Rules vs. Statistics in Implicit Learning of Biconditional Grammars

Bert Timmermans* & Axel Cleeremans**
*Vrije Universiteit Brussel **Université Libre de Bruxelles

Abstract

A significant part of everyday learning occurs incidentally — a process typically described as implicit learning. A central issue in this domain and others, such as language acquisition, is the extent to which performance depends on the acquisition and deployment of abstract rules. Shanks and colleagues [22], [11] have suggested (1) that discrimination between grammatical and ungrammatical instances of a biconditional grammar *requires* the acquisition and use of abstract rules, and (2) that training conditions — in particular whether instructions orient participants to identify the relevant rules or not — strongly influence the extent to which such rules will be learned. In this paper, we show (1) that a Simple Recurrent Network can in fact, under some conditions, learn a biconditional grammar, (2) that training conditions indeed influence learning in simple auto-associators networks and (3) that such networks can likewise learn about biconditional grammars, albeit to a lesser extent than human participants. These findings suggest that mastering biconditional grammars does not require the acquisition of abstract rules to the extent implied by Shanks and colleagues, and that performance on such material may in fact be based, at least in part, on simple associative learning mechanisms.

1. Introduction

Over development and learning, we acquire a considerable amount of information incidentally. Natural language offers perhaps the most striking example of such incidental learning: Infants do not need to be explained grammar rules in order to be able to communicate effectively and are presumably unaware of the fact that they are learning something at all. Adult speakers likewise "know" whether expressions of their native language are grammatically correct but can seldom explain why. Such dissociations between performance and ability to verbalize the relevant knowledge are often described as being subtented by processes of *implicit learning* (IL). Thus, the notion of *"implicit learning"* (IL) usually designates cases in which a person learns about the structure of a fairly complex stimulus environment, without necessarily intending to do so, and in such a way that the resulting knowledge is difficult to express [1]. IL is the ability to learn without awareness, as opposed to explicit learning, which is strategy- and/or hypothesis-driven, and of which one tends to be consciously aware. A considerable body of empirical evidence now suggests that people can indeed acquire information about the underlying structure of ensembles of stimuli in an incidental manner [5]. For instance, in a typical artificial grammar learning situation (e.g., [16]), Ss are asked to memorize a set of meaningless letter strings generated based on a simple finite-state grammar that specifies legal transitions between successive letters. Reber's main finding, now replicated hundreds of times, is that Ss are subsequently able to discriminate novel instances of grammatical strings from ungrammatical strings somewhat better than chance, despite remaining unable to verbalize the rules of the grammar. Based on these and other findings, Reber accordingly suggested that Ss must have unconsciously acquired abstract knowledge about the grammar. This early *abstractionist* account, however, has now become largely obsolete, based on (1) the fact that successful performance in this sort of task can be achieved without knowledge of the rule system (e.g., [2]), and on (2) the fact that when probed directly about the relevant knowledge, Ss often turn out to be able to express this knowledge [9], [23].

1.1 Implicit Learning and Abstraction

A central issue in this context is the question of whether the mechanisms through which implicit and explicit knowledge are acquired are best viewed as being subtended by separate processing systems or as being different manifestations of a single set of learning mechanisms. Early theories of IL (e.g. [16]) have tended to assume that it involves independent rule-based unconscious learning mechanisms. Today, based on issues raised by the complex measurement challenges associated with the assessment of awareness, as well as on the fact that many computational mechanisms can in fact perform in a *rule-like* manner without necessarily having acquired *rule-based* knowledge [17], many authors have proposed instead that performance in typical IL tasks is in fact best accounted for by assuming that Ss consciously learn either specific exemplars or fragments thereof during training. Performance at test can then be explained by simple mechanisms that compute the similarity between training and test exemplars, or that are sensitive to the overlap between fragments of the training items and the test items. From this perspective, the main distinction between implicit and explicit learning should thus not be one of awareness, but one of information-processing: Implicit learning would essentially involve incidental or episodic memory-based processes and result in conscious knowledge of exemplars and or fragments, whereas explicit learning would essentially involve active hypothesis testing and result in conscious knowledge of abstract rules. This position has been expressed most clearly by Shanks and colleagues [23], [22], [24]. However, while it is undeniable that humans are capable of abstract thought, the extent to which such processes are rooted in dedicated mechanisms remains unclear. Indeed, the debate about the nature of knowledge acquired in implicit learning situations finds an echo in recent research dedicated to various aspects of natural language learning, which we briefly discuss in the next section.

1.2 Implicit Learning and Natural Language

In a series of experiments modeled after the artificial grammar learning paradigm, Saffran et al. [18] exposed 6-7 years old children and adult Ss to a continuous speech flow such as *bupadapatubitutibudutabapidabu*. Ss were told that the experiment was about the influence of auditory stimuli on creativity. The only cues to word boundaries were the transitional probabilities between pairs of syllables (e.g., *bu-pa*), which were higher within words than between words. Afterwards, Ss heard two sets of sounds, each consisting of three syllable pairs, and were told to decide which one sounded more like the tape they had heard. Both adult and child Ss managed to perform well above chance, suggesting that learning about the deep structure of the material might proceed in the absence of intention to do so, and after only short exposure to the relevant material. Saffran et al. concluded that sensitivity to statistical structure is a fundamental process in language acquisition.

Marcus et al. [13], in stark contrast, claim that sensitivity to statistical structure is not sufficient to account for their data, and that 7-month-old infants can "represent, extract, and generalise abstract algebraic rules." The infants were exposed to artificial auditory "sentences" during a training phase, and were subsequently presented with test items instantiated with a novel set of sounds, half of which shared their abstract structure with the training items and half of which did not. For instance, infants habituated to *gatiti* or *linana* (both sharing an *ABB* structure) were subsequently presented with test sentences such as *wofefe* (familiar *ABB* structure) or *wofewo* (novel *ABA* structure). Despite the test material being instantiated over completely novel features, infants tended to listen more to the sentences instantiating a novel abstract structure. Marcus et al. concluded that infants had the capacity to represent "algebraic" rules, and that simple associative learning models such as connectionist networks would be unable to generalize as infants do. However, Marcus et al.'s claim that networks could not model the observed effect was disputed by several authors (e.g., [19], [15], [8]), essentially based on the fact that successful transfer need not necessarily be based on the overlap between features of the input patterns themselves. Instead "the relevant overlap of representations required for generalisation […] can arise over internal representations that are subject to learning." ([15], p.2). Transfer and generalization therefore remain complex issues, in part because of the challenges associated with designing stimulus material that can *only* be learned through abstractive mechanisms. Shanks and colleagues [22], [11] have attempted to address precisely this issue in an interesting series of experiments described in the following section.

1.3 Biconditional AGL: Shanks et al. (1997)

As mentioned before, Shanks and St John [23] proposed to abandon the idea of a conscious/unconscious dichotomy in favour of a rule-based/instance-based dichotomy. The basic idea is that humans possess two learning systems capable of creating distinct forms of mental representation, one system consisting of symbolic rule-abstraction mechanisms and the other involving subsymbolic, memory-based, connectionist mechanisms (see [21] for a discussion). In this context, Shanks et al. considered transfer in AGL tasks to be at least to some extent mediated by abstract (rule-) knowledge and claimed that people systematically become aware of the relevant regularities in *those AGL tasks where only rule learning is possible*. To demonstrate, Shanks et al. exposed Ss to artificial grammar strings generated by a biconditional grammar (see also [14]). Biconditional grammars involve cross-dependency recursion (see [3]) such that letters that appear at each position before and after a central dot depend on each other. An example is given in Figure 1, where letter D is paired with F, G with L, and so on. Shanks et al. constructed biconditional grammar training strings as well as a set of grammatical and ungrammatical and test strings, in such a way that grammatical and ungrammatical test items could not be

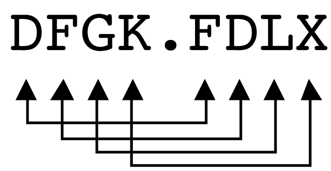

Figure 1: A biconditional grammar string as used by Shanks et al. (1997). Possible letters in each position before the dot are linked biconditionally with the letters that may appear after the dot.

distinguished — in contrast with the typical transitional grammars used in artificial grammar learning experiments — on the basis of their overlap with the training strings in terms of bigrams or trigrams (or any other *n*-gram). During training, two groups of Ss were shown strings one at a time on a computer screen and had to perform one of two tasks on each trial.

The *match* group Ss, who had been told that the task was about memory, were exclusively exposed to *grammatical* strings. On each trial, they first had to memorise a string displayed on the computer screen for a few seconds. Immediately thereafter, they had to identify this string among three possibilites (the string they had just memorized and two foils). The *edit* group Ss, in contrast, were exclusively exposed to *ungrammatical* strings. They were told that the strings had been constructed according to rules and that their task was to find them. On each trial, edit Ss were shown an *ungrammatical* string, and they had to indicate which letters they thought violated *(N)* or confirmed *(Y)* the rules. They were then given the correct string *and* the correct Y/N sequence as feedback. Shanks et al. showed a dissociation between the two groups: While the edit group performed well and most Ss extracted the rules, the match group performed at chance level, thus suggesting that "instance-memorisation and hypothesis-testing instructions recruit partially separate learning processes." ([22], p.243). Their basic claim is thus that discriminating between grammatical and ungrammatical biconditional strings *requires* abstract knowledge of the rule system, and that such knowledge *cannot* be learned by associative learning mechanisms such as instantiated in connectionist networks. In this paper, our goal is to explore the extent to which such networks can learn about biconditional grammars. To do so, we report on two simulation studies. Our first simulation study suggests that biconditional grammar learning can, under some conditions, be performed by networks developing representations based on frequency statistics. Our second simulation study was dedicated to exploring how differences between *match* and *edit* learning could be modelled without explicitly invoking a memory- versus rule-based distinction.

2. Simulation Study 1

In this first simulation study (see also [25]), we simply aimed to determine whether the Simple Recurrent Network (SRN; see Figure 2) was able to learn material from the Shanks et al. [22] experiments. The SRN, initially proposed by Elman (e.g. [10]; see also [6]) is one of the most influential connectionist models in the implicit learning and psycholinguistic literatures. SRNs are typically trained to predict the next element of sequences presented one element at a time to the network, and are therefore particularly appropriate to explore tasks involving sensitivity to sequential structure. To perform this prediction task, the network is presented, on each time

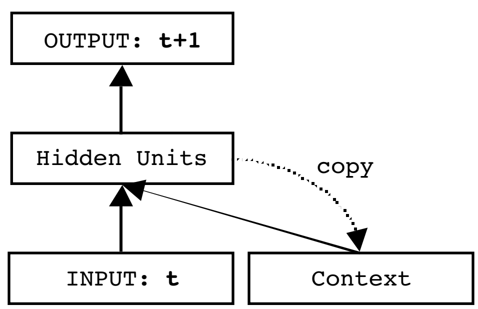

Figure 2: The Simple Recurrent Network as conceptualised by Elman (1990).

step, with element t of a sequence, and with a copy of its own internal state (i.e. the vector of hidden units activations) at time step t-1. Based on these inputs, the network has to predict element t+1 of the sequence. During training, the network's prediction responses are compared to the actual successor of the sequence, and the resulting error signal is then used to modify its connection weights using the back-propagation algorithm. As described in [20] and [4], the network progressively learns to base its predictions on the constraints set by an increasingly large and self-developed temporal window. This progressive incorporation of he statistical dependencies between successive elements of the sequence in the internal representations of the network eventually enables it to behave *as though* it had learned the relevant sequential rules. The SRN can thus, for instance, exhibit *perfect* generalization to an infinite number of novel sequences after (necessarily finite) training on a set of sequences generated from a finite-state automaton.

2.1 Network Architecture and Parameters

An SRN with 100 hidden units and local representations on its pools of input and output units was trained using backpropagation on the biconditional strings designed by Shanks et al. [22]. Strings were presented one element at a time to the network by activating the corresponding input unit (each of the 9 input units represented the letters D, F, G, L, K, X, the dot, the beginning, and the end of a string respectively) The learning rate was set to 0.15, and momentum was to 0.9. Context units were reset to zero after each complete string presentation.

2.2 Training Material

The training material consisted of the set of 18 strings designed by Shanks et al. [22] (List 1). The test material consisted of 18 novel grammatical and 18 ungrammatical strings respecting the following constraints: (1) Grammatical strings had to conform to the biconditional grammar: Letter position 1 is linked to 5, 2 to 6 and so on, with the linked letters being D–F, G–L, and K–X. (2) The use of the 6 letters was balanced, so that each letter appeared 3 times in each of the 8 letter locations. (3) Each training string differed from all other training strings by at least 4 letter locations. (4) Each training item had a grammatical similar item and an ungrammatical similar item that each differed from the training item by only 2 letter positions. Each training item was different from all other test items by at least 3 letter locations. The simulation was carried out on exactly these strings. A training epoch consisted of all 18 strings being presented once to the network, in a random fashion.

2.3 Procedure

Each of 9 networks initialized with different random weights was trained on the 18 training strings designed by Shanks et al. [22] for 3000 epochs. The networks were tested on seven different occasions during training. On each test, the networks were exposed to 18 novel grammatical strings and on 18 ungrammatical strings. Performance during test was assessed by recording the relative strength of the output unit corresponding to the actual successor of each element of each string. Different measurements of accuracy exist, of which we used the *Luce ratio* [12] — a simple measure of relative strength in which the activation of the target output unit is divided by the sum of the activations of all output units. These prediction responses were then averaged separately for each string so as to obtain a single measure of how well the networks were able to process each string. A high average luce ratio thus indicates that the network is successful in predicting each element of the corresponding string. Global measures of performance for each of the seven tests were obtained by averaging the mean luce ratios separately for grammatical and ungrammatical strings over the 9 networks.

2.4 Results

Figure 3 represents global prediction performance obtained during each of the 7 tests, and separately for training, novel grammatical, and ungrammatical strings. The figure clearly shows that the networks were able to discriminate between novel grammatical and ungrammatical strings. The training strings were learned almost perfectly from 100 epochs onwards. Further, the network clearly discriminates

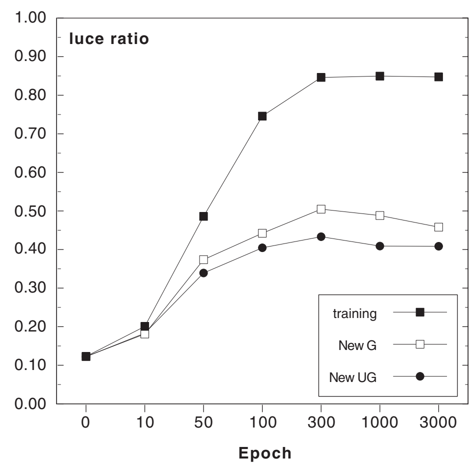

Figure 3: Average SRN prediction response strength, represented at various points during training, and plotted separately for training, novel grammatical and non-grammatical strings.

between novel grammatical and ungrammatical strings (i.e., better predictions for grammatical strings), even *before* it is completely successful in mastering the training strings. A MANOVA applied on these data confirmed that the networks successfully discriminated between novel grammatical and non-grammatical strings, $F(1, 8) = 97.08$, $p<.0001$. Further analyses aimed at ruling out that the networks had merely learned to predict the central dot or the end of the strings confirmed that letter-by-letter predictions were indeed better for grammatical than for ungrammatical strings, particularly for letters occurring after the central dot. Based on these findings, we can therefore conclude that contrary to what Shanks et al. claimed, the SRN can in fact distinguish between novel grammatical and ungrammatical strings generated by a biconditional grammar without making use of explicit rules.

It is important to note, however, that this result depends on the specific set of training strings used by Shanks et al. [22]. Indeed, the SRN exhibits well known specific difficulties in learning material that involves maintaining information across several times steps [20], as is the case here, and it would have failed had the stimulus material be perfectly balanced in terms of how frequently the different biconditional pairs occur in the stimulus set.

3. Simulation Study 2: The Match/Edit Distinction

In this second simulation study, our goal was to explore the effects of different training conditions on network performance. Recall that in [11], match Ss were given incidental learning instructions and were only exposed to grammatical strings. In contrast, edit Ss were informed that the strings instantiated a simple rule system and that their task consisted of uncovering this structure. Edit Ss were shown only ungrammatical strings, and had to indicate, on a letter-by-letter basis, which of the letters of each string they thought violated the rules. To do so, they typed a string of Y/Ns, endorsing or rejecting each letter of the string as grammatical. They were then shown the correct string, as well as the correct string of Y/N judgments.

3.1 Network Architecture and Parameters

To capture the match/edit distinction, we designed two simple feedforward networks. The *Match* networks were simple autoassociators that were trained exclusively on grammatical strings. The *Edit* networks in contrast, were exposed exclusively to ungrammatical strings during training, and were trained to produce both the correct string and the Y/N sequence as output, just as human participants.

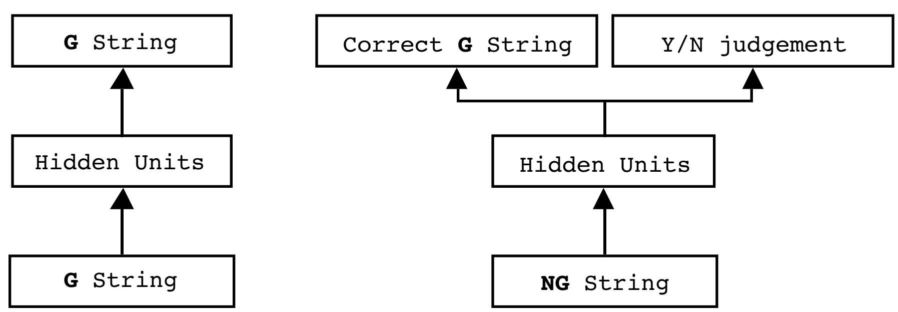

Figure 4: Match (left panel) and Edit (right panel) networks used in Simulation Study 2.

The networks are shown in Figure 4. Both networks used local representations on their pools of input and output units. Strings were presented by activating one of 6 units in each of 8 pools of units, each corresponding to the 6 letters that could occur in each of 8 positions within a string (the central dot was not represented). Edit networks were endowed with an additional pool of 8 units corresponding to the judgements about the grammaticality of each letter.

3.2 Procedure

A total of 9 networks in each condition were trained and tested in the same manner as described for Simulation Study 1. However, to assess performance in a way that more closely corresponds to human performance, we followed the procedure used by [7] so as to obtain percentages of correct classifications based on the networks' responses. To do so, we first computed average Luce Ratios for each test string, as described before (the activation of the nodes representing the Y/N input was not taken into account). Next, we computed the probability that each string would be classified as grammatical by entering its Luce Ratio in the following expression:

(1) $$p(\text{"grammatical"}) = 1/\,1 + e^{-k\,luce\,-\,T}$$

where *k* is a scaling parameter, *luce* is the average luce ratio for the string, and *T* is a threshold that was adjusted manually so as to yield equal numbers of "grammatical" and "ungrammatical" responses. The resulting individual probabilities were then averaged separately over grammatical and ungrammatical strings for each of the set of networks trained under match or edit conditions to yield global endorsement rates broken down by string type. Finally, based on these global endorsement rates, we computed the percentages of correct classifications expected for grammatical and ungrammatical strings in each condition.

3.3 Results

Results are shown in Figure 5. Following Shanks's analyses, Edit networks were classified as 'learners' and 'nonlearners' on the basis of the % correct responses at 1000, 2000 and 3000 epochs. The left panel shows the percentage of novel grammatical and ungrammatical strings that were endorsed by the networks as grammatical. The figure clearly shows that the Match networks fail to discriminate between G and NG strings, endorsing about 57% of each as grammatical. Edit Nonlearner networks perform better for most of the training period, but eventually likewise end up failing to discriminate between G and NG strings. In contrast, Edit Learner networks very quickly discriminate between G and NG strings, eventually endorsing about 57% of grammatical strings as grammatical, and correctly rejecting about 52% of the ungrammatical strings. The right panel of Figure 4 shows these data in a more compact form, representing the percentage of correct classifications produced by each type of network. Edit learner networks manage to achieve 57% of correct classifications overall. This result is well in line with standard results in the artificial grammar learning literature, but falls far short of the 95% correct classifications reported by [11]. Further manipulations of the simulation parameters and architecture will explore the extent to which this significant discrepancy can be reduced, but at this point, one can nevertheless conclude the following: First, the simulations were successful in showing that training the networks under "match" conditions indeed results in their failing to learn the biconditional grammar. Mere exposure to grammatical instances of biconditional strings does not seem to be sufficient for the auto-associator networks to become sensitive to the structure of the grammar. Second, we observed, like [11], that some Edit networks fail to learn while others succeed. Third, the simulations again suggest, consistently with Simulation Study 1, that biconditional grammars can be learned to some extent

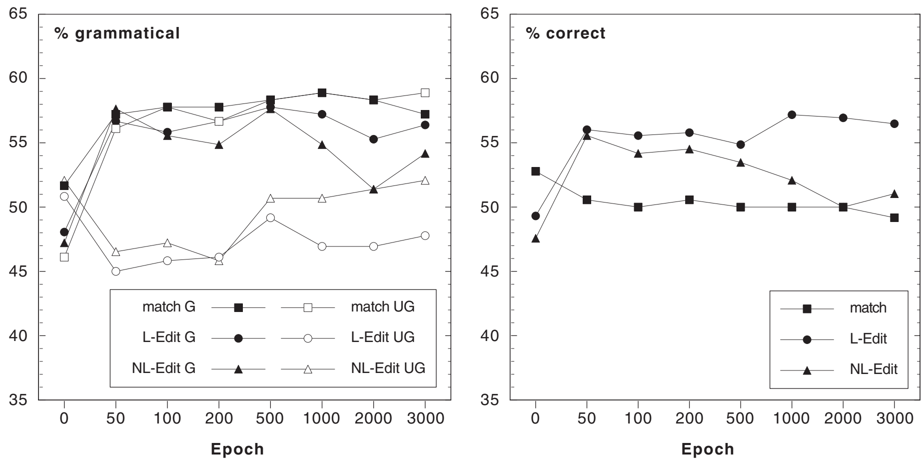

Figure 5: Left panel: Classification performance, plotted separately for grammatical (filled symbols) and ungrammatical (open symbols) strings, and for match (diamonds), non-learning edit (triangles) and learning edit (circles) networks. Right panel: Percentage of strings classified correctly, plotted separately for match (diamonds), non-learning edit (triangles) and learning edit (circles) networks.

through purely associative learning mechanisms. We discuss the implications of these findings in the general discussion that follows.

4. General Discussion

The goal of this paper was to explore the extent to which simple networks can learn about biconditional grammars. These grammars, in contrast to typical finite-state grammars, cannot be learned based on surface similarity, to the extent that neither memorized instances or fragments of the training strings contain cues about the grammatical status of a test item. Our main finding is that simple networks such as the SRN or some of the auto-associators networks used in Simulation Study 2 can actually learn to discriminate between novel grammatical and ungrammatical instances of biconditonal grammar strings. This outcome does not entail that rule-based learning never occurs (as it obviously does for some Ss in Shanks et al.'s experiments), but simply (1) that biconditional grammars might not address all the issues involved in efforts to dissociate rule-based vs. memory-based learning processes in the implicit learning literature, and (2) that abstraction might, at least on the larger portion of a representational continuum extending from pure instance-based representations to fully abstract, propositional representations, be a *graded* dimension. In this respect, connectionist models are particularly striking examples of the graded character of abstraction, to the extent that their internal representations can span most of the underlying continuum depending on task demands. Hence, while genuine abstraction may ultimately involve dedicated mechanisms closely tied to awareness and language, we believe that simple learning mechanisms based on functional similarity are often surprisingly powerful in the critical steps of developing ensembles of relevant sub-symbolic representations upon which further processes can then operate.

Acknowledgments

Axel Cleeremans is a Research Associate of the National Fund for Scientific Research (Belgium). This work was supported by a grant from the Université Libre de Bruxelles in support of IUAP program #P/4-19 and by European Commission Grant HPRN-CT-1999-00065. Bert Timmermans is supported by a grant from the Vrije Universiteit Brussel.

References

1. Berry, D.C. and Dienes, Z., (1993). *Implicit learning: Theoretical and empirical issues*. Hillsdale: Lawrence Erlbaum Associates.
2. Brooks, L.R. (1978). *Non-analytic concept formation and memory for instances* In *Cognition and Concepts*, (E. Rosch and B. Lloyd, eds). Hillsdale: Lawrence Erlbaum Associates, pp. 16-211.
3. Christiansen, M. and Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, *23*, 157-205.
4. Cleeremans, A., (1993). *Mechanisms of implicit learning: Connectionist models of sequence processing*. Cambridge: MIT Press.
5. Cleeremans, A., Destrebecqz, A., and Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*, 406-416.
6. Cleeremans, A., Servan-Schreiber, D., and McClelland, J.L. (1989). Finite state automata and simple recurrent networks. *Neural Computation*, *1*, 372-381.
7. Dienes, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science*, *16*, 41-79.
8. Dienes, Z., Altmann, G., and Gao, S.-J. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science*, *23*, 53-82.
9. Dulany, D.E., Carlson, R.A., and Dewey, G.I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General*, *113*, 541-555.
10. Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-212.
11. Johnstone, T. and Shanks, D.R. (submitted). Abstractionist and processing accounts of implicit learning. .
12. Luce, R.D. (1963). Detection and Recognition. In *Handbook of Mathematical Psychology*, (R.D. Luce, R.R. Bush and E. Galanter, eds). Wiley: New York.

13. Marcus, G.F., Vijayan, S., Bandi Rao, S., and Vishton, P.M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77-80.
14. Mathews, R.C., Buss, R.R., Stanley, W.B., Blanchard-Fields, F., Cho, J.R., and Druhan, B. (1989). Role of implicit and explicit process in learning from examples: A synergistic effet. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *15*, 1083-1100.
15. McClelland, J.L. and Plaut, D. (1999). Does generalization in infant learning implicit abstract algrebra-like rules? *Trends in Cognitive Sciences*, *3*, 166-168.
16. Reber, A.S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 855-863.
17. Redington, M. and Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. *Journal of Experimental Psychology: General*, *125*, 123-138.
18. Saffran, J.R., Newport, E.L., Aslin, R.N., and Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological Science*, *8*, 101-105.
19. Seidenberg, M.S. and Elman, J.L. (1999). Do infants learn grammar with algebra or statistics? *Science*, *14*, 433.
20. Servan-Schreiber, D., Cleeremans, A., and McClelland, J.L. (1991). Graded State Machines: The representation of temporal contingencies in simple recurrent networks. *Machine Learning*, *7*, 161-193.
21. Shanks, D.R. (1998). Distributed representations and implicit knowledge. In *Knowledge, Concepts and Categories*, (K. Lamberts and D.R. Shanks, Eds). Hove: Psychology Press, pp. 197-214.
22. Shanks, D.R., Johnstone, T., and Staggs, L. (1997). Abstraction processes in artificial grammar learning. *The Quarterly Journal of Experimental Psychology*, *50*, 216-252.
23. Shanks, D.R. and St. John, M.F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367-447.
24. St John, M.F. and Shanks, D.R. (1997). *Implicit learning from an information-processing standpoint* In *How implicit is implicit learning?*, (D.C. Berry, Ed.). Oxford: Oxford University Press, pp. 124-161.
25. Timmermans, B. and Cleeremans, A. (2000). Rules versus statistics in biconditional grammar learning: A simulation based on Shanks el al. (1997) In *Proceedings of the Twenty-Second Annual Conference of Cognitive Science Society*. Hillsdale: Lawrence Erlbaum Associates, pp. 947-952.