# Reaction Times and Predictions in Sequence Learning: A Comparison

Ingmar Visser and Maartje E.J. Raijmakers and Peter C.M. Molenaar[1]
{op_visser,op_raijmakers,op_molenaar}@macmail.psy.uva.nl
Developmental Processes Research Group
Department of Psychology, University of Amsterdam
Roetersstraat 15, 1018 WB Amsterdam
The Netherlands

### Abstract

In the simple recurrent network (SRN) model, proposed by Cleeremans and McClelland (1991) to describe implicit sequence learning, the distinction between reaction time and prediction of the next trial is somewhat blurred. That is, the reaction time of the network is taken to be inversely proportional to the activation value of the corresponding node. In a prediction task the prediction would also be directly derived from the activities of the output nodes. In order to investigate the difference between ability to predict following stimuli and reaction times, we study implicit sequence learning in a similar vein as done by Cleeremans and McClelland (1991), using a slightly less complex grammar than they did. In addition we ask subjects to guess where the next stimulus will be at randomly chosen trials during the learning process. Results show a direct correspondence between fast reaction times and correct predictions.

## Introduction

Implicit learning has been studied for over thirty years starting with Reber (1967). Only recently attention has been given to modeling this kind of learning behavior in detail, mainly using neural networks. Specifically simple recurrent networks have been used successfully by Cleeremans and McClelland (1991) to model subjects' behavior on learning sequences that are generated by a finite state automaton, in fact the very same automaton that was used by Reber (1967).

Many different paradigms have been developed for studying implicit learning behavior. One characteristic that divides those paradigms is the way in which they assess the possession of implicit knowledge. In this paper two such measures, reaction times and predictions, are studied. In implicit learning research the sequential implicit learning paradigm has become increasingly popular and with that the use of reaction time as the primary measure of performance (see for example Nissen & Bullemer, 1987; Cleeremans & McClelland, 1991; Seger, 1997). We used an augmented sequence learning paradigm in which a direct comparison between reaction times and predictions was possible.

[1]The authors wish to thank their students Sander van Duyn, Wanda Toxopeus, Stijn Gooskens, Thijs de Jongh & Edibe Tali for valuable help in setting up this experiment and collecting and analyzing the data.

## Sequence learning

One of the more recent paradigms to study implicit learning is so-called sequence learning. Subjects are typically offered sequences of stimuli that are formed according to some (formal) rule(s). The only thing subjects have to do is press some key that corresponds to the current stimulus. For example when the stimuli are just zeros and ones, the current stimulus could be formed by taking the xor of the preceding two stimuli. It is now interesting to see if subjects implicitly learn this rule. This is measured by comparing RTs on correct trials, that is trials on which the current stimulus is in fact the xor of the two preceding trials, with RTs on incorrect trials, where the current stimulus is *not* the xor of the two preceding trials.

Cleeremans and McClelland (1991), using this paradigm, had their subjects learn an endless sequence of stimuli genrated by a finite state grammar. To determine the effects of implicit learning, they assessed reaction times, and found these to be decreasing as subjects got more training. Similar studies have been done where, instead of measuring reaction times, performance was assessed by asking subjects to predict the next stimulus after having seen an initial segment of a string. However, few studies have investigated the exact relation between RTs and prediction performance in implicit learning. The present study aims to gain insight into this relation by analyzing RTs and prediction performance simultaneously.

In this context the work of Cleeremans and McClelland (1991) on the SRN model for implicit learning, is of interest. They use the SRN model to predict RT performance of subjects by taking the reaction time of the network to be inversely proportional to the activity of the output unit corresponding with the correct response[2]. The activity of the 'correct' output unit can thus be interpreted as a measure of anticipation of the position of the next stimulus. This anticipation in turn can be used to make *predictions* of the next stimulus as well; in this case the position corresponding with the output unit with the highest activity has the highest probability of being predicted. This means that the

[2]Note that this doesn't leave the possibility for incorrect responses. This is not a big problem, however, since typically incorrect responses are very seldom because of the simplicity of the task.

SRN model predicts a negative relation between prediction performance and RTs, with the RTs decreasing as prediction performance gets better. The aim of the present study is to test this hypothesis empirically.

## Experiment

To assess the relation between RTs and prediction of stimuli directly we did a sequence learning experiment in which the standard series of RT trials was interspersed with prediction trials at which subjects had to guess where the next stimulus would come. A similar procedure is proposed by Jimenez, Mendez, and Cleeremans (1996) which they named the *continuous* generation task. The main difference between this procedure and other generation tasks is that no feedback is given on the correctness of the prediction; rather, after subjects have made their prediction the next stimulus of the sequence is presented with the same response-stimulus interval as between consecutive RT trials.

Subjects were given a four-choice RT task, consisting of a total of 4800 trials divided in twenty blocks of 240 trials each. The blocks were split into two sessions that were presented on two consecutive days. Unknown to subjects the sequence of stimuli followed a pattern that was generated using the finite state grammar which is described below. Because of the rather complex structure of the sequences generated with such a grammar subjects were presented with 4800 trials. There were two types of stimuli: RT trials and prediction trials. At the RT trials subjects were asked merely to reproduce the current stimulus by pressing the appropriate key. At the prediction trials subjects were asked to predict the next stimulus by pressing the appropriate key. Each block of 240 trials was divided into subblocks of four types: grammatical RT, random RT, grammatical prediction and random prediction. The switch from one subblock to the next was not marked so subjects were unaware of the existence of these subblocks. The sequence of stimuli in the random subblocks was unrestricted but for the fact that no two consecutive stimuli could be the same, which would lead to undesired speed-up of responses due to priming.

The random trials are used as a control condition, accomodating for possible effects of motor training, as well as for additional effects of subjects gaining implicit knowledge of the grammar. This design provides the possibility to assess the effects of implicit learning, by comparing RTs and prediction performance in the grammatical trials to those obtained in the random trials. Note that this is a within subjects design, so that each subject is his own control group (i.e., the performance of each subject on the grammatical trials is compared with that same subject's performance on random trials). The prediction of an inversely proportional relation between RTs and prediction performance, as derived from the SRN model (Cleeremans & McClelland, 1991), translates into three statistical hypotheses. The first is an interaction effect of condition and time on the RTs: If implicit learning occurs, RTs should decrease more for the grammatical trials than for the random trials. The second is an interaction effect of condition and time on prediction performance: over time, prediction should improve for the grammatical trials, but not for the random trials. Finally, on trials leading to correct predictions, RTs

### Method

**Subjects** Twenty-four subjects, undergraduates at the Department of Psychology of the University of Amsterdam, participated in this experiment. They received both course credits and money for participation. On top of that they could earn bonusses for fast and accurate responses.

**Procedure** At the start of the experiment subjects were told that in this task both accuracy and speed were important. The experiment started with two small blocks of trials that were not recorded to familiarize the subjects with the task. Each block consisted of four subblocks: 20 random RT trials, 100 grammatical RT trials, 100 grammatical prediction trials and 20 random prediction trials. In the RT subblocks only reproduction of the stimuli was asked of the subjects; in the prediction subblocks RT trials were interspersed with prediction trials. At the end an extra block was added in which the order of the random and grammatical trials was reversed to test whether the order of the subblocks could influence the results.

To enable a more direct comparison between prediction and reaction times, an extra block was added in which the series of trials for both the RT subblock and the prediction block was identical. In this way it is possible to directly compare the RT on a given trial with the prediction made on the very same trial.

**Stimulus material** The sequence of stimuli in the grammatical subblocks was generated from the finite state grammar in figure 1. Sequences are produced by this grammar in the following manner:

1. Start in state #1 and randomly choose one of the arcs leaving that state while noting the letter corresponding to the followed arc.

2. In the next state repeat this process of choosing an arc and noting the corresponding letter

3. The process ends when state #7 is reached and the process starts over again to create strings of unbounded length.

**Display** As can be seen in figure 1 the alphabet of the grammar consists of four letters. The letters were translated into screen positions as shown in figure 2. In the grammatical RT subblock subjects were exposed to 100 trials; in each trial the $\times$-symbol appeared in one of the quadrants of the computer display and the subjects had to press the corresponding key on the numerical keypad on the keyboard. The keys 1,2,4 and 5 on the numerical keypad were used to ensure that the spatial configuration of the response keys matched the spatial configuration of the stimulus positions

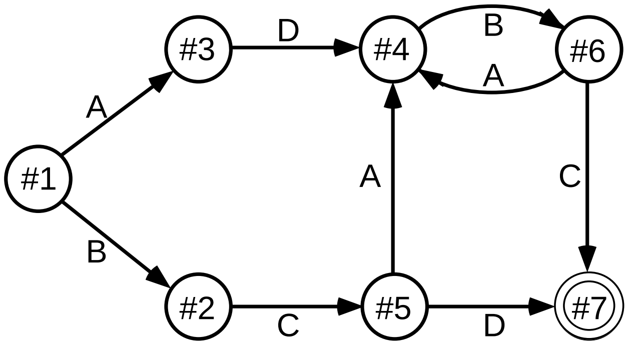


Figure 1: Finite state automaton used to generate strings for sequence learning experiments. A string is formed by starting in state #1 and then randomly choosing one of the arcs leaving that state meanwhile noting the letter corresponding to that arc. Continue stepping from state to state until the end state #7 is reached; from there the process starts over again from state # 1.

on the display. Subjects were instructed to hold their index finger over the middle of the four keys and press the appropriate key only with the index finger.

**Exit interviews** All subjects were asked a series of questions after the experiment was completed to assess whether subjects had acquired any explicit knowledge of the grammatical sequence.

## Results

The data of one of the subjects was was not included in the analyses, because the subject had too many errors in three consecutive blocks due to misplacing the index finger over the numerical keypad. Comparison of the last two blocks revealed that the order of the subblocks, random before grammatical or vice versa, did not significantly influence reaction times.

**RT trials** Grammatical RTs decreased from 404.7 ms at the beginning of the experiment to 342.6 ms at the end; random RTs decreased from 414.2 ms to 370.3 ms. The mean RTs are displayed in Figure 3.

The first hypothesis predicts that RTs decrease more for the grammatical trials than for the random trials. In order to test this hypothesis, RTs were averaged over subjects and over two consecutive blocks. A repeated measures ANOVA with two within factors, block (10 levels) $\times$ grammaticality (2 levels), indicates a significant interaction between grammaticality and blocks: as predicted, grammatical trial RTs decreased more over time than did random trial RTs, $F(9, 198) = 3.87; p < 0.001$. The analysis also yielded significant main effects for grammaticality and training: grammatical trial RTs were significantly smaller than the random trial RTs, $F(1, 22) = 59.49; p < 0.001$, and RTs became faster over blocks for both grammatical

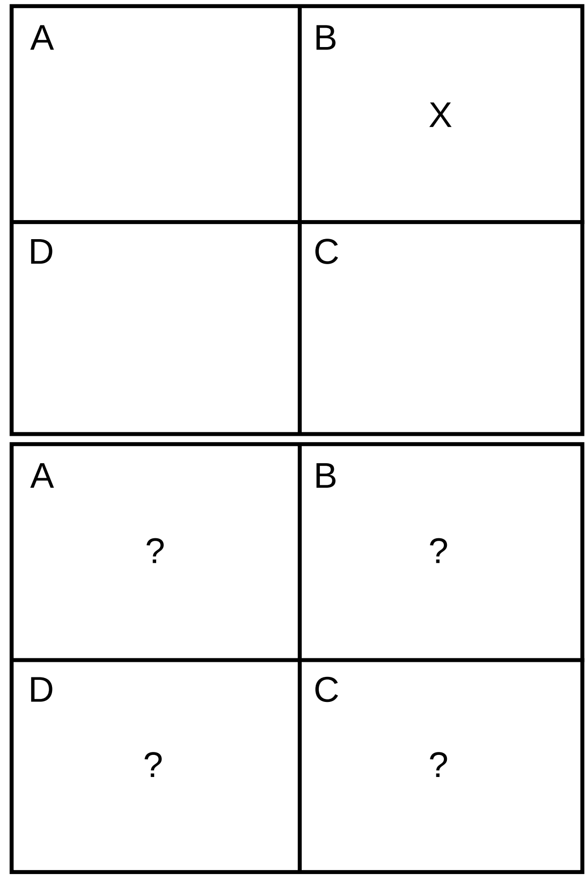


Figure 2: The top panel shows the computer display for the RT trials. Subjects have to press the key corresponding to the quadrant of the screen where the $\times$ is shown. In the bottom panel the screen lay-out for a prediction trial: all quadrants have a question mark and subjects have to choose whatever letter they think will occur next. The letters in the top-left corner of the quadrants were not part of the actual display.

and random trials, $F(3.78, 198) = 25.75; p < 0.001$ with Greenhouse-Geisser correction for non-homogeneous variances.

**Prediction trials** The percentage of correct predictions in grammatical subblocks increased from 33.6 % at the beginning to 52.2 % at the end of the experiment. The corresponding percentages for the random predictions are 30 and 34 % respectively. The proportions of correct responses on predictions for both randam and grammatical subblocks are displayed in Figure 4.

The second hypothesis states that prediction performance should improve over time for the grammatical trials, but not for the random trials. In line with this prediction, a significant interaction between blocks (time) and

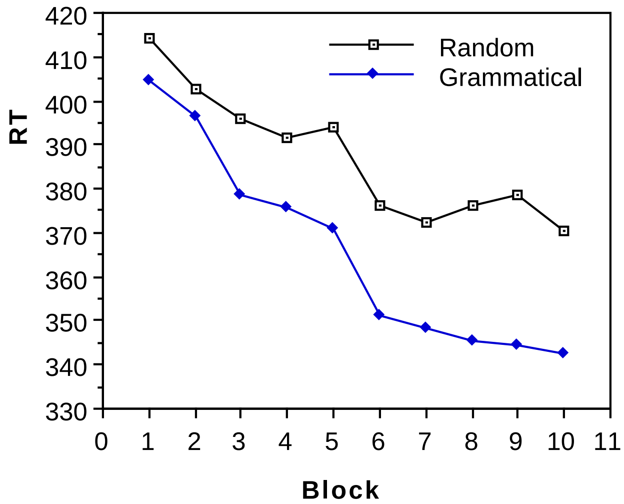


Figure 3: Mean reaction times for grammatical and random trials. Means are averaged over two consecutive blocks, $N = 23$.

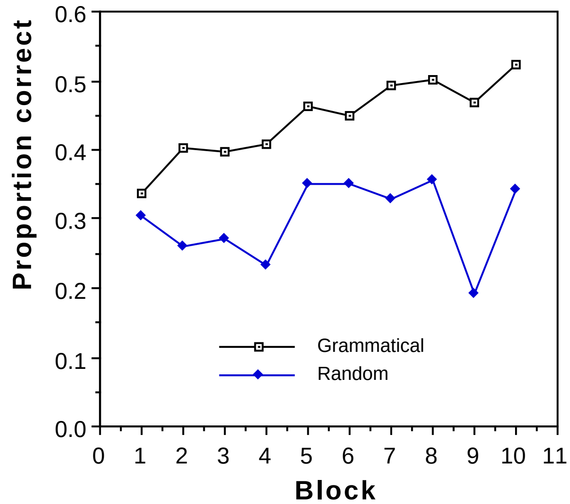


Figure 4: Proportion correct predictions of grammatical and random prediction trials, $N = 23$.

grammaticality was found, $F(7.9, 198) = 2.25; p = 0.027$, showing that the grammatical predictions did show more improvement over time than did the random predictions. More specifically, there was no improvement over time for the random trial predictions when analyzed seperately, $F(1, 228) = 0.845; p = 0.359$, as was to be expected.

**Prediction and RT trials: comparison** To compare performance on prediction and RT trials directly we added a block of trials in which the strings used for the RT trials and for the prediction trials were identical. Table 1 shows the mean RTs for correctly and incorrectly predicted items in this added block of trials. An anova with one within factor (correct vs. incorrect) confirms that correct predictions correspond to fast RTs, $F(1, 22) = 6.44; p = 0.019$.

Table 1: Mean reaction times for correctly and incorrectly predicted trials.

| Prediction | mean | sd |
|---|---|---|
| correct | 360.96 | 49.64 |
| incorrect | 389.97 | 30.30 |

**Exit interviews** Subjects were asked whether they noticed anything particular in the sequence of stimuli. Although some subjects felt there was some ‘regularity’ in the seqeunce, none of the subjects could specify this, except for three subjects that said that the subsequence $AB$ occured rather frequently. This is the subsequence in the grammar which corresponds with the loop between the two top right nodes in Figure 1.

### Discussion

The results show that implicit learning occurs: subjects give faster responses on grammatical trials than on random trials and this effect becomes larger towards the end of the experiment. Secondly, subjects gradually get better at predicting following stimuli due to training as well. Thirdly, as expected, smaller RTs correspond with a better ability to predict the following stimulus.

## Models of sequence learning

Cleeremans and McClelland (1991) applied the SRN to implicit sequence learning. The SRN successfully describes subjects’ growing sensitivity to dependencies between successive stimuli. The success of the SRN model is due to its ability to capture the ‘statistical constraints’ inherent in the sequence of stimuli. The SRN model also correctly predicts, at least in a qualitative manner, the inverse relation between RTs and the proportion of correct predictions as we have shown above. A drawback of the SRN model is that it is not very well suited for describing individual differences. The SRN model construes implicit sequence learning in subjects as statistical learning. Subjects first grow sensitive to first order frequencies of symbols, then to second order freqeuncies, that is bigram frequencies, then third order frequencies et cetera. Individual differences in both the learning process and the resulting implicit knowledge base, that is knowledge of frequency constraints, are not brought out by the model. Below we will describe how hidden Markov models can be used to model individual behavior of subjects.

### The hidden Markov model

Hidden Markov models, henceforth HMMs, are also called stochastic finite automata since they are equivalent to finite automata where the arcs between states have probabilities corresponding to them. The only restriction is that the probabilities on the arcs leaving a particular state should sum to one. This resemblance to finite automata is the reason for exploring the possibility of applying HMMs to implicit learning. Before presenting results of fitting HMMs to subjects' data we give a short introduction to HMMs.

Hidden Markov models have mainly been used in speech recognition applications such as Schmidbauer, Casacuberta, Castro, and Hegerl (1993), Chien and Wang (1997) although recently more psychologically oriented applications have come up as well such as in action learning (Yang, Xu, & Chen, 1997). The main reason that HMMs are used in speech recognition is that they are espially well suited for capturing temporal dependencies in a series of utterances which then helps in identifying phonemes. This feature can be used to model the temporal dependencies that are inherent in the series of stimuli that are typically used in implicit learning.

More formally a HMM consists of a the following elements (notations adapted from Rabiner (1989)), also see figure 5 for clarification:

1. a set of states $S_i, \; i = 1, \ldots, N$
2. a set $V$ of observation symbols $V_k, \; k = 1, \ldots, M$
3. a matrix $A$ of transition probabilities $a_{ij}$ for moving from state $S_i$ to state $S_j$
4. a matrix $B$ of observation probabilities $b_j(k)$ of observing symbol $V_k$ while being in state $S_j$
5. a vector $\pi$ of initial state probabilities $\pi_i$ corresponding to the probability of starting in state $S_i$ at $t = 1$

The equations describing the dynamics of the model are as follows:

$$
\begin{aligned}
S_{t+1} &= A\,S_t + \zeta_{t+1} \\
O_{t+1} &= B\,S_t + \xi_{t+1},
\end{aligned}
$$

where $S_t$ is the hidden process and $O_t$ is the observed process; $\zeta_{t+1}$ and $\xi_{t+1}$ are zero mean martingale increment processes, cf. Elliott, Aggoun, and Moore (1995, p. 20) for further details. A hidden Markov process then is a Markov process with multiple indicators for each (hidden) state. By substituting $S_t$ by its definition in terms of $S_{t-1}$ in the defining equation for $O_{t+1}$ it is easily seen that in fact $O_{t+1}$ is dependent on all foregoing observations back to $O_1$. Hence, at any given point observations can depend on all foregoing observations. This is in contrast with a normal Markov model where the next observation only depends on the current observation.

## Characterizing sequence learning behavior

Fitting a hidden Markov model is in fact the inverse of producing a sequence of stimuli from a finite state automaton: finding the best automaton to describe a given sequence of observations. This procedure can be applied to any kind of sequence of categorical observations and hence also to a sequence of responses in a sequence learning experiment. In simulation studies we have shown that in fitting a HMM the right automaton can be induced from the data (Visser, Raijmakers, & Molenaar, accepted for publication). That is, having generated a sequence from the grammar used in the experiment we found the HMM in Figure 5 exploratively.

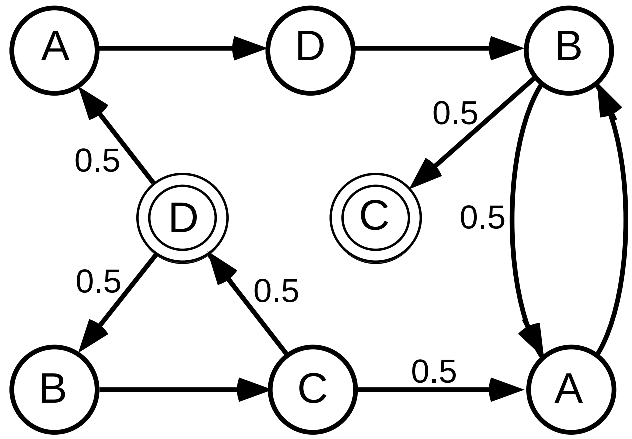


Figure 5: Representation of a hidden Markov model. This model produces exactly the same sequences as the grammar we used in the experiment with equal probabilities. Sequences are generated in the same manner as in FSAs: start in one of the states on the left with letter A or B, then follow the arcs leading from those states. A sequence ends when one of the accepting states is reached, that is the two states with the double circle around them. From there the process continues by going to one of the starting states again. For the accepting state with the letter D the arcs are drawn to the start states. For reasons of clarity the arcs from the accepting state with the C are left out. The arcs leading from one state to the next have probabilities corresponding to them which are given in the figure for some of the arcs.

**Sequence learning data** In the prediction subbblocks of the experiment subjects were presented with question marks on the screen at random points in the sequence of stimuli. In between the prediction trials normal RT trials were presented. For each subject this resulted in a sequence of responses consisting of the trials that were presented on the screen interspersed with their own predictions about the position of the next stimulus.

In order to characterize sequence learning we fitted HMMs on these sequences of responses. To bring out the learning we fitted separate HMMs on the initial and final segments of the sequence of responses. Both segments consisted of 500 trials. We expected to see a rise in number of hidden states of the model from beginning to end; that is, we expected subjects to gradually build a more complex model of the grammar underlying the sequence of stimuli. A rise in number of states would reflect subjects' growing sensitivity to the structure of the sequence. For two subjects we indeed found such a rise in the number of states from two states at the start of learning to four states at the end of learning. Overall however, results were inconclusive. This is, we think, mainly due to the fact that only a small proportion of the series of responses that were analyzed were actually produced by the subject. Of the series of 500 trials that the HMMs were fitted on, only 125 were produced by the subjects, the others were generated by the finite state automaton and only *reproduced* by the subjects. As a consequence, of all the responses only a quarter could were useful in discriminating between beginning and end of the learning phase. Hence the low power of the test. In future research it would be useful to have longer sequences of freely generated responses to which HMMs can be fitted more reliably.

## Conclusion

In sequence learning both RTs and prediction have been used as a measure of performance. The results of this experiment show that when measured simultaneously it is possible to relate directly improvement in prediction performance and improvement in RT performance on grammatical trials. The direct comparison shows what is to be expected: fast RTs are indicative of the subjects' level of anticipation of the next trial and on the same count result in correct predictions. With this study it is also shown that prediction is possible even in a fairly complex rule system, that can not be verbalized by subjects.

The SRN model has proved to be a valuable model for describing the learning processes inherent in implicit sequence learning. However the model does not seem especially suitable to describe individual subjects' behavior. Therefor we introduced the hidden Markov model as a stochastic counterpart of the FSA to characterize individual learning behavior. Since hidden Markov models are an excellent means of describing temporal denpencies between responses they are in principle well suited for describing implicit learning behavior. Our results with fitting HMMs are promising in that we can reliably estimate them on the kind of sequences that are generally used in implicit sequence learning. It would be interesting to do experiments where subjects generate longer sequences of responses instead of the single predictions they made in the experiment described in this paper.

## References


Chien, J. T., & Wang, H. C. (1997). Telephone speech recognition based on bayesian adaptation of hidden Markov models. *Speech Communication*, *22(4)*, 369–384.

Cleeremans, A., & Jimenez, L. (1998). Implicit sequence learning: The truth is in the details. In M. Stadler & P. Freuch (Eds.), *Handbook of Implicit Learning* (pp. 323–364). Thousand Oaks (Ca): Sage Publications.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *JEP: General*, *120*, 235–253.

Elliott, R. J., Aggoun, L., & Moore, J. B. (1995). *Hidden Markov models: Estimation and control.* New York: Springer Verlag.

Jimenez, L., Mendez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *JEP: Learning, Memory and Cognition*, *22–4*, 948–969.

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

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of IEEE*, *77(2)*, 267–295.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 317-327.

Schmidbauer, O., Casacuberta, F., Castro, M. J., & Hegerl, G. (1993). Articulatory representation and speech technology. *Language and Speech*, *36(2)*, 331–351.

Seger, C. A. (1997). Two forms of sequential implicit learning. *Consciousness and Cognition: An International Journal*, *6*(1), 108-131.

Visser, I., Raijmakers, M. E., & Molenaar, P. C. (accepted for publication). Confidence intervals for hidden Markov model parameters. *British journal of mathematical and statistical psychology*.

Yang, J., Xu, Y., & Chen, C. S. (1997). Human action learning via hidden Markov model. *IEEE Transactions on Systems, Man and Cybernetics*, *27(1)*, 34–44.