LEARNING THE STRUCTURE OF EVENT SEQUENCES

Axel Cleeremans
James L. McClelland

Department of Psychology
Carnegie Mellon University

ABSTRACT

How is complex sequential material acquired, processed, and represented when there is no intention to learn ? Recent research (Lewicki, Hill & Bizot, 1988) has demonstrated that subjects placed in a choice reaction time task progressively become sensitive to the sequential structure of the stimulus material despite their unawareness of its existence. This paper aims to provide a detailed information-processing model of this phenomenon in an experimental situation involving complex and probabilistic temporal contingencies. We report on two experiments exploring a 6-choice serial reaction time task. Unbeknownst to subjects, successive stimuli followed a sequence derived from "noisy" finite-state grammars. After considerable practice (60,000 exposures), subjects acquired a body of procedural knowledge about the sequential structure of the material, although they were unaware of the manipulation, and displayed little or no verbalizable knowledge about it. Experiment 2 attempted to identify limits on subjects' ability to encode the temporal context by using more distant contingencies that spanned irrelevant material. Taken together, the results indicate that subjects become progressively more sensitive to the temporal context set by previous elements of the sequence, up to three elements. Responses are also affected by carry-over effects from recent trials. A PDP model that incorporates sensitivity to the sequential structure and carry-over effects is shown to capture key aspects of both acquisition and processing of the material.

INTRODUCTION

In many situations, learning does not proceed in the explicit and goal-directed way characteristic of traditional models of cognition (Newell & Simon, 1972). Rather, it appears that some of our knowledge and skills are acquired in an incidental and unintentional manner. Indeed, many studies (see Reber, 1989; for a review) have documented dissociations between task performance and reportable knowledge. The classic result in these experimental situations is that "subjects are able to acquire specific procedural knowledge (i.e. processing rules) not only without being able to articulate what they have learned, but even without being aware that they had learned anything" (Lewicki, Czyzewska & Hoffman, 1987).

Although controversy still pervades the field, at least two different implicit learning paradigms have yielded consistent and robust results : artificial language learning (Reber, 1967, 1989), and sequential pattern acquisition (Nissen & Bullemer, 1987; Lewicki, Czyzewska & Hoffman, 1987; Lewicki, Hill & Bizot, 1988; Willingham, Nissen & Bullemer, 1989; Cohen, Ivry & Keele, 1990). Related research with neurologically impaired patients also provides strong evidence for the existence of learning processes that do not entail or require awareness of the results or of the learning experience itself (see Schacter, 1987, for a review).

Despite this wealth of evidence documenting implicit learning phenomena, few models of the mechanisms involved have been proposed. This lack of formalization can doubtless be attributed to the difficulty of assessing subject's knowledge when it does not lend itself easily to verbalization.

Nevertheless, an important first step in the direction of understanding the relationship between awareness, attention and learning consists of attempting to identify mechanisms that account for subjects's performance in implicit learning situations. In the present paper, we report on a series of experiments inspired by Lewicki, Hill & Bizot's (1988) paradigm, and propose a detailed information processing model of the task. These experiments placed subjects in a choice reaction time task, and manipulated the sequential contingencies of the material in a novel way that allows detailed data about subject's representations of the temporal structure to be obtained.

The main results of our experiments indicate that subjects unintentionally acquire a complex body of knowledge about the temporal structure of the material. We describe a PDP model that implements a proposed mechanism to account for performance in this task. The model -- trained in exactly the same conditions as subjects -- captures key aspects of both acquisition and performance in this task. Its core mechanism implements the hypothesis that sequential structure gets induced as a direct result of an encoding of events *together* with an internal representation of the temporal context.

EXPERIMENT 1

Subjects were exposed to a six-choice reaction time task. The entire experiment was divided in 20 sessions. Each session consisted of 20 blocks of 150

This research was supported by a grant from the National Fund for Scientific Research (Belgium) to the first author and by an NIMH RSDA to the second author. Address correspondence to : Axel Cleeremans, Department of Psychology, Carnegie Mellon University, Pittsburgh PA 15213: or email to *cleeremans@psy.cmu.edu.*

trials. On any of the 60,000 trials, a stimulus could appear at one of six positions arranged in a horizontal line on a computer screen. The task consisted of pressing as fast and as accurately as possible on one of six corresponding keys. Unbeknownst to subjects, the sequential structure of the stimulus material was manipulated. Stimuli were generated using a small finite-state grammar that defined legal transitions between successive trials. Some of the stimuli, however, were not "grammatical" . Indeed, on each trial, there was a 15% chance of substituting a random stimulus to the one prescribed by the grammar. This "noise" served two purposes. First, it ensured that subjects could not simply memorize the sequence of stimuli, and hindered their ability of detecting regularities in an explicit way. Second, since each stimulus was possible on every trial (if only in a small proportion of the trials), we could obtain detailed information about what stimuli subjects did or did not expect at each step.

If subjects become increasingly sensitive to the sequential structure of the material over training, one would thus predict an increasingly large difference in the reaction times elicited by predictable and unpredictable stimuli. Further, detailed analyses of the RTs to particular stimuli in different temporal contexts should reveal differences that reflect subject's encoding of the sequential structure of the material.

Method

Subjects. Six subjects (CMU staff and students) aged 17-42 participated in the experiment. Each subject was paid $100 for his participation in the 20 sessions of the experiment, and received a bonus of up to $50 based on speed and accuracy.

Apparatus and display. The experiment was run on a Macintosh II computer. The display consisted of six dots arranged in a horizontal line on the computer's screen. Each screen position was paired with a key on the computer's keyboard, also arranged in a line ('Z', 'X', 'C', 'B', 'N', 'M'). The stimulus was a small black circle that appeared immediately below one of the six dots. The timer was started at the onset of the stimulus and stopped by the subject's response. The RSI was 120 msec.

Procedure. Subjects received detailed instructions during the first meeting. They were told that the purpose of the experiment was to "learn more about the effect of practice on motor performance". Both speed and accuracy were stressed as being important. After receiving the instructions, subjects were given 3 practice blocks of 15 random trials each at the task. A schedule for the 20 experimental sessions was then elaborated. Most subjects followed a regular schedule of two sessions a day.

Stimulus Material. Stimuli were generated on the basis of the small finite-state grammar shown in Figure 1. Finite-State grammars consist of nodes connected by labeled arcs. Expressions of the language are generated by starting at node #0, choosing an arc, recording its label, and repeating this process with the next node. The vocabulary associated with the grammar we used consists of six letters ('T', 'S', 'X', 'V', 'P', and 'Q'), each represented twice on different arcs of the grammar (as denoted by the subscript of each letter). This results in highly context-dependent transitions, as identical letters can be followed by different sets of successors as a function of their position in the grammar (For instance, 'S_1' can only be followed by 'Q', but 'S_2' can be followed by either 'V' or 'P'). The grammar was constructed so as to avoid direct repetitions of a particular letter, since it is known (Hyman, 1953; Bertelson, 1961) that repeated stimuli elicit shorter reaction times independently of their probability of presentation. Finally, note that the grammar loops onto itself: the first and last nodes, both denoted by the digit 0, are actually the same.

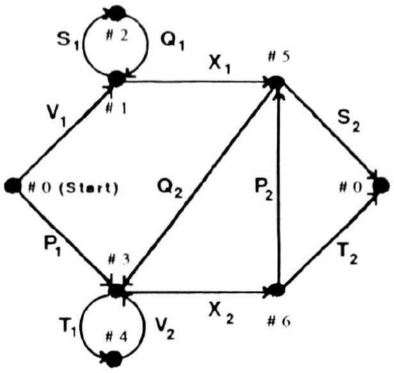

Figure 1 : The Finite State Grammar used to generate the stimulus sequence in Experiment 1.

Stimulus generation proceeded as follows. On each trial, three steps were executed in sequence. First, an arc was selected at random among the possible arcs coming out of the current node, and its corresponding letter recorded. The current node was set to be node #0 on the very first trial of any block, and was updated on each trial to be the node pointed to by the selected arc. Second, in 15 % of the cases, another letter was substituted to the letter recorded at step 1 by choosing it at random among the five remaining letters in the grammar. Third, the selected letter was used to determine the screen position at which the stimulus would appear. A 6 x 6 Latin Square design was used, so that each letter corresponded to each screen position for exactly one of the six subjects.

Design. The experiment consisted of 20 sessions of 20 blocks of 155 trials each. Each block was initiated by a "Get ready" message and a warning beep. After a short delay, 155 trials were presented to the subject. The first five trials of each block were entirely random so as to eliminate initial variability in the responses. These data points were not recorded. The next 150 trials were generated according to the procedure described above. After each block, the computer paused for approximately 30 seconds. The message "Rest Break" was displayed on the screen, along with information about subjects's performance (mean RT and accuracy for the last block, and amount earned).

Results & Discussion

Figure 2 shows the average RTs on correct responses for each of the 20 experimental sessions, plotted separately for predictable and unpredictable trials. A general practice effect is readily apparent, as well as an increasingly large difference between predictable and unpredictable trials. A two-way ANOVA with repeated measures on both factors (practice [20 levels] X trial type [grammatical vs.

ungrammatical]) revealed significant main effects of practice, $F(19,95) = 9.491$, $p < .001$; and of trial type, $F(1,5) = 105.293$, $p < .001$; as well as a significant interaction, $F(19,95) = 3.022$, $p < .001$. It appears that subjects become increasingly sensitive to the sequential structure of the material. Yet, when interviewed after the task, all subjects reported feeling that the sequence was random, and failed to report noticing any pattern in the data but small alternations (e.g. the loops on nodes #2 and #4).

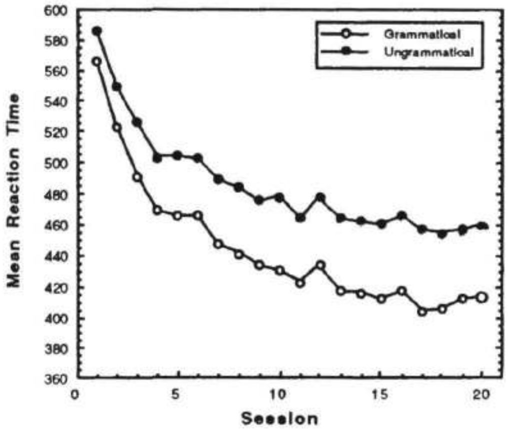

Figure 2 : Mean RTs for grammatical and ungrammatical trials for each of the 20 sessions of Experiment 1.

Accuracy averaged 98.12% over all trials. The small difference between accuracy on predictable (98.35%) and unpredictable (96.68%) trials was not significant.

One mechanism that would account for the progressive differentiation between predictable and unpredictable trials consists of assuming that subjects, in attempting to optimize their reaction times, progressively come to anticipate successive events on the basis of an increasingly large temporal context set by previous elements of the sequence. In the grammar we used, most elements can be perfectly anticipated on the basis of two elements of temporal context, but some of them require three or even four elements of temporal context to be maximally disambiguated. For instance, the path 'SQ' (leading to node #2) occurs only once in the grammar and can only be legally followed by 'S' or by 'X'. In contrast, the path 'TVX' can lead to either node #5 or node #6, and is therefore not sufficient to perfectly distinguish between stimuli that occur only at node #5 ('S' or 'Q') and stimuli that occur only at node #6 ('T' or 'P'). One would assume that subjects initially respond to the predictions entailed by the shortest paths, and progressively become sensitive to the higher-order contingencies as they encode more and more temporal context.

A simple analysis that would reveal whether or not subjects are indeed basing their performance on an encoding of an increasingly large temporal context was conducted. Its general principle consists of comparing the data with the probability of occurrence of the stimuli *given* different amounts of temporal context.

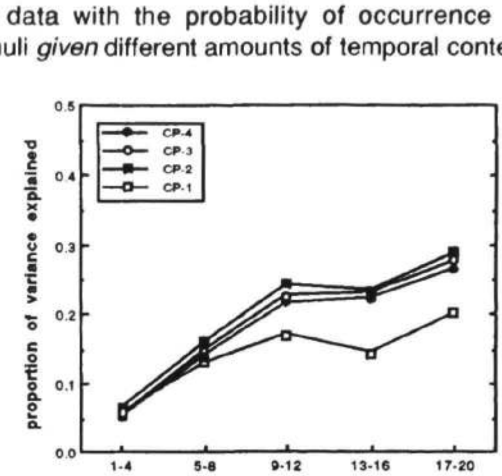

Figure 3 : Correspondence between the human responses and CPs after paths of length 1-4 during successive blocks of four simulated sessions.

First, we estimated the conditional probabilities (CPs) of observing each letter as the successor of every grammatical path of length 1, 2, 3 and 4 respectively. Next, the average RT for each successor to paths of length 4 were computed, separately for successive blocks of four experimental sessions. Finally, 20 separate regression analyses were conducted, using each of the four sets of CPs as predictor, and each of the five sets of mean RTs as dependent variable. If subjects are encoding increasingly large amounts of temporal context, we would expect the variance in the distribution of their responses at successive points in training to be better explained by CPs of increasingly higher statistical orders.

Figure 3 illustrates the results of these analyses. Each point on the figure represents the r-squared coefficient of a specific regression analysis. Points corresponding to analyses conducted with the same amount of temporal context (1 - 4 elements) are linked together. Although the overall fit is rather low (Note that the vertical axis only extends to 0.5), the figure nevertheless reveals that subjects become increasingly sensitive to the temporal context set by previous elements of the sequence. One can see that the correspondence with the first-order CPs tends to level off below the fits for the second, third and fourth orders. The fits to the second, third and fourth order paths are highly similar in part because their associated CPs are themselves highly similar.

In order to assess more directly whether subjects are able to encode three or four letters of temporal context, several analyses on specific successors of specific paths were conducted. One such analysis involved several paths of length 3. These paths were the same in their last two elements, but differed in their first element as well as in their legal successors. For example, we compared 'XTV' versus 'PTV' and 'QTV', and examined RTs for the letters 'S' (legal only after 'XTV') and 'T' (legal only after 'PTV' or 'QTV'). If subjects are sensitive to three letters of context, their response to an 'S' should be relatively faster after 'XTV' than in the other cases, and their response to a 'T' should be relatively faster after 'PTV' or 'QTV' than after 'XTV'. Averaging over all candidate contexts of this type, we found that a slight advantage for the legal successors emerged in sessions 8-12 and remained present over sessions 13-16 and 17-20 ($p < .05$). Thus there appears to be evidence of sensitivity to at least three elements of temporal context. However, no sensitivity to the first element of otherwise identical paths of length 4 (e.g. 'XTVX' vs. 'PTVX and 'QTVX') was found, even during sessions 17-20.

EXPERIMENT 2

Experiment 1 demonstrated that subjects progressively become sensitive to the sequential structure of the material and seem to be able to maintain information about the temporal context for up to three steps. The temporal contingencies characterizing this grammar were relatively simple, however, since in most cases, only two elements of temporal context are needed to disambiguate the next event perfectly.

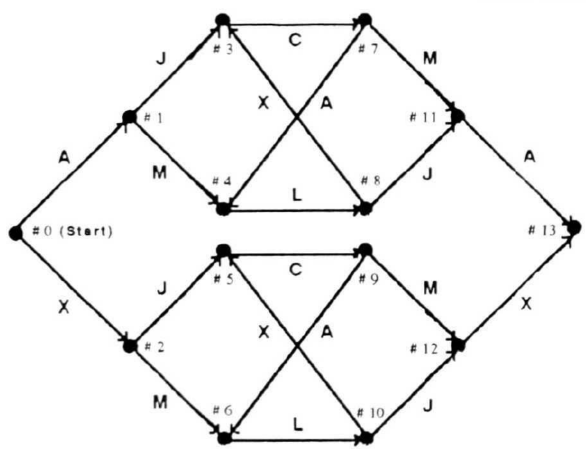

Figure 4 : The Finite State Grammar used to generate the stimulus sequence in Experiment 2.

Further, contrasting long-distance dependencies were not controlled for their overall frequency. In Experiment 2, a more complex grammar (Figure 4) was used in an attempt to identify limits on subjects' ability to maintain information about more distant elements of the sequence. In this grammar, the last element ('A' or 'X') is contingent on the first one (also 'A' or 'X'). Information about the first element, however, has to be maintained across either of the two embeddings in the grammar, and is totally irrelevant for predicting the elements of the embeddings. Further, the two embeddings are identical. Thus, in order to accurately predict the last element at nodes #11 or #12, one needs to maintain information across a minimum of three intervening steps. Accurate expectations about the nature of the last element would be revealed by a difference in the RT elicited by the letters 'A' and 'X' at nodes #11 and #12 ('A' should be faster than 'X' at node #11, and vice-versa). Naturally, there was again a 15% chance of substituting another letter to the one prescribed by the grammar. Further, in order to avoid direct repetitions between the letters that precede and follow node #13, a small loop was inserted at node #13. One random letter was always presented at this point; after which there was a 40% chance of staying in the loop on subsequent steps.

Method

Six new subjects (CMU undergraduates and graduates, aged 19-35) participated in Experiment 2. The design of Experiment 2 was otherwise identical to that of Experiment 1.

Results & Discussion

Figure 5 shows the main results of Experiment 2. They closely replicate the general results of Experiment 1, although subjects were a little bit faster overall in Experiment 2. A two-way ANOVA with repeated measures on both factors (practice [20 levels] X trial type [grammatical vs. ungrammatical]) again revealed significant main effects of practice, $F(19,95) = 32.011$, $p < .001$; and of trial type, $F(1,5) = 253.813$, $p < .001$; as well as a significant interaction, $F(19,95) = 4.670$, $p < .001$.

Accuracy was 97.00% over all trials. The difference between grammatical (97.60%) and ungrammatical (95.40%) was significant; $t(5) = 2.294$, $p < 0.5$.

Of greater interest are the results of analyses conducted on the responses elicited by the successors of the four shortest paths starting at node #0 and leading to either node #11 or node #12 ('AJCM', 'AMLJ', 'XJCM' & 'XMLJ'). Among those paths, those beginning with 'A' predict 'A' as their only possible successor, and vice-versa for paths starting with 'X'. This only holds, though, if all four letters of each path are encoded. Indeed, the sub-paths 'JCM' and 'MLJ' undifferentially predict 'A' or 'X' as their possible successors. The RTs on legal successors of each of these four paths (i.e. 'A' for 'AJCM' and 'AMLJ'; and 'X' for 'XJCM' and 'XMLJ') were averaged together and compared to the average RT on their illegal successors (i.e. 'X' for 'AJCM' and 'AMLJ'; and 'A' for 'XJCM' and 'XMLJ'), thus yielding two scores. Any significant difference between these two scores would mean that subjects are discriminating between legal and illegal successors of these four paths, thereby suggesting that they have been able to maintain information about the first letter of each path over three irrelevant steps. The mean RT on legal successors over the last four sessions of the experiment was 384.896, and the corresponding score for illegal successors was 387.847. A paired t-test on this difference failed to reach significance ($t(5) = 0.571$, $p > 0.05$).

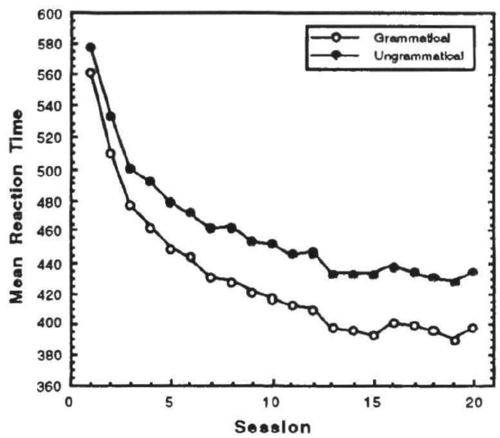

Figure 5 : Mean RTs for grammatical and ungrammatical trials for each of the 20 sessions of Experiment 2.

To sum up, subjects do not appear to be able to encode long-distance dependencies when they span 3 items of embedded independent material; at least, they cannot do so in the amount of practice used here. However, there is clear evidence of sensitivity to at least the previous two elements of temporal context.

SIMPLE RECURRENT NETWORKS

Early models of sequence processing (e.g. Estes' "statistical learning theory") have typically assumed that subjects somehow compute the conditional probabilities for all relevant statistical orders, but failed to show *how* subjects might come to represent or compute them. In the following, we present a model of sequence processing that comes to elaborate its own internal representations of the temporal context despite very limited processing resources. The model consists of a Simple Recurrent back-propagation Network ('SRN', see Elman, 1988; Cleeremans, Servan-Schreiber & McClelland, 1989).

In the SRN (Figure 6), the hidden unit layer is allowed to feed back on itself, so that the intermediate results of processing at time **t-1** can influence the intermediate results of processing at time **t**. In practice, the SRN is implemented by copying the pattern of activation on the hidden units onto a set of "context units" which feed into the hidden layer, along with the input units. All the forward-going connections in this architecture are modified by back-propagation. The recurrent connections from the hidden layer to the context layer implement a simple copy operation and are *not* subject to training.

As reported elsewhere (Cleeremans, Servan-Schreiber & McClelland, 1989), we have shown that an SRN trained to *predict* the successor of each element of a sequence presented one element at a time can learn to perform this "prediction task" perfectly on simple finite-state grammars like the one used in Experiment 1. Following training, the network produces the conditional probabilities of presentation of all possible successors of the sequence. Since all letters of the grammar were inherently ambiguous (i.e. predicting them requires more than the immediate predecessor to be encoded), the network must have developed representations of entire subsequences of events. Note that the network is never presented with more than one element of the sequence at a time. Thus, it has to elaborate its own internal representations of as much temporal context as is needed to achieve optimal predictions.

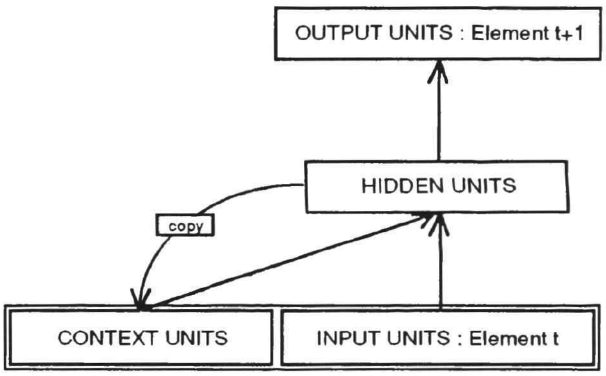

Figure 6 : The simple recurrent network.

A complete analysis of the learning process is too long to be presented here (a full account is given in Servan-Schreiber, Cleeremans & McClelland, 1988), but the key points are as follows : As the initial papers about back-propagation (e.g. Rumelhart, Hinton & Williams, 1986) pointed out, the hidden unit patterns of activation represent an "encoding" of the features of the input patterns that are relevant to the task. In the SRN, the hidden layer is presented with information about the current letter, but also -- on the context layer -- with an encoding of the relevant features of the previous letter. Thus, a given hidden layer pattern can come to encode information about the relevant features of two consecutive letters. When this pattern is fed back on the context layer, the new pattern of activation over the hidden units can come to encode information about three consecutive letters, and so on. In this manner, the context layer patterns can allow the network to learn to maintain prediction-relevant features of an entire sequence.

To model our experimental situation, we used an

SRN with 15 hidden units and local representations on both the input and output pools (i.e. each unit corresponded to one of the 6 stimuli). The network was trained to *predict* each element of a continuous sequence of stimuli generated in exactly the same conditions as for the human subjects. On each step, a letter was generated from the grammar as described above, and presented to the network by setting the activation of the corresponding input unit to 1.0. Activation was then allowed to spread to the other units of the network, and the error between its response and the *actual successor* of the current stimulus was then used to modify the weights.

During training, the activation of each output unit was recorded on every trial and transformed into Luce ratios to normalize the responses. For the purpose of comparing the model's and the subject's responses, we assumed 1) that the normalized activations of the output units represent response tendencies, and 2) that there is a linear reduction in RT proportional to the relative strength of the unit corresponding to the correct response[2].

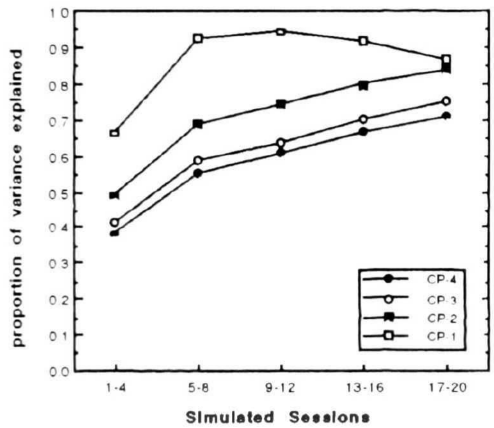

Figure 7 : Correspondence between the SRN's responses and CPs after paths of length 1-4 during successive blocks of four simulated sessions.

This data was first analyzed in the same way as for Experiment 1 subjects, and compared to the CPs of increasingly higher statistical orders in 20 separate regression analyses. The results are illustrated in Figure 7. In stark contrast with the human data (Figure 3, note the scale difference), the variability in the model's responses appears to be very strongly determined by the probabilities of particular successor letters given the temporal context. The figure also reveals that the model's behavior is dominated by the first-order CPs for most of the training, but that it becomes progressively more sensitive to the second and higher order CPs. If training was to be continued beyond 60,000 exposures, the model's responses would come to approximate increasingly higher CPs.

Figure 8 illustrates a more direct comparison between the model's responses at successive points in training with the corresponding human data. First, we computed the average RT of each letter *at each node of the grammar*. This yields a set of 42 data points (Due to noise, each of the six letters may occur at any of the seven different nodes). This analysis was conducted on RTs averaged over blocks of four successive experimental sessions, thus yielding five different sets of data. Next, a similar analysis was conducted on the model's responses. Finally, we conducted 25 separate regression analyses on these data. Each point in Figure 8 represents the r-squared coefficient of a regression analysis using the model's responses at a particular point in training as predictor and the human data as dependent variable. One would expect the model's early performance to be a better predictor of the subjects's early behavior, and vice-versa for later points in training.

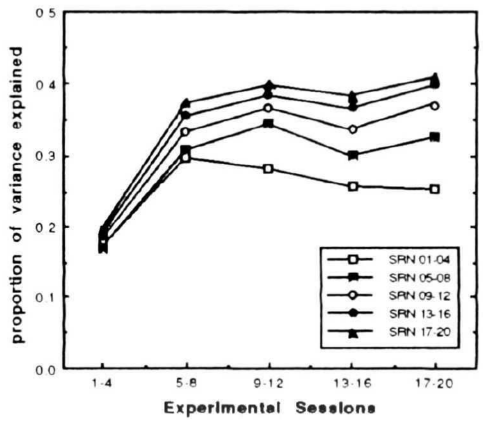

Figure 8 : Correspondence between the SRN's responses and the corresponding human data during successive blocks of four sessions of training.

It is obvious that the model is not very good at capturing subjects's behavior : the overall fit is relatively low (note that the vertical axis only goes up to 0.5), and reflects only weakly the expected progressions. Basically, too much of the variance in the model's performance is accounted for by sensitivity to the temporal context.

However, exploratory examination of the data revealed that performance in this task depends on three other factors (in addition to the conditional probability of appearance of a stimulus) :

First of all, it appears that a response that is actually executed remains primed for a number of subsequent trials (Remington, 1969). If it follows itself immediately, there is about 60 to 90 msec of facilitation, depending on other factors. If it follows after a single intervening response (as in 'VT-V', for example), there is about 25 msec of facilitation if the letter is grammatical at the second occurrence, and 45 msec if it is ungrammatical.

[2] Naturally, the second assumption is a simplification. We are currently in the process of exploring more realistic versions of this assumption.

The second factor may be related: responses that are grammatical at trial **t** but do not actually occur remain primed at trial **t+1**; the effect is somewhat weaker, averaging about 30 msec. The first two factors may be summarized by assuming that activations at time **t** decay gradually over subsequent trials, and responses that are actually executed become fully activated, while those that are not executed are only partially primed.

The third factor is a priming, not of a particular response, but of a particular sequential pairing of responses. This can best be illustrated by a contrasting example, in which the response to the second 'X' is compared in 'QXQ-X' and 'VXQ-X'. The response to the second X tends to be about 10 msec faster in cases like 'QXQ-X', where the 'X' follows the same predecessor twice in a row, than it is in cases like 'VXQ-X', in which the first 'X' follows one letter and the second follows a different letter.

This third factor can perhaps be accounted for in several ways. We have explored the possibility that it results from a rapidly decaying component to the increment to the connection weights mediating the associative activation of a letter by its predecessor. Such "fast" weights have been proposed by a number of investigators (McClelland & Rumelhart, 1985; Hinton & Plaut, 1987). The idea is that when 'X' follows 'Q', the connection weights underlying the prediction that 'X' will follow 'Q' receives an increment which has a short-term component in addition to the standard long-term component. This short-term increment is still present in sufficient force to influence the response to a subsequent 'X' that follows a immediately subsequent 'Q'.

In light of these analyses, one possibility for the relative failure of the original model to account for the data is that the SRN is partially correct, but that human responses are also affected by rapidly decaying activations and adjustments to connection weights from preceding trials. To test this idea, we incorporated both kinds of mechanisms into a second model.

This new simulation model was exactly the same as before, except for the following two changes :

First, it was assumed that pre-activation of a particular response was based, not only on activation coming from the network but also on a decaying trace of the previous activation:

$$\text{respact}[i](\mathbf{t}) = \text{act}[i](\mathbf{t}) + (1 - \text{act}[i](\mathbf{t})) * \mathbf{k} * \text{respact}[i](\mathbf{t} - \mathbf{1})$$

where act(**t**) is the activation of the unit based on the network at time **t**, and respact(**t**) is a kind of non-linear running average that remains bounded between 0 and 1. When a particular response is executed, the corresponding respact is set to 1.0. The constant **k** is set to 0.5, so that the half-life of a response activation is one time step.

The second change is simply to assume that when weights are changed by the back-propagation learning procedure, there are two components, one of which is a small (epsilon = 0.15) but effectively permanent change (i.e., a decay rate slow enough to ignore for present purposes) and the other of which is a larger (epsilon = 0.2) change that has a half-life of a single time-step.

With these changes in place, we observed that, of course, the proportion of the variance in the model accounted for by predictions based on one to four letters of temporal context is dramatically reduced (Figure 9). More interestingly, the pattern of change in these measures, as well as the overall fit, is now quite similar to that seen in the data (Figure 3).

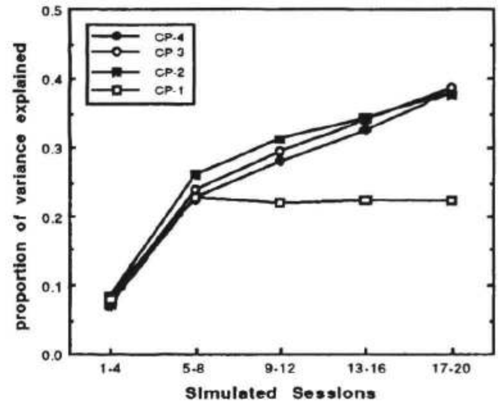

Figure 9 : Correspondence between the augmented SRN's responses and CPs after paths of length 1-4 during successive blocks of four simulated sessions.

Indeed, there is a similar progressive increase in the correspondence with the higher-order CPs, with the curve for the first-order CPs leveling off relatively early, as in the human data.

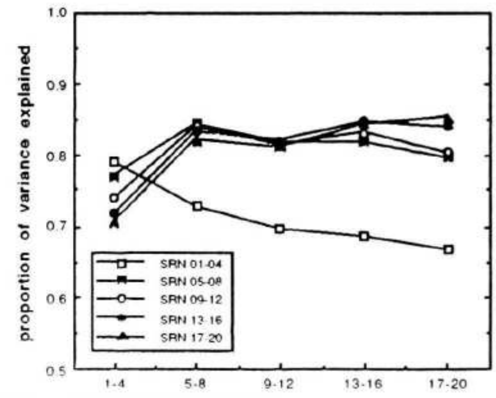

Figure 10 : Correspondence between the augmented SRN's responses and the corresponding human data during successive blocks of four sessions of training.

A more direct indication of the good fit provided by the current version of the model is given by the fact that it now correlates extremely well with the performance of the subjects (Figure 10; compare with the same analysis illustrated in Figure 8). Late in training, the model explains about 86% of the variance of the corresponding human data. Close inspection of the figure also reveals, that, as expected, the SRN's early distribution of responses is a better predictor of the corresponding early human data. This correspondence gets inverted later on, thereby suggesting that the model now captures key aspects of acquisition as well. Indeed, at every point, the best prediction of the human data is the simulation of the corresponding point in training.

GENERAL DISCUSSION

In Experiment 1, subjects were exposed to a 6-choice serial reaction time task for 60,000 trials. The sequential structure of the material was manipulated by generating successive stimuli on the basis of a small finite-state grammar. On some of the trials, random stimuli were substituted to those prescribed by the grammar. The results clearly support the idea that subjects become increasingly sensitive to the sequential structure of the material. Indeed, the smooth differentiation between predictable and unpredictable trials can only be explained by assuming that the temporal context set by previous elements of the sequence facilitates or interferes with the processing of the current event. Experiment 2 showed that subjects were relatively unable to maintain information about long-distance contingencies that span irrelevant material.

Taken together, these results suggest that in this task, subjects gradually acquire a complex body of procedural knowledge about the sequential structure of the material. They are clearly sensitive to more than just the immediate predecessor of the letter; indeed, there is evidence of sensitivity to differential predictions based on two and even three elements of context. However, sensitivity to temporal context is clearly limited: even after 60,000 trials of practice, there is no evidence of sensitivity to fourth-order temporal context. Of course, it remains possible that the subjects would eventually discover the fourth-order structure, just as the model can do.

The augmented SRN model provides a detailed, mechanistic , and fairly good account of the data. At this point it is difficult to be certain whether the model is capable of offering a complete account of all of the structure in the data. First, we have not explored the parameter space very extensively to discover whether it is possible to improve on the existing fit; and second, it is not clear just how much more systematic (as opposed to random) variance there is in the data to be accounted for.

It is often claimed that learning can proceed without explicit awareness (e.g. Reber, 1989; Willingham, Nissen & Bullemer, 1989). In our case, it appears that subjects do become aware of the alternations that occur in the grammar (e.g 'SQSQ' and 'VTVT' in Experiment 1), but have little reportable knowledge of any other contingencies. Given the fairly close correspondence of the augmented SRN with the subjects's performance, this class of model would appear to offer a viable framework for modeling this type of implicit learning.

REFERENCES

Bertelson, P. (1961). Sequential redundancy and speed in a serial two-choice responding task. *QJEP*, **13**, 90-102.

Cleeremans, A., Servan-Schreiber, D., & McClelland, J.L. (1989). Finite state automata and simple recurrent networks. *Neural Computation*, **1**, 372-381.

Cohen, A. Ivry, R.I., & Keele, S.W. (1990). Attention and structure in sequence learning. *JEP : LMC*, **16**, 17-30.

Elman, J.L. (1988). *Finding structure in time*. CRL Technical report 8801, Center for research in language, University of California, San Diego.

Hinton, G.E., & Plaut, D.C. (1987). Using fast weights to deblur old memories. In *Proc. 9th. Int. Conf. Cog. Sci. Society*.

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *JEP*, **45**, 188-196.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *JEP : LMC*, **13**, 523-530.

Lewicki, P., Hill, T., & E. Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, **20**, 24-37.

McClelland, J.L., & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *JEP : General*, **114**, 159-188.

Newell, A., & Simon, H.A. (1972). *Human problem solving*. Englewood Cliffs, N.J. : Prentice-Hall.

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

Reber, A.S. (1967). Implicit learning of artificial grammars. *JVLVB*, **6**, 855-863.

Reber, A.S. (1989). Implicit learning and tacit knowledge. *JEP : General*, **118**, 219-235.

Remington, R.J. (1969). Analysis of sequential effects in choice reaction times. *JEP*, **82**, 250-257.

Rumelhart, D.E., Hinton, G., & Willians, R.J. (1986). Learning internal representations by error propagation. In Rumelhart, D.E. and McClelland, J.L. (Eds.), *Parallel Distributed Processing, I : Foundations*. Cambridge, MA : MIT Press.

Schacter, D.L. (1987). Implicit memory : History and current status. *JEP : LMC*, **13**, 501-518.

Servan-Schreiber, D., Cleeremans, A., & McClelland, J.L. (1988). *Encoding sequential structure in simple recurrent networks*, Technical Report CMU-CS-88-183, Department of Computer Science, Carnegie Mellon University.

Willingham, D.B., Nissen, M.J., & Bullemer, P. (1989). On the development of procedural knowledge. *JEP : LMC*, **15**, 1047-1060.