

Implicit Learning of Serial Reaction Time Tasks: Connectionist vs. Symbolic Models

Ron Sun (rsun@cecs.missouri.edu)

Department of CECS, University of Missouri, Columbia, MO

Chris Terry (cterry@cs.ua.edu)

Department of Computer Science, University of Alabama, Tuscaloosa, AL

Abstract

This paper describes simulations of implicit learning experiments. It compares simulations using connectionist models with existing simulations using symbolic models. It addresses an interesting issue raised by proponents of symbolic models, namely, the claim that implicit learning is better modeled by symbolic rule learning programs. This paper revisits such an issue by quantitatively comparing connectionist simulations with symbolic ones, in the context of the serial reaction time task of Lewicki et al (1987). This comparison is interesting because it helps to clarify, to some extent, some long standing confusions compounded by many claims and counter-claims. It also points to the idea of hybrid connectionist and symbolic models.

Introduction

There have been a variety of simulations of implicit learning experiments. The majority of them are connectionist, while some are symbolic. Proponents of symbolic models, however, raised some interesting issues. They claim that implicit learning is “better modeled by symbolic rule learning programs” (Ling and Marinov 1994), and symbolic models are better for “not only conscious processing but also unconscious processing”, based on some limited success of modeling Lewicki’s experiments (Lewicki et al 1987) using C4.5, a decision tree learning algorithm developed by Ross Quinlan. In this paper, we will revisit such claims by quantitatively comparing connectionist simulations with symbolic ones, especially in the context of the serial reaction time (SRT) task of Lewicki et al (1987). This comparison is interesting because it helps to clarify, to some extent, some long standing confusions compounded by many similar claims and counter-claims.

Some background is in order here. With regard to the serial reaction time task specifically, Cleeremans and McClelland (1991) simulated an SRT task involving nondeterministic grammars. They employed a recurrent backpropagation network that saw one position at a time but developed an internal context representation over time that helped to predict next positions. The model succeeded in matching human data in terms of degrees of dependency on preceding

segments in a sequence (i.e., conditional probabilities). However, their success was obtained through introducing additional mechanisms for several types of priming (e.g., short-term weight changes and accumulating activations). They did not deal with capturing directly the reaction time data of their subjects.

Ling and Marinov (1994) simulated the SRT data from Lewicki et al (1987), using a symbolic decision tree learning algorithm (i.e., C4.5). Their model produced data on quadrant prediction accuracy and, based on the data, they succeeded in matching the human reaction time data, using a transformation that included a power function (for capturing un-specific learning). However, they did not attempt the match without such a power function.

Similarly, Lebiere et al (1998) simulated data on SRT using ACT-R. The simulation was based on a combination of instance-based learning implemented in ACT-R and a set of hand-coded, symbolic, a priori rules. A fit with data was found.

It has been claimed, on the connectionist side, that a vast majority of human cognitive activities (i.e., implicit processes), including “perception, motor behavior, fluent linguistic behavior, intuition in problem solving and game playing — in short, practically all skilled performance”, can only be modeled by subsymbolic computation (connectionist models), and symbolic models can give only an imprecise and approximate explanation to these processes (Smolensky 1988). It has also been claimed, on the symbolicist side, that “one and the same algorithm can be responsible for conscious and nonconscious processes alike”, or even that implicit learning “should be better modeled by symbolic rule learning programs” (Ling and Marinov 1994). See also Fodor and Pylyshyn (1988).

This argument is in a way similar to what has been happening in relation to modeling past-tense acquisition in children (including how to capture the U-shaped curves in the process). For arguments and counter-arguments concerning advantages or disadvantages of connectionist and symbolic models in relation to past-tense acquisition, see, for example, Christiansen et al (1999). In this paper, let us look into the simulation of implicit learning specifically.

Simulating Lewicki et al (1987)

The Model. CLARION is a general cognitive architecture capable of simulating a variety of cognitive data (see Sun 1999, Sun et al 2001). The model consists of two levels: an implicit learning level (the bottom level) that learns using trial-and-error processes through a combination of backpropagation and reinforcement learning (i.e., Q-learning) algorithms (Watkins 1989, Sun and Peterson 1998); an explicit learning level (the top level) that learns explicit rules through on-line hypothesis testing based on information from the implicit level (the bottom level), which was termed “bottom-up learning” in Sun et al (2001). Bottom-up learning proceeds by first constructing the most specific rule that corresponds to a “good” decision made by the bottom level, and then refining (generalizing) it through examining the result of applying the rule, mainly through the use of an “information gain” measure that compares the success ratios of various modifications of the current rule.

Note that for this type of task, there is no significant amount of explicit learning in human subjects, as shown by Lewicki et al (1987). (This point can be controversial; more discussions later.) Correspondingly, in the model, the top level is not relevant (practically speaking). A parameter in the model is set in accordance with domain characteristics, which prevents explicit learning from occurring. The parameter concerns the minimum frequency of repetitions of a pattern in order for the afore-mentioned explicit learning to occur (see Sun et al 2001 for details of explicit learning).

In the bottom level, a simplified learning process was employed, again in accordance with domain characteristics, in which the backpropagation algorithm was used but temporal credit assignment (Q-learning) was not. This was because in this task, subjects predicted one position at a time, with immediate feedback, and thus there was no role for temporal credit assignment (Q-learning) to play.


Specifically, in the model, $Q(x, a)$ computes the likelihood of the next position a , given the information concerning the current and past positions x . The actual probability of choosing a as the current prediction (of the next position) is determined based on the Boltzmann distribution, as is common for Q-learning:

$$p(a|x) = \frac{e^{Q(x,a)/\alpha}}{\sum_i e^{Q(x,a_i)/\alpha}}$$

where α controls the degree of randomness (temperature) of the decision-making process. This method is also known as Luce’s choice axiom (Watkins 1989).

The error signal used in the learning algorithm is as follows:

$$\Delta Q(x, a) = \alpha(r + \gamma \max_b(y, b) - Q(x, a)) = \alpha(r - Q(x, a))$$



(1) identification trials (2) matrix scanning trials

Figure 1: Identification trials and matrix scanning trials.

where x is the input, a is one of the outputs (the predictions), $r = 1$ if a is the correct prediction, $r = 0$ if a is not the correct prediction, and $\gamma \max_b(y, b)$ (which represents Q-learning of Watkins 1989) is set to zero because of the fact that only one-step prediction is involved in this task. This process amounts to a variant of the backpropagation algorithm. Based on the error measure, the backpropagation algorithm is applied to adjust internal weights (which are randomly initialized before training) as usual.

The Task. The task was based on matrix scanning: Subjects were to scan a matrix, determine the quadrant of a target digit (the digit 6) and respond by pressing the key corresponding to that quadrant. Each block consisted of six identification trials followed by one matrix scanning trial. In identification trials, the target appeared in one of the quadrants and the subject was to press the corresponding keys. In matrix scanning trials, the target was embedded among 36 digits in a matrix, but the subject’s task was the same. See Figure 1. In each block of 7 trials, the actual location of the target in the 7th (matrix scanning) trial was determined by the sequence of the 6 preceding identification trials (out of which 4 were relevant). 24 rules were used to determine the location of the target on the 7th trial. Each of these rules mapped the target quadrants in the 6 identification trials to the target location on the 7th trials in each block. 24 (out of a total of 36) locations were possible for the target to appear. The major dependent variable was the reaction time on the 7th trial in each block.

The whole experiment consisted of 48 segments, each of which consisted of 96 blocks of 7 trials (so there were a total of 4,608 blocks). During the first 42 segments, the afore-mentioned rules were used to determine target locations. However, on the 43rd segment, a switch occurred that reversed the outcomes of the rule set: the upper left was replaced by the lower right, the lower left was replaced by the upper right, and so on. The purpose was to separate unspecific learning (e.g., motor learning) from prediction learning (i.e., learning to predict the target location on the 7th trial).

The Data. The reaction time data of three sub-

jects were obtained by Lewicki et al (1987). See Figure 2. Each curve showed a steady decrease of reaction times up until the switch point. At that point, there was a significant increase of reaction times. After that, the curve gradually lowered again.

The Model Setup. The input to the model contained (a sequence of) 6 elements, with each element having 4 possible values (for 4 different quadrants).¹ The output contained the prediction of the 7th element in a sequence. Thus, 24 input units (representing 6 elements, with 4 values each), 24 output units (one for each possible location of the 7th element of a sequence), and 18 hidden units were used. We tried various parameter settings. The best learning rate was 0.5, with a momentum term of 0.2. The model was trained by presenting the stimulus materials in the same way as the human experiment described by Lewicki et al (1987), without any further embellishments or repetitions of the materials.

Because in this experiment there were a total of 6^4 sequences with each consisting of 7 elements, the setting was too complex for subjects to discern the sequence structures explicitly, as demonstrated in human experiments by various explicit tests done by Lewicki et al (1987). Computationally, no explicit representation of knowledge could be extracted in the model, because the large number of sequences entailed that there were no sufficient repetitions of any particular sequence throughout the experiment, which prevented the model from coming up with any rule. The density parameter was set to be $1/50$; that is, at least one repetition (of a sequence) was necessary every 50 blocks in order to maintain an explicit rule (Sun and Peterson 1998). In this task, there were 4,608 presentations of sequences and there were $6^4 = 1296$ different sequences, and thus on average the repetition rate of any sequence was only 0.0007716.

A note concerning the existence of explicit knowledge in implicit learning tasks in general is in order here. It has been hotly debated whether there is a significant amount of explicit knowledge involved in implicit learning tasks and whether explicit knowledge can account for the performance in such tasks (see Sun et al 2001, 2002 for reviews). Without getting into details of such debates, we can reasonably believe that, although explicit knowledge may be present in many implicit learning tasks, the existence of a significant amount of such knowledge is highly unlikely in the task of Lewicki et al (1987), given the complexity of their task setting (Sun et al 2002). This was the assumption made in our simulation, although it would not change our main points even if this assumption was dropped.

The Match. We were able to create an error rate curve going downwards (averaged over 10 runs to

¹A sequence of 6 elements was assumed to be within the capacity of the short-term working memory.




Figure 2: Matching Lewicki’s data using linear transformation. See Figure 4 for parameter values.

ensure its representativeness), resembling Lewicki et al’s reaction time curves. The model reached 100% accuracy at a point before the switch. See Figure 3. The question is how we should translate the error rates into reaction times.

One way of translation is through a linear transformation from error rates to reaction times (as commonly used in existing simulation work); that is, $RT_i = a * e_i + b$, where RT_i is the reaction time (in ms), e_i is the error rate, and a and b are two free parameters. For each set of human data, we adjust the parameters to match the data. One possible interpretation of linear transformation is that it specifies the time needed by a subject to search for a target item as well as the time needed by a subject to respond to a target item without searching (through correctly predicting its location); that is,

$$RT_i = ae_i + b = b(1 - e_i) + (a + b)e_i$$

where b is interpreted as the time needed to respond to an item without searching (since $1 - e_i$ is the probability of successfully predicting the location of a target item), and $a + b$ is interpreted as the time needed to respond to an item by first searching for it and then responding to it (after finding it). So, instead of relying on additional functions (as in e.g. Ling and Marinov 1994), this method relies only on error

rates to account for human performance in terms of reaction times.

Another way of generating reaction time from prediction accuracy is through a formula used by Ling and Marinov (1994):

$$RT_i = t_1(1 - e_i) + t_2e_i + B\alpha^{-t}$$

where t_1 is the time needed to respond when there is no search (using correct predictions), t_2 is the time needed to respond when search is necessary, B is the initial motor response time, and α is the rate at which the motor response time decreases. The third term is meant to capture unspecific practice effects (mostly resulting from motor response learning). In other words, in this formula, we separate the motor response time from the search time and the prediction time (as represented by t_1 and t_2 respectively). Note that, if we set $B = 0$, we have $t_1 = b$ and $t_2 = a + b$ and thus this equation becomes the same as the previous one. This formula takes into account the independent nature of motor learning, as separate from the learning of prediction of target locations. However, it involves two more free parameters.

Using the linear transformation (without the power function), we generated three sets of data from the error rate curve reported earlier,² one for matching each human subject in Lewicki et al (1987), using different a and b values for each subject.³ As shown in Figure 2, the model outcome fit the human data well up to the point of switching (segment 42). When the switch to a random sequence happened, the model's reaction times became much worsened whereas the subject's reaction times suffered only slightly (although in a statistically significant way).⁴

²When curve fitting, we used Microsoft Excel Solver to find the best parameter values (e.g., a and b in $a * x + b$) such that the difference between the model data and the subjects data was minimized. Microsoft Excel Solver uses the Generalized Reduced Gradient nonlinear optimization algorithm.

³The error rate curve reported earlier was the best curve and happened to match all three subjects approximately equally well after the transformation (with different parameters for each subject). Note that Ling and Marinov (1994) also used a single error rate curve to match different subjects with different parameters for transformation.

⁴We tried many different parameters but discovered that the size of the jump tended to vary little (unless the match as a whole was bad). It is clear, from our experiments with different settings of the parameters, that, if the model learns the sequences perfectly before the switch (as is the case with our model), the model data inevitably have huge jumps. However, the more of the sequences it does not learn, the flatter the curve and the less the jump. Although this may model Subjects 1 and 3 satisfactorily, Subject 2 has a large drop in reaction time early on which is best matched by having the model increase its accuracy in a rapid manner.




Figure 3: The model prediction errors.

Subject 1:	$a = 320$	$b = 430$
Subject 2:	$a = 860$	$b = 620$
Subject 3:	$a = 170$	$b = 530$

Figure 4: Parameters for matching Lewicki's data.

Next, we added the power function, in order to compare our simulation with Ling and Marinov (1994). After adding the power function, we re-fit the parameters. The effect of adding the power function was that we reduced the contribution from the model prediction (i.e., the error rate e_i) while we took into consideration the contribution from the power function. In this way, we obtained a much better match after the switch and a good match before the switch at the same time.⁵ See Figure 6. Note that in the figure, we used exactly the same parameter settings (for a, b, B , and α) as in Ling and Marinov (1994). These parameters might be further optimized, which led to a slightly better fit but the difference was not significant.

The match between our model and the human data was excellent as measured by the sum-squared errors. Compared with Ling and Marinov (1994), CLARION (with the power function) did better on two of the three subjects, using the same parameters for transformation as Ling and Marinov did. See Figure 7 for a comparison.

So, what conclusion can we draw concerning the relative merits of the two models? The next section attempts to answer this question in a more theoretically oriented way.

Connectionist vs. Symbolic Models

Revisiting the argument of whether connectionist or symbolic models are better models (see Introduction), what do the above simulations have to say about it? To put it simply, we believe that this issue is a red herring. Being able to simulate some limited

⁵By adding the power function, we were able to reduce the total difference between the jumps in our curves and the corresponding jumps in the subject data by half. This comparison suggested that the amount of benefit the human subjects got from their predictions (i.e., by lowering e_i) was, although significant, relatively small. Significant benefit was gained through the improvement of motor responses as represented by the power function.

Subj.1:	$t_1 = 150$	$t_2 = 350$	$B = 700$	$\alpha = 0.33$
Subj.2:	$t_1 = 150$	$t_2 = 350$	$B = 1600$	$\alpha = 0.33$
Subj.3:	$t_1 = 100$	$t_2 = 210$	$B = 800$	$\alpha = 0.19$

Figure 5: Parameters for matching Lewicki’s data with power functions added (as in Ling and Marinov 1994).

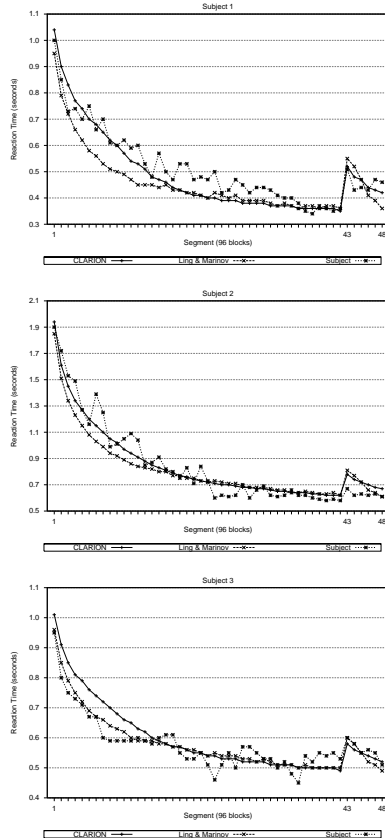


Figure 6: Matching Lewicki’s data with power functions added. See Figure 5 for parameter values.

data of implicit learning amounts to very little, in that any Turing equivalent computational process, that is, any generic computational model, should be able to simulate these data. Thus, the simulation of data by itself does not prove whether a particular model is a suitable one or not (cf. Roberts and Pashler 2000). Other considerations need to be brought in to justify a model.

We would suggest that one such issue is accessibility (Sun 1995, 1999, Cleeremans et al 1998). While symbolic models of implicit learning lead to explicit symbolic representation of implicit knowledge (e.g., Ling and Marinov 1994, Lebiere et al 1998, Anderson 1993) that is evidently accessible (without using any add-on auxiliary assumptions), connectionist models of implicit learning lead to implicit (subsymbolic, distributed) representation of knowledge that is in-

model	subj.1	subj.2	subj.3
CLARION w/o power f.	0.30	1.85	0.14
CLARION w/ power f.	0.07	0.25	0.05
Ling & Marinov (1994)	0.14	0.43	0.04

Figure 7: Comparing the goodness of fit in terms of SSEs.

herently inaccessible (such as in the bottom level of CLARION).

Note that it is generally not the case that distributed representations (as in the bottom level of CLARION) are absolutely inaccessible, but they are not as immediate as localist representations. Distributed representations may be accessed through indirect, transformational processes. As Kirsh (1990) put it, “explicitness [of representation] really concerns how quickly information can be accessed..... It has more to do with what is present in a process sense, than with what is present in a structural sense”. The accessibility difference between the two levels should be understood in this way.

Thus, connectionist models have a clear advantage: Being able to match human implicit learning data (at least) as well as symbolic models, they also account for the inaccessibility of implicit knowledge better and more naturally than symbolic models (Cleeremans et al 1998, Sun 1999). In this sense, they are better models.

On the other hand, it is generally agreed upon that symbolic/localist models have their roles to play too. They are better at capturing explicit processes, including their accessibility (Smolensky 1988, Sun 1995, 1999).

This contrast lends support to the belief that, since connectionist models are good for implicit processes and symbolic models for explicit processes, the combination of the two types should be adopted in modeling cognition (Sun 1995, 1999). There have been many philosophical and psychological theories related to this point (Sun et al 2001): See, e.g., James (1890), Schacter (1987), Reber (1989), Stanley et al (1989), Clark and Karmiloff-Smith (1993), and Sun (1999). This combination is exemplified by the CLARION model (Sun et al 2001).

This combination may also shed some light on the issue of consciousness, because the implicit/explicit difference involves, in its core, the issue of awareness, which is the key to consciousness (Cleeremans et al 1998). The representational distinction provides a plausible grounding for the notion of awareness (see Sun 1999 for details of theoretical arguments).

Simulating Other Tasks

Beside modeling the data from Lewicki et al (1987), CLARION has also been applied to model a variety of other SRT experiments, including Curran and Keele (1993) and Willingham et al (1989). Notably, in

these tasks, due to sufficient repetitions of sequential patterns, explicit learning of these patterns was involved, although implicit learning was dominant. Therefore, the top level of CLARION was utilized. Together, the model demonstrates the interaction between implicit and explicit learning (Sun et al 2002).

Beside SRT simulations, CLARION can capture data from many other implicit learning tasks. These tasks include artificial grammar learning (Reber 1989) and process control (Stanley et al 1989) (see Sun et al 2002). In addition, CLARION has also simulated extremely complex skill learning tasks as well, such as the minefield navigation task (see Sun et al 2001). The generality of CLARION has been amply demonstrated, on top of its cognitive validity.

Acknowledgments

This work was supported in part by Office of Naval Research grant N00014-95-1-0440 and by Army Research Institute grant DASW01-00-K-0012.

References

- J. R. Anderson, (1993). *Rules of the Mind*. Lawrence Erlbaum Associates. Hillsdale, NJ.
- M. Christiansen, N. Chater, and M. Seidenberg, (eds.) (1999). Connectionist models of human language processing: progress and prospects, special issue of *Cognitive Science*, 23(4), 415-634.
- A. Clark and A. Karmiloff-Smith, (1993). The cognizer's innards: a psychological and philosophical perspective on the development of thought. *Mind and Language*. 8 (4), 487-519.
- A. Cleeremans, A. Destrebecqz and M. Boyer, (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, 2, 10, 406-416.
- A. Cleeremans and J. McClelland, (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*. 120. 235-253.
- T. Curran and S. Keele, (1993). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 19, 189-202.
- J. Fodor and Z. Pylyshyn, (1988). Connectionism and cognitive architecture: A critical analysis. In: Pinker and Mehler (eds.) *Connections and Symbols*, MIT Press, Cambridge, MA. 1988
- W. James, (1890). *The Principles of Psychology*. Dover, NY.
- D. Kirsh, (1990). When is information explicitly represented. In: P. Hanson, (ed.) *Information, Language, and Cognition*. University of British Columbia Press, Vancouver, Canada.
- C. Lebiere, D. Wallach, and N. Taatgen, (1998). Implicit and explicit learning in ACT-R. *Proc. of ECCM'98*, pp.183-189. Nottingham University Press.
- P. Lewicki, M. Czyzewska, and H. Hoffman, (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 13 (4), 523-530.
- C.X. Ling and M. Marinov, (1994). A symbolic model of the nonconscious acquisition of information. *Cognitive Science*, 18(4), 595-621.
- A. Reber, (1989). Implicit learning and tacit knowledge. *Journal of Exp Psychology: General*. 118 (3), 219-235.
- S. Roberts and H. Pashler, (2000). How persuasive is a good fit? *Psychological Review*, 107 (2), 358-367.
- D. Schacter, (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 501-518.
- P. Smolensky, (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1):1-74.
- W. Stanley, et al, (1989). Insight without awareness *Quarterly Journal of Experimental Psychology*. 41A (3), 553-577.
- R. Sun, (1995). Robust reasoning. *Artificial Intelligence*. 75, 2. 241-296.
- R. Sun, (1999). Accounting for the computational basis of consciousness: A connectionist approach. *Consciousness and Cognition*, Vol.8, 529-565.
- R. Sun and T. Peterson, (1998). Autonomous learning of sequential tasks: experiments and analyses. *IEEE Transactions on Neural Networks*, Vol.9, No.6, pp.1217-1234.
- R. Sun, E. Merrill, and T. Peterson, (2001). From implicit skill to explicit knowledge: a bottom-up model of skill learning. *Cognitive Science*, Vol.25, No.2, pp.203-244.
- R. Sun, P. Slusarz, and C. Terry, (2002). The interaction between the implicit and explicit processes: a dual process approach. Submitted.
- C. Watkins, (1989). *Learning with Delayed Rewards*. Ph.D Thesis, Cambridge University, UK.
- D. Willingham, M. Nissen, and P. Bullemer, (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 15, 1047-1060.