Human Sequence Learning:
Can Associations Explain Everything?

Rainer Spiegel (RS272@CAM.AC.UK)
University of Cambridge, Department of Experimental Psychology,
Downing Street, Cambridge, CB2 3EB, UK

IPL McLaren (IPLM2@CUS.CAM.AC.UK)
University of Cambridge, Department of Experimental Psychology,
Downing Street, Cambridge, CB2 3EB, UK

Abstract

It will be shown that whilst a popular connectionist model, the simple recurrent network (SRN) as introduced by Elman (1990), is a very good first approximation in modeling human sequence learning, it is not, in itself, sufficient. At CogSci 2000, all five papers referring to the SRN tried to provide evidence that it is an adequate model of human performance. We will take on a more moderate position. The results of a human experiment followed by a structured interview reveal that human sequence learning is not always the kind of statistical process captured by the SRN alone.

Introduction

In cognitive science, there is an ongoing debate whether human learning should be modeled by the explicit use of rules or by figuring out statistical regularities. In addition, it has been emphasized that human learning consists of both rule and associatively-based processes (e.g. Jones & McLaren, 1999; McLaren et al. 1994).

Perhaps the most popular connectionist model used to study sequence learning is the previously mentioned SRN developed by Elman, which has become ubiquitous in the literature (Cleeremans, 1993; Elman, 1990; Elman et al. 1996; McLeod et al. 1998). A diagram of the SRN is shown in Figure 1.

The SRN is capable of learning any sequence, even sentences with hierarchical structure (McLeod et al. 1998). However, as will be seen later, this is not to be confused with learning to master any kind of sequential problem. Among other connectionist models, the SRN does not implement rules and therefore learns sequences in an associative way. Therefore, the results of the SRN should resemble human learning if humans also learn sequences associatively. In contrast, the results of the SRN might differ from human performance either if human sequence learning incorporates something more than an associative process, or if the associative mechanisms used in human sequence learning are not those employed in the SRN.

In the SRN, the network receives input from the input units and is made to predict the next step of the sequence at the output level. The SRN has connections from the hidden units to the so-called context units, which are exact copies of the hidden units one time step ago. All other connections in the network are adjustable. The context units provide the SRN with a dynamic memory, i.e. depending on the sequence position, the very same inputs can result in different predictions of the network. Each time step, the network is trained by adjusting the weights on the connections according to the backpropagation learning algorithm (first introduced by Werbos, Le Cun, Amari, Parker and now most widely accessible in Rumelhart et al. 1986). The SRN with a supervised learning algorithm was chosen for this paper, because it was considered appropriate in modeling the human experiment (as will be seen later, the subjects in the human experiment received a signal if they had made an error and no signal if they had made the correct response).

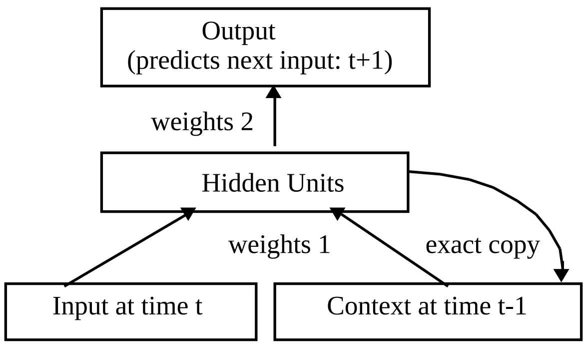

Figure 1: The simple recurrent network (SRN).

The serial reaction time task (Nissen & Bullemer, 1987) is a particularly successful paradigm used to test human sequence learning. In this type of task, the subject sits in front of a screen on which one of a number of lights flashes on at different locations. The subject is asked to press the key below the flashing light as fast as possible and the reaction time of each response is measured. While the subjects are not informed about any sequence in the stimulus material, the lights flash in a particular order. Therefore, there is a contingency in the way that the preceding stimuli predict the current one. If subjects speed up their reaction times on the sequences but are not able to provide verbal information concerning the contingencies, then their learning could be considered associative. If subjects can verbalize the contingencies, their behavior can be seen as cognitive and/or associative, depending on the degree to which subjects are able to state the underlying rules governing the sequences. A particular advantage of the task is that the stimuli of the sequence are presented one after another and as a result, the reaction times to each separate stimulus can be measured. Consequently, it can be compared to each separate output activation of the SRN used to model the task. Another advantage is that there are lights flashing on at different screen locations rather than symbols. Symbols would have a semantic meaning for subjects, which would have the consequence that subjects would probably not start the experiment without preconceptions.

Simulation studies with the SRN

The SRN has been assumed to be able to learn any kind of sequence learning task with the exception of the catastrophic interference problem, i.e. when the SRN learns a set of sequences and is then trained on a new set of sequences similar to the old ones and tested again on the old sequences, its performance is very poor. The reason for this problem, however, seems to lie in the fact that the SRN uses backpropagation as the learning algorithm which does not have an adaptive learning rate. However, McLaren (1993) has shown that an adaptively parametrised error correcting system (APECS) can avoid catastophic interference. As a result, there seems to be some chance that all sequential problems could possibly be modeled successfully with neural networks employing the SRN architecture. Although there is published work about weaknesses of the SRN (Servan-Schreiber et al., 1991; Cleeremans, 1993; Maskara & Noetzel, 1993; Timmermans & Cleeremans, 2000), those and other reported failures did not turn out to be due to the SRN architecture *per se* (Spiegel, Jones & McLaren, 2001; Spiegel & McLaren, 2001). Our new papers may lead some connectionists to argue that the SRN models human performance entirely. In order to prevent this seemingly wrong conclusion from happening, we present a sequential problem that can be solved by humans, but not by the SRN. Furthermore, a detailed network analysis about how the SRN tackles the task will be necessary to prove that this problem cannot be completely solved by the SRN. We start with an analysis of SRN performance on the task that we eventually settled on.

Procedure

The SRN was implemented using the C programming language. The task can be represented by the following grammar: (ab(c*1 ^ 3)ba) ^ (abb(c*1 ^ 3)bba). Here, the symbol ^ stands for the word *or*. For connectionist models, those letters have no semantic meaning. Expressed in words, those grammars mean: The sequence starts with the letter A. After that, the letter B follows either once or twice. Then, the letter C follows either once or three times, before the letter B appears the same number of times it had appeared earlier (once or twice) and finally, the sequence ends with the letter A. The input and output layers have a local representation for each symbol in the sequence, i.e. a = (1,0,0), b = (0,1,0) and c = (0,0,1). The network is trained with a learning rate of 0.1 and 300 hidden units.

Results

As it turned out, the network could learn this problem, a not inconsiderable achievement as there had been claims that it could not (for an overview, see Spiegel et al. 2001), but never generalized to both novel sequences with the same structure, but two c's, i.e. (abccb**a**) ^ (abbccb**b**a). It failed to predict an a after the final b in the first sequence type and failed to predict the second b after the second last b in the second sequence type (the bold letters). An entire simulation experiment with eight networks having been trained for at least half a million trials never resulted in the case where the SRN generalized to both novel sequences, even when the lowest stringent criterion was set, i.e. best match.

Modifying network parameters After this simulation experiment, 300 other SRNs were run with different numbers of hidden units (ranging from 5 to 500) on the same problem, but none of them reached sufficient generalization performance. Moreover, with less than 6 and more than 450 hidden units, the network completely lost the ability to learn the task.

Network analysis A separate network analysis focused on the ability to generalize to varying numbers of c fillers in the range between two and eleven. An interesting discovery was made: the SRN would never stabilize in generalizing to any even number of c fillers in both sequence types of the grammar displayed earlier, but it would generalize after an odd number of c fillers. In essence, the network appeared to exploit the fact that it was only trained on an odd number of c filler items by adopting cyclical patterns of activity tuned to the 1 and 3 c cases. These would also apply to other odd numbers of c items (e.g. 101), but never to even numbers of the c fillers. Moreover, whilst the performance on the trained patterns would remain stable, the ability to generalize to novel patterns would fluctuate over training trials.

The possibility of generalization during transitions between these stable states remains, however, so further tests were carried out. To assess this possibility, a very sensitive network session had to be run, with different numbers of hidden units and a test of generalization performance after each single trial. As a consequence, an SRN was implemented that would do 100,000 generalization tests during 100,000 training trials. Finally, two cases were found: after trial 39956 a 400 hidden units SRN fulfilled the best match criterion on generalizing to both two c sequences while also mastering the sequences it had been trained on. One trial later, it had lost its ability to generalize. On trial 39967 it regained this ability for one single trial, but lost it immediately from the next trial onwards and never regained it. Hence, we thought this was enough evidence to state that the SRN does not stabilize in generalizing to the sequences with two c's. Based on those findings, an experiment with human subjects was carried out to explore whether they were able to generalize to the two c case.

Human Experiment

The experiment comprised a three choice serial reaction time task. The stimulus was a circle flashing in different locations on a computer screen. The circles were arranged as a triangle, i.e. lower left corner, upper middle corner, lower right corner. The subjects were asked to press the key that corresponded to the stimulus location as fast and as accurately as possible. They were divided into an Experimental and a Control group. In both groups the order of presentation during training blocks as well as during testing followed a sequence. The sequences for the Experimental group were the same as those that the SRN had been trained on. They shall be called *consistent* sequences from now on:

(ab (c*1 ^ 3) b**a**) ^ (abb (c*1 ^ 3) b**b**a)

In the human experiment, all three letters corresponded to a particular circle, i.e. circle flashes were what the subjects saw, not letters. In the first sequence type, subjects should learn to predict the final a (bold letter) once the c had stopped and the letter b had appeared. In the second sequence type, subjects should be able to predict the second b (bold letter) once the letter c had stopped and the first b had appeared. The Control group received the same sequences as the Experimental group in 50 percent of the cases, and the following ones in the other 50 percent of the cases. They shall be called *inconsistent* sequences from now on:

(ab (c*1 ^ 3) b**b**) ^ (abb (c*1 ^ 3) b**a**a)

Because the final letter in the first sequence type and the letter before the final letter in the second sequence type had an alternative letter in 50 percent of the cases, the Control group should never be able to predict the location of the last circle in the single b case and the location of the circle before the last circle in the double b case. There were four training sessions of equal length for both Experimental and Control groups. Following that, there were two testing sessions of equal length in which both groups received 50 percent of the following consistent sequences:

(ab (c*1 ^ 2 ^ 3) b**a**) ^ (abb (c*1 ^ 2 ^ 3) b**b**a)

In addition, both groups received 50 percent of the following inconsistent sequences:

(ab (c*1 ^ 2 ^ 3) b**b**) ^ (abb (c*1 ^ 2 ^ 3) b**a**a)

The difference between the training trials and the testing trials lies in the fact that both groups receive the same sequences during testing and both groups receive the two c case which is used to test their performance to generalize to novel sequences.

The experiment aimed to investigate the following hypotheses: subjects in the Experimental group should perform faster on the critical positions (=bold letters) in the consistent sequences than in the inconsistent sequences and they should generalize to the novel sequences, because they were constructed according to the same underlying grammar. On the other hand, subjects in the Control group should show no real difference between the reaction times on consistent and inconsistent sequences. The same holds for accuracy. Subjects in the Experimental group should be more accurate on the critical positions of the consistent sequences, whereas subjects in the Control group should show more or less equal accuracy on consistent and inconsistent sequences. As a result, the (RTinconsistent-RTconsistent) differences as well as the (Errors_inconsistent-Errors_consistent) differences should be significantly higher in the Experimental group than in the Control group.

Method

Subjects The experiment comprised 30 subjects aged 18 to 40 years who were either graduate or undergraduate students at the University of Cambridge. The subjects were randomly assigned to each condition.

Apparatus The experiment was run on a Macintosh Quadra 630 computer. The subjects were seated approximately 80cm from the screen, which was roughly at eye level. The diagonal of the screen was 30cm in size. The light in the room was dimmed to a constant level.

Procedure After detailed instructions, the circles appeared on the screen. The display consisted of white outlines of three triangularly placed circles in the middle of a black background. They were two centimeters in diameter and the centers of the circles on the bottom of the triangle were approximately 5.5cm apart. The center of the upper circle was approximately 4cm apart from the centers of the two other circles. Each trial, one of the outlines would flash in such a way that it would become a solid white circle that remained on the screen until the subject responded or had not pressed a key within 4.25 seconds of the stimulus onset. After each response or after 4.25 seconds the solid circle was immediately cleared leaving only the outlines remaining. The response keys were arranged in the following way: the lower left circle corresponded to the 'v' key, the upper middle circle to the 'b' key and the lower right circle to the 'n' key. Subjects were requested to use their index-, middle-, and right finger of their preferred hand. If subjects took longer than 4.25 seconds, pressed the wrong key or a different key than the three designated, an acoustic signal indicated that the subject had made an error. Reaction time was measured in milliseconds from the stimulus onset until the key press, and the interval between a response and the onset of the next stimulus was 180ms. When one sequence finished, the screen was cleared (to the black background) and then the three outlines reappeared for 600ms until the first circle filled with white.

Block characteristics Both Experimental and Control group started with one block of 9 random circle locations in order to assess the subjects' baseline reaction time and accuracy. Then both Experimental group and Control group received 78 sequences in each of the six following blocks. Out of those six blocks, the first four blocks comprised the training trials and the last two blocks comprised the testing trials. Out of the 78 sequences in both training and testing phase, the first six of each block were not taken into the final analysis because concentration at the beginning of each block may be influenced by the preceding pause. Of the remaining 72 sequences in each block of the training phase, the Experimental group randomly received eighteen of all possible combinations of the consistent sequences (see above). The Control group randomly received nine of all possible combinations of both consistent and inconsistent sequences.

Of the 72 sequences in the testing phase, both Experimental and Control group received six of all possible combinations of the earlier mentioned consistent and inconsistent sequences per block, i.e. 12 as a whole (in order to determine the average reaction time, only the reaction times where the subjects made a correct response were counted):

The difference between these trials and those of the Control group in the training phase was that it tested whether the subjects were able to generalize to novel sequences, i.e. the ones containing the letter c twice.

Interview A structured questionnaire immediately followed the last block of the experiment. In this questionnaire, subjects were asked whether the circles had flashed on in a particular sequential order, and if so, what they can tell about the sequences.

Results

It was necessary to assess both average reaction time and number of error differences because a significant result for one of the measures does not necessarily mean much, as there could be a significant opposite trend in the other. If this was the case, the significant result in the expected direction would not reveal evidence for learning. This effect is called speed-accuracy tradeoff.

Reaction times The results of the average reaction times are considered first, i.e. the dependent variable was: inconsistent minus consistent average reaction time. An analysis of variance was carried out with the between subjects factor *group* (Experimental vs. Control group) and the within subjects factors *type* (single vs. double b case) and *number of c's* (one vs. two vs. three). Before this analysis was carried out, we tested whether the underlying assumptions for an analysis of variance with repeated measurements were met. Cochran's C test to check the equality of variances as well as the Mauchly test of sphericity (Norusis, 1994) revealed that the assumptions were entirely fulfilled. The analysis of variance revealed a significant main effect for the between subjects factor *group*, $F(1,28)=9.35$, $p<.01$, $f=.57$. The Experimental group ($M_e=34.71$, $\pm SE_e=7.58$) reveals a significantly higher reaction time difference when compared with the Control group ($M_c=1.49$, $\pm SE_c=7.79$). The size of this effect is expressed in terms of the index f (Cohen, 1988). Because f-values greater than .4 are considered large, this effect can be regarded as very strong. Cohen's f can be set in relation to more traditional effect size measures, such as the amount of variance explained by this effect ($\eta^2=.25$).

Thereafter, it was crucial to know whether subjects show a reliable effect on the one c case and the three c case and in particular whether they are able to generalize to the novel sequences with two Cs. This last will form the basis for the critical comparison with the SRN. In this experiment, the two c case showed a result on the borderline between significant and marginally significant, $F(1,28)=2.82$, $p=.05$, ($M_{e\_2c}=21.65$, $\pm SE_{e\_2c}=9.22$ vs. $M_{c\_2c}=-4.0$, $\pm SE_{c\_2c}=12.14$), providing some evidence that people do generalize to novel sequences. It is interesting to note that the sequences the subjects had been trained on obviously show larger effects, which is partly reflected in the strong main effect of the ANOVA. In order to provide a better comparison between trained and novel sequences, here are the results for the sequences the subjects were trained on:

The one c case revealed a significant effect in favor of the Experimental Group, $F(1,28)=5.29$, $p=.01$,

(M_{e_1c}=47.23, $\pm SE_{e\_1c}$=15.77 vs. M_{c_1c}=-4.75, $\pm SE_{c\_1c}$=16.24). Similarly, the three c case showed a significant effect in the same direction $F(1,28)=4.62$, $p=.02$, (M_{e_3c}=35.26, $\pm SE_{e\_3c}$=7.84 vs. M_{c_3c}=13.22, $\pm SE_{c\_3c}$=6.61). The full results are displayed in Figure 2.

Number of Errors So far, however, it could still be that the effects on the reaction times are due to a speed-accuracy tradeoff. Therefore, it was necessary to focus on the second performance measurement, i.e. the *number of errors* subjects made with consistent and inconsistent sequences. The same kind of ANOVA with the dependent variable *error differences* revealed no significant difference between Experimental and Control group, $F(1,28)=2.28$, ns., nor did any of the individual comparisons for all three numbers of c even show a descriptive trend in the opposite direction, which entirely excludes the possibility of a speed-accuracy tradeoff.

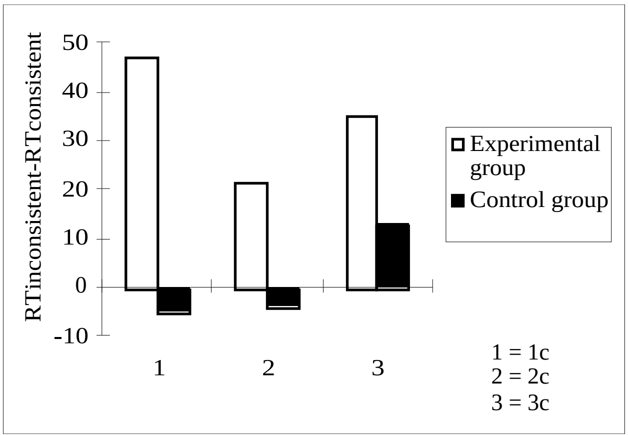

Figure 2: Average reaction time differences in humans.

Computational Model The Experimental SRN was trained with the same sequences as the Experimental group until it first reached the earlier defined performance criteria. The Control SRN was trained for 40,000 trials on the same sequences as the Control group. Both SRN's were trained with a learning rate of .1 and had 300 hidden units. The results of the network performance are displayed in Figure 3.

As can be seen in terms of the output activation differences (activation corresponding to the critical target value of the consistent sequences minus activation corresponding to the critical target value of the inconsistent sequences), the Experimental SRN has learned the task, but is not able to generalize to the two c case in any way. The Control SRN more or less resembles the human Control group in a way that it is not able to predict the next position, because it equally favors consistent and inconsistent sequences.

Structured Interview In order to get a better idea of how people solved this task, it was necessary to find out to what extent subjects verbalize the sequences. Here it was crucial to find out how many people in the Experimental group verbalized the rule that the number of b's after the c's was dependent on the number of b's before the c's. Only fourteen out of fifteen subjects in the Experimental group were able to take part in the interview. All fifteen subjects in the Control group answered the questions.

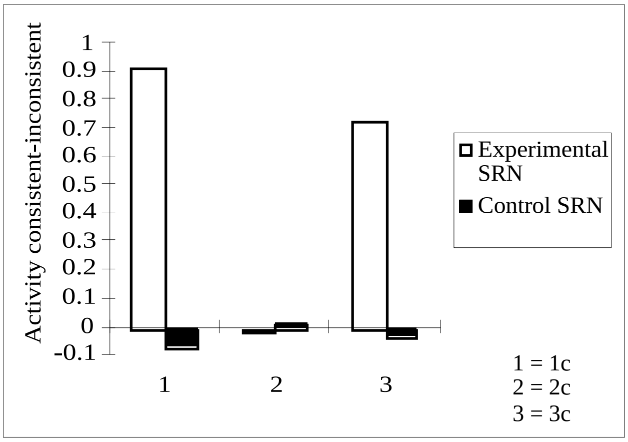

Figure 3: Activity differences between consistent and inconsistent output units on the critical letters.

Almost all of the subjects in the Experimental as well as the Control group verbalized that the circles flashed in a sequential order, while none of the subjects in either group was able to verbalize how many times a circle flash corresponding to the grammatical letter C had appeared.

None of the subjects in the Control group verbalized any dependency between the b flashes before and after the c flashes, which was expected, because they were independent of each other in the Control group. In the Experimental group, two out of fourteen subjects verbalized this dependency. A close look at the reaction time differences of those two subjects revealed that their reaction time differences were more pronounced in their effect (1c=107.32, 2c=51.37, 3c=82.55) than the average RT-differences of the remaining twelve subjects in the Experimental group (1c=32.56, 2c=16.72, 3c=25.12) and none of the error differences point in the wrong direction, which excludes the possibility of a speed-accuracy tradeoff. Interestingly, when performing the same ANOVA with the exception of the two subjects who verbalized the rule, there is still evidence of learning the trained sequences: $F(1,25)=4.71$, $p<.05$, $f=.43$, $\eta^2=.16$, but no longer evidence for generalization to the novel sequences, of $F(1,25)=1.54$, ns., (M_{e_2c}=16.71, $\pm SE_{e\_2c}$=10.91 vs. M_{c_2c}=-4.0, $\pm SE_{c\_2c}$=12.14). In other words: Successful learning but generalization failure occurs when the subjects who represented the rule are left out of the analysis, which corresponds to the results of the associative SRN. As a result, one tentative conclusion is possible here: the human ability to represent the rule (which is absent in the SRN) may have led to successful generalization in humans. There is, of course, another equally valid interpretation of this finding, however, and that is that those subjects who had learned the sequences (associatively) best were the ones who subsequently became aware of them and were able to induce the rules governing them.

Discussion

This experiment provides evidence that humans and the SRN may differ when dealing with a particular sequential problem. Whilst the SRN is capable of learning all of the sequences presented in the training set, it cannot generalize to particular sequences that were constructed according to the same underlying grammar. Furthermore, a logical analysis of the inner representations of the network revealed the reason why it does not learn the problem: the way the network represents the temporal order of the sequences in its context layer makes it impossible to solve the complete generalization problem.

There is some evidence that humans approach the problem in a different way. The structured interview suggested that some humans can induce the underlying rules of the sequences and the results of those subjects in the experiment provide evidence that they may make explicit use of them when generalizing to novel sequences.

However, it would be hard, and possibly premature to uncouple the rule-based and the associative component of humans who have participated in this task. Those two subjects who represented the task in a rule-based way have probably started with associative learning and later somehow induced the rule. The suggestion here is that there is a real possibility that there are associative mechanisms available to humans which interact with cognitive processing to determine task performance. On the basis of our results, we consider the purely associative SRN a very powerful model that may be able to learn many kinds of sequence, but does *not* simulate the human ability to generalize in this experiment.

Acknowledgments

We gratefully acknowledge the support from Emmanuel College Cambridge and the Cambridge European Trust.

References

Cleeremans, A. (1993) *Mechanisms of Implicit Learning. Connectionist Models of Sequence Processing.* Cam- bridge, MA: MIT-Press.

Cohen, J. (1988) *Statistical Power Analysis for the Be- havioral Sciences* (2nd ed.). Hillsdale, N.J: Erlbaum.

Elman, J.L. (1990) Finding structure in time. *Cognitive Science*, *14*, 179-211.

Elman, J.L., Bates, E.A., Johnson, M.H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996*) Rethink- ing Innateness: A connectionist perspective on de- velopment.* Cambridge, MA: MIT-Press.

Jones, F.W. & McLaren, I.P.L. (1999) Rules and Asso- ciations. *Proceedings of the Twenty-First Annual Con- ference of the Cognitive Science Society, pp. 240-245.* Mahwah, NJ: Erlbaum.

Maskara, A. & Noetzel, A. (1993) Sequence Recogni- tion with Recurrent Neural Networks. *Connection Science*, *5*, 139-152.

McLaren, I.P.L. (1993) APECS: a solution to the sequen- tial learning problem. *Proceedings of the Fifteenth Annual Convention of the Cognitive Science Society, pp. 717-722.* University of Colorado at Boulder.

McLaren, I.P.L., Green, R.E.A. & Mackintosh, N.J. (1994) Animal Learning and the Explicit/Implicit Distinction. In N.C. Ellis (Ed.), *Implicit and Explicit Learning of Languages.* London: Academic Press.

McLeod, P., Plunkett, K. & Rolls, E.T. (1998) *Introduction to Connectionist Modelling of Cognitive Processes*. Oxford: Oxford University Press.

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

Norusis, M.J. (1994) *SPSS Advanced Statistics 6.1.* [Handbook]. Chicago: SPSS Inc.

Rumelhart, D.E., Hinton, G.E. & Williams, R.J. (1986). Learning Internal Representations by Error Propagation. In D.E. Rumelhart & J.L. McClelland (Eds.), *Parallel distributed processing*, Vol. 2. Cambridge, MA: MIT-Press.

Spiegel, R. & McLaren, IPL. (2001). Recurrent Neural Networks and Symbol Grounding. *Proceedings of the International Joint INNS/IEEE Conference on Neural Networks*, Washington, D.C.

Spiegel, R., Jones, F.W. & McLaren, IPL. (2001). The Prediction-Irrelevance Problem in Grammar Learning. *Proceedings of the International Joint INNS/IEEE Conference on Neural Networks*, Washington, D.C.

Servan-Schreiber, D., Cleeremans, A. & McClelland, J.L. (1991) Graded State Machines: The representa- tion of temporal contingencies in simple recurrent networks. *Machine Learning, 7,* 161-193.

Timmermans, B. & Cleeremans, A. (2000) Rules versus Statistics in Biconditional Grammar Learning: A Simulation based on Shanks et al. (1997). In L.R. Gleitman & A.K. Joshi (Eds.), *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*. Mahwah, N.J.: Erlbaum.