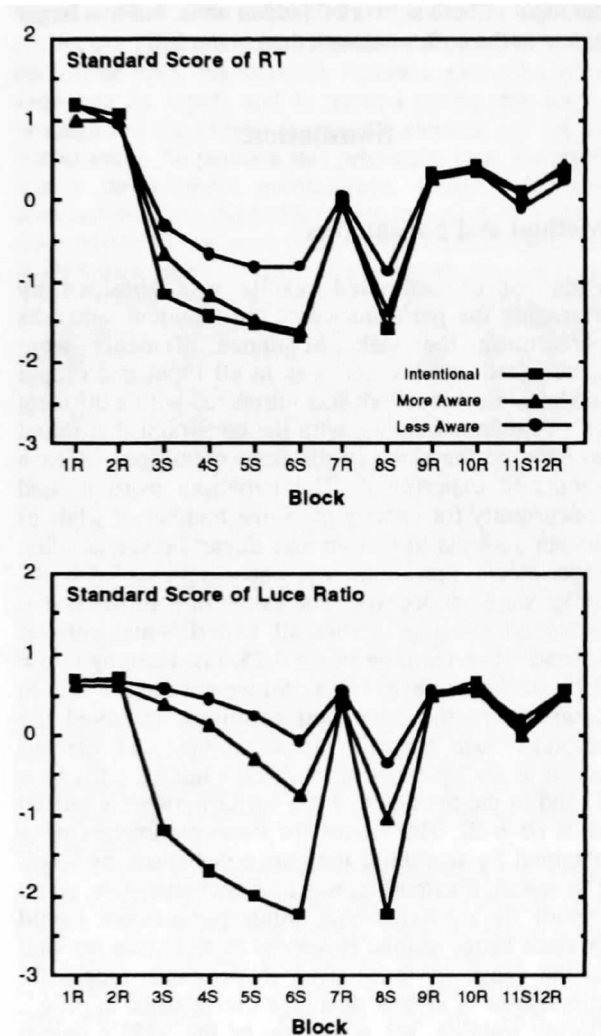
### Method and parameters

Each set of simulated results was obtained by averaging the performance of ten identical networks performing the task. Sequence elements were represented in a localist way in all input and output modules. Each network was initialized with a different set of random weights, with the constraint that initial weights be the same in different conditions of each simulated experiment. The networks were trained concurrently for exactly the same number of trials as human subjects in Curran and Keele's corresponding experiments. For simplicity, traces represented in the buffer were not decayed. The values of parameters that remained constant across all experiments were as follows: slow learning rate = 0.25, fast learning rate = 0.45, activation decay = 0.5, fast weight decay = 0.5. In cases where the simulated condition involved the secondary task, normally distributed random noise was added to the net input of the SRN's hidden units ( $\sigma = 2$ ) and to the net input of the buffer network's hidden units ( $\sigma = 8$ ). The values for these parameters were obtained by searching the parameter space by hand. The search, if extensive, was thus not exhaustive, and it cannot be excluded that other parameters would produce better results. However, these values are well in the range of those used in previous successful simulations of similar data (see Cleeremans, in press). During training, the activation of the SRN's output units was recorded. After training of a set of networks was completed, the responses corresponding to the stimuli that were actually presented were transformed into Luce ratios and averaged over the ten replications

of each simulation, to yield a single data point for each block in every condition. These averages were then subtracted from one (because stronger prediction responses correspond to faster reaction times). To facilitate comparison between human and simulated data sets, both were then transformed into standard scores with respect to the entire distribution, i.e., over all three experiments.

## Results

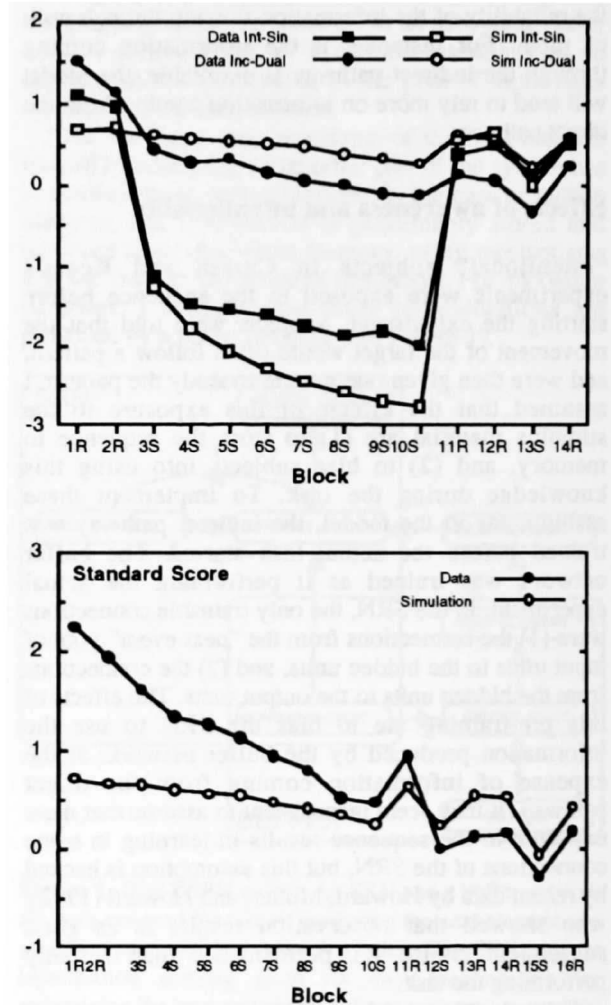
Figure 2 shows human (top) and simulated (bottom) data for Experiment 1. Both human and simulated subjects exhibit large performance differences in the single-task phase (up to block 8). These differences, that result from different degrees of explicit knowledge of the sequence in the different groups, vanish in the second, dual-task phase.



**Figure 2.** Human (top) and simulated (bottom) performance in Experiment 1 of Curran and Keele (1993). Letters following block numbers on X axis indicate type of material (Random or Sequence).

This pattern of results also obtains in the simulation because networks that were most sensitive to information coming from the indirect pathway suffer most when this information becomes unreliable because of the noise. An analysis of the quality of the information produced by the buffer network in each group revealed differences in the expected directions. The simulations are far from perfect, however. One major discrepancy with the data concerns the size of the effects in the "Intentional" and "More aware" groups. Although the fit can be improved by choosing different parameters and by limiting the standard score transformation to a single experiment, "More aware" networks tended to systematically resemble "Less aware" networks more than "Intentional" networks.

Figure 3 (top) shows results for both human and simulated subjects in Experiment 2. Here, intentional subjects learn under single-task conditions, and incidental subjects learn under dual-task conditions.



**Figure 3.** Human and simulated performance in Experiments 2 (top) and 3 (bottom) of Curran and Keele (1993). Letters following block numbers on X axis indicate material (Random or Sequence). Int.: Intentional subjects; Inc: Incidental subjects; Sin: Single-task condition; Dual: Dual-task condition.

Both groups transfer to a dual task-phase, where the large initial differences between the groups again disappear. Similar effects are exhibited by the model. Finally, the results from Experiment 3, shown in the bottom panel of Figure 3, indicate that incidental subjects who first learn under dual-task conditions and then transfer to a single-task phase exhibit about the same facilitation in both phases. The model tends to exhibit a slightly larger effect in single- than in dual-task conditions, but this was not systematic over different sets of parameters. The large discrepancy between the simulation and human data in early training blocks (also present to a lesser extent in the other simulations) results from unspecific practice effects that the model fails to simulate (see Cleeremans, 1993). Over all three experiments, the model accounts for about 75% of the variance in the data using a linear fit ( $r^2 = 0.743$ ), and for about 88% when using an exponential fit ( $r^2 = 0.876$ ) that better reflects the somewhat curvilinear relationship between the model's responses and human RTs. Other factors, such as unspecific practice effects and strategic adjustments resulting from subject's adaptation to their own changes in performance, are not represented in the model and would all tend to lower the fit even if the model was in fact perfect in simulating sequence learning.

### Discussion

In this article, I introduced a new model of sequence processing that implements the notion that learning of sequential material in choice reaction tasks may tap on two different sources of knowledge: Knowledge based on the temporal context established by previous elements of the sequence, and knowledge based on an encoding of the next element in some other part of the system. The model was found to be able to account for the major effects reported by Curran and Keele (1993). Their experimental data suggest (1) that explicit knowledge may facilitate implicit learning, (2) that subjects will tend to develop and use explicit knowledge of the material whenever possible, and (3) that both acquisition and expression of such knowledge depend on the availability of attentional resources. Other factors may also be involved. One such factor is sequence complexity. For instance, in sharp contrast with the simple repeating six-elements sequences used here, Cleeremans and McClelland (1991) used probabilistic material that was generated from a noisy finite-state grammar. Memorizing this kind of material is impossible, because it is not based on the repetition of a specific pattern. Thus, although some simple patterns tended to recur, subjects could presumably not use an explicit trace of the sequence to tell themselves what the next element would be, and indeed exhibited very little explicit knowledge of the sequence. This outcome is predicted by the DSRN model, because there would little benefit of using information coming from the indirect pathway (relative to information coming from the direct pathway) when the material is

probabilistic in nature. Further research will explore the impact of such factors on the relationship between explicit and implicit performance in sequence learning situations.

### Acknowledgments

I thank Tim Curran and Steven Keele for giving me access to their data and to their thoughts about the issues. Alain Content had a crucial suggestion and many useful comments on an previous version of this paper.

### References

- Cleeremans, A. (1993). Mechanisms of implicit learning: Connectionist models of sequence processing. Cambridge, MA: MIT Press.
- Cleeremans, A. & McClelland, J.L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, 120, 235-253.
- Cohen, A. Ivry, R.I., & Keele, S.W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 17-30.
- Curran, T., & Keele, S.W. (1993). Attentional and non-attentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 189-202.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Hayes, N. A. & Broadbent, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, 28, 249-276.
- Howard, J.H., Mutter, S.A., & Howard, D.V. (1992). Serial pattern learning by event observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1029-1039.
- Knowlton, B., Ramus, S.J., & Squire, L. (1992). Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science*, 3, 172-177.
- Lewicki, P., Hill, T., & Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, 20, 24-37.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19, 1-32.
- Perruchet, P., & Amorim, M.-A. (1992). Conscious knowledge and changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 785-800.
- Reber, A.S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219-235.
- Servan-Schreiber, D., Cleeremans, A. & McClelland, J.L. (1991). Graded State Machines: The representation of temporal contingencies in simple recurrent networks. *Machine Learning*, 7, 161-193.