CHAPTER 14

IMPLICIT LEARNING IN SECOND LANGUAGE ACQUISITION

John N. Williams

I. INTRODUCTION

We constantly use implicit knowledge in everyday action and perception. The father who tries to teach his child to ride a bicycle immediately realises that, despite riding to work every day, he cannot explain how to turn a corner. A tennis player's backhands might always land out despite their conscious efforts to aim them in. When we listen to music we will instantly recognise a note that violates the principles of musical structure to which we have become accustomed in our culture, even if we have never had any musical training. People have fluent and productive command of their native language and are able to instantly detect grammatical irregularities, without being able to explain the underlying rules. Implicit knowledge 'can be causally efficacious in the absence of awareness that this knowledge was acquired or that it is currently influencing processing' (Cleeremans, Destrebecqz, & Boyer, 1998, p. 406). Examples of the use of explicit knowledge are when a student selects a particular theorem to solve a geometry problem, when the learner driver follows their instructor's step-by-step commands in order to change gear or when the language learner consults a grammar book in order to find the first person singular form of a particular verb. Explicit knowledge is knowledge that we know that we know (Dienes & Perner, 1999) and that we are aware of using.

Within second language acquisition (SLA), the contrast between implicit 'acquisition' and explicit 'learning' was brought to the fore by Krashen (1981, 1994), but his Acquisition–Learning Hypothesis is compatible with any theory of the putative implicit learning mechanism. At one extreme, generative linguists would appeal to processes that operate with reference to universal grammar (UG). At the other, emergentists would appeal to domain-general principles of associative learning, as exemplified perhaps by connectionism. All studies of 'acquisition' are studies of implicit learning.

The New Handbook of Second Language Acquisition

What appears to define ‘implicit learning research’ is a type of methodology, rather than a theoretical orientation. It involves control over the learning task, control over the input, measurement of learning and, in the best cases, rigorous attempts to establish whether test performance is a reflection of a properly operationalised concept of implicit knowledge. But this is simply the methodology one needs in order to unambiguously establish acquisition, in Krashen’s sense, of anything, whatever one’s theoretical orientation. Thus, although researchers with a more emergentist view of learning are naturally drawn to implicit learning research, this is a field within which it is possible to rigorously explore all learning processes. Indeed, evidence of limitations on implicit learning can, in principle, provide a firmer basis for an appeal to innate constraints on learning than the traditional and, for some, questionable (e.g. Elman et al., 1996) theoretical arguments from learnability theory.

In this overview we will consider a number of theoretical and methodological issues. First, how are implicit and explicit knowledge to be operationalised so that they can be measured, and what is the evidence for implicit second language knowledge according to these criteria? We then consider the learning process, what can be learned and what cannot, and what this might tell us about the nature of the implicit learning mechanism. This raises the issue of constrains on implicit learning, which we consider further in the context of the influence of attention.

But first some terminology. As noted by Hulstijn (2003), it is important to maintain a distinction between incidental learning and implicit learning. In its strictly methodological sense, *incidental learning* refers to an experimental arrangement in which the participants are not informed that there will be a test of learning. This is also true of implicit learning experiments. Within SLA research, the term *incidental* has also come to be used in relation to the actual learning process to mean that people learn something without intending to. For example, they might learn a rule of grammar in the course of performing a meaning-focused task, or they might learn some regularity in the sequencing of forms whilst performing a short-term memory task. The term *implicit learning* refers to the above situations, with the added condition that there is no awareness of the regularity to be learned at the point of learning. In contrast, *explicit learning* involves an intention to learn (which may or may not result from advance warning of a test of learning) as well as the use of conscious knowledge at the point of learning. For example, the learner might engage in hypothesis formation and testing in an attempt to discover underlying structure.

The learning process itself can be characterised as either inductive or deductive. *Inductive* learning involves forming generalisations on the basis of examples, whereas *deductive learning* is guided or constrained by additional knowledge (e.g. parametric options provided by UG). Implicit learning is usually regarded as inductive, but if UG is involved, it could be regarded as deductive (DeKeyser, 2003).

It is useful to distinguish the nature of the learning process from the status of the resulting knowledge as assessed by a test of learning. Implicit learning could lead to explicit knowledge, since a person may become spontaneously aware of regularities in the input. To borrow a term from the problem-solving literature, this might be referred to as ‘insight’. At the same time, explicit learning might result in implicit knowledge. With increased practice, explicit knowledge may become automatised and may come to influence behaviour without awareness. Thus, the issue of the existence of implicit or explicit knowledge in the mind of the learner is distinct from the issue of how it got there. We start with a discussion of how implicit and explicit knowledge may be measured, regardless of how it was acquired.

II. IMPLICIT AND EXPLICIT KNOWLEDGE

As Ellis (2005) remarks, there is a 'data problem' in SLA research. Although there are competing theories of the acquisition process, it is difficult to adjudicate between them because of the difficulty of accurately measuring acquisition, as opposed to learning. 'Thus, SLA as a field of inquiry has been characterized by both theoretical controversy and by a data problem concerning how to obtain reliable and valid evidence of learners' linguistic knowledge' (Ellis, 2005, p. 142). What criteria might be used to determine whether a person's behaviour is determined by implicit knowledge?

A. Influences Behaviour Without Awareness

Awareness is the most commonly used criterion of implicitness within psychology. Explicit knowledge is knowledge that a person knows that they know (Dienes & Perner, 1999). If we characterise a first-order state as simply having knowledge of something, a person can be said to have explicit knowledge when they are in a higher order state of knowing that they know something. They should be able to intentionally use this higher order knowledge to control actions, including verbal report. Conversely, implicit knowledge is defined as knowledge that a person has without knowing that they have it.

Two commonly used implicit learning paradigms in psychology make use of verbalisation as an operationalisation of implicit knowledge. The artificial grammar (AG) learning paradigm was introduced by Reber (1967). The learning materials consist of letter sequences such as VXXVS and TPPPTS that are generated by a finite state grammar.[1] Participants are typically exposed to these sequences in the context of what appears to be a short-term memory test. They are then told that the letter sequences, in fact, followed a rule system, and they perform a grammaticality judgement test (GJT) on new grammatical and ungrammatical letter strings. Their performance is above chance, yet they are completely unable to verbally describe the underlying system (Reber, 1967; Reber & Allen, 1978). Reber and Allen (1978) conclude that it is possible to implicitly acquire an abstract representation of the structure of the grammar. Whilst some disagree that the knowledge acquired can be properly described as abstract, or that it is wholly unconscious (see below), there is no doubt that such experiments demonstrate incidental acquisition of statistical properties of stimulus sequences.

In serial reaction time (SRT) tasks, a stimulus moves between different screen positions (typically 4 or 6) and the participant indicates each position using corresponding response keys. What the participant is not told is that the majority of the sequences

[1]Strings are produced by tracing a path through a state diagram. Starting from an initial state, each transition to a subsequent state generates a specific letter. Some states can lead to a variety of alternative states. Strings are grammatical if a path can be traced through the diagram from the start state to the end state. The grammar generates a finite set of strings.

follow a regular pattern, generated by either a finite state grammar (e.g. Cleeremans & McClelland, 1991) or, more simply, a repeating sequence of positions (Curran & Keele, 1993; Destrebecqz & Cleeremans, 2001, 2003). With practice, responses to stimuli in the regular sequences gradually get faster, but if the stimuli suddenly appear in random sequences, responses slow down markedly, indicating sensitivity to the structure of the regular sequences. Yet when asked afterwards if they noticed any pattern to the sequences, participants can provide only minimal valid information (Cleeremans & McClelland, 1991; Norman, Price, Duff, & Mentzoni, 2007). In Cleeremans and McClelland (1991), they even felt that explicit knowledge was detrimental to their performance and so avoided using it. Subjects may also fail to distinguish fragments of the trained sequence from novel fragments in a recognition memory test (Destrebecqz & Cleeremans, 2003; Norman et al., 2007). Thus, it appears that the slowdown for random sequences is due to the violation of expectancies based on implicit knowledge.

Green and Hecht (1992) showed a striking dissociation between verbal report and a sentence correction task within naturalistic SLA. They found that the ability to correct grammatical errors lagged well behind the ability to provide explanations for the corrections and that correct corrections were often associated with incorrect explanations. They argue that the ability to correct sentences is driven by implicit knowledge and that whilst explicit knowledge provided through instruction might facilitate the development of implicit knowledge, learners rely on the latter in the correction task.

The above studies appear to demonstrate sensitivity to regularities in the absence of verbalisable knowledge. But do they demonstrate the existence of implicit knowledge? According to critics such as Shanks and St. John (1994) the evidence is not compelling. AG and SRT experiments tend to use regularities that are intrinsically difficult to verbalise; there is a delay between training and debriefing and a lack of detailed questioning. Therefore, verbal reports are not a reliable indicator of awareness.

An alternative to verbal report is to require participants to make subjective judgements of their mental state when making each decision. For example in a GJT, learners might be asked to rate their confidence in each judgement that they make. If the accuracy of their decisions is above chance when they say they are guessing, then they can be said to be using implicit knowledge. In an AG experiment, Dienes and Scott (2005) found exactly this, providing compelling evidence of implicit learning of letter sequences. But they also found that the average confidence level of correct decisions was significantly higher than that of incorrect ones, suggesting that people were basing some of their decisions on conscious knowledge.

But do correct high-confidence judgements necessarily imply explicit knowledge? Not necessarily. Consider making grammaticality judgements in one's native language. One may well be highly confident, even though the judgements might be based on implicit knowledge. In such cases, judgements appear to be a reflection of *intuition*. Dienes and Scott (2005) therefore argue that we must separate out the part of the mental state concerned with whether the judgement is based on a conscious intention (i.e. confidence) from the part concerned with conscious knowledge of the structure of the domain (i.e. structural knowledge). Intuition would be when one has conscious judgement knowledge (not guessing), but no conscious structural knowledge. Norman et al. (2007) refer to this state as 'fringe consciousness' and define it as 'A situation in which behaviour is driven in a flexible manner by consciously accessible feelings, but where there is no conscious access to the antecedents of those feelings' (p. 833).

How can we assess this state of intuition? In Dienes and Scott's (2005) AG learning experiment, in addition to making confidence judgements, subjects were also asked to say whether each judgement was based on a guess, intuition, memory (for items received in training) or rule. Judgements based on memory and rule were above chance in accuracy, reflecting explicit knowledge; so too were judgements based on guess and intuition, suggesting a contribution of implicit knowledge. Rebuschat (2008) also used a GJT supplemented with confidence and source judgements in a study of learning German verb position rules under different training conditions. Under incidental training conditions involving a focus on meaning, there was a correlation between confidence and accuracy, and responses based on memory and rule were significantly above chance, indicating a contribution of explicit knowledge. But whilst guess responses were at chance, moderately confident responses based on intuition were significantly above chance, indicating a contribution of unconscious structural knowledge. Interestingly, under training conditions that required participants to intentionally search for rules, there was rather stronger evidence for unconscious knowledge since above-chance responding was even found for guess responses. Thus, even intentional induction can lead to implicit knowledge.

We can draw two conclusions from this work on the assessment of subjective mental states. First, no test of knowledge is likely to be process pure. Grammaticality judgements will reflect contributions of both implicit and explicit knowledge, perhaps depending on the specific item involved. For example, in an AG experiment, ungrammatical items that contain violations in the salient beginning and end positions might lead to high-confidence judgements based on memory or rule, whereas violations in the middle part of the string might lead to moderately confident responses based on intuition, or even guesses. In SLA studies, combining such measures with a linguistic analysis of test structures could provide valuable information about the kinds of regularities that are more or less likely to be associated with explicit and implicit knowledge.

Second, limiting our interest to situations in which knowledge is applied completely unconsciously is perhaps too severe and unrealistic. Indeed, for some sceptics there are no such situations in any case, and all claims to the contrary are based on flawed methodology (Lovibond & Shanks, 2002; Shanks & St. John, 1994). Knowledge structures conscious perception, and learning involving cognitive representations will usually lead to changes in conscious experience of one kind or another (Perruchet & Gallego, 1997; Perruchet & Vinter, 1998). For example, in a GJT, grammatical items might be processed with greater perceptual fluency than ungrammatical ones, and awareness of this fact can bias towards judging them as grammatical (Buchner, 1994; Kinder & Shanks, 2003). Or learners may come to consciously perceive the input as segmented into chunks, such as bigrams in AG experiments or words and phrases in natural language. The underlying learning mechanisms producing these effects may be implicit, but their effect is to structure conscious perception. Therefore, we should not be surprised that it is difficult to isolate cases where knowledge has absolutely no conscious effects and behaviour is a result of guessing. What is perhaps more important is that in the moment of use, knowledge influences behaviour in the absence of conscious or intentional recollection of previous experiences or explicit rules. This leads us towards measures of awareness that allow for confident responses based on intuition. But it also opens the way to other diagnostics of implicit knowledge, such as automaticity.

B. Influences Behaviour Automatically

Explicit knowledge guides intentional actions, whereas implicit knowledge is deployed automatically (Cleeremans & Jiménez, 2002; Dienes & Perner, 1999). Thus, automaticity can be used as a diagnostic of implicitness. Of course, how automaticity is to be identified is an issue in itself (Segalowitz, 2003). Within SLA a speed diagnostic is prevalent, presumably because it relates to the notion of fluency, and fluency is seen as a reflection of acquisition, as opposed to learning (consider Krashen's, 1981, Monitor Hypothesis).

Ellis (2005) examined the correlations between performance by learners of English on 17 target structures in five language tests: oral imitation, oral narrative, a timed GJT (responding before a deadline), an untimed GJT and a metalinguistic knowledge test. A principal components factor analysis showed that the two oral tests and the timed GJT loaded on one common factor, whilst the untimed GJT and metalinguistic tests loaded on a second.[2] In terms of Ellis's task analysis, the factor that distinguishes the oral and timed GJT is time pressure. Assuming that speeded performance primarily reflects implicit knowledge, these two sets of tasks distinguish implicit and explicit knowledge. Considering the very different tasks involved, it is impressive that such a clear division between them emerged. In a detailed by-item analysis of the timed and untimed GJTs, Ellis (2006) found no relationship between the level of performance for individual structures on the two tasks, reinforcing the idea that they tap different types of knowledge.

The Ellis (2005, 2006) studies clearly show a distinction between knowledge that can be applied quickly and knowledge that takes longer to access. But this need not correspond to any differentiation in the form of the underlying knowledge because the same (explicit) knowledge could just be used more quickly with practice. Another approach is to combine various aspects of automaticity within one test, such as speed and freedom from attentional control. Oral production and imitation tasks are often regarded as relatively good measures of implicit knowledge because they divert attention away from form whilst imposing time pressure. Ellis (2005) found that out of his battery of tasks, elicited imitation loaded most heavily on the 'implicit' factor. Erlam (2006) developed a version of an elicited imitation task that involved hearing a statement (that may or may not involve a grammatical error), judging its truth value and repeating it in *correct* English. Performance correlated moderately well with an oral narration task and fairly strongly with IELTS listening and speaking scores. However, it is unclear just how critical time pressure is in such tasks. Hulstijn and Hulstijn (1984) found that accuracy in a story-retelling task was not affected by time pressure, only by focus on form. As Erlam (2006, p. 487) notes, the most direct evidence for the use of implicit knowledge in elicited imitation would come from

[2]See Isemonger (2007) for a critique of the statistical procedures used by Ellis (2005), and Ellis and Loewen (2007) for an even clearer separation of the two factors when more appropriate procedures are employed.

spontaneous, fluent and unconscious corrections of ungrammatical sentences in the input. Unfortunately, her study did not contain a rigorous assessment of awareness, but the idea is reminiscent of the notion of 'fluent restoration' in speech shadowing (Marslen-Wilson & Welsh, 1978). Whilst shadowing speech in their native language, subjects will often spontaneously correct mispronounced words with no disfluency. The ability to fluently restore grammatical errors in shadowing tasks that also involve a meaning-based component would perhaps provide a more stringent test of implicit grammatical knowledge than imitation.

Neurological measures perhaps provide the most promising approach to the identification of automatic processing. Event-related potential (ERP) responses like the P600, N400 and especially the early left anterior negativity (ELAN) are produced within a few hundred milliseconds of semantic and syntactic violations and so are not likely to be the result of conscious thought processes. This is especially true of the ELAN, which is assumed to reflect immediate and automatic structure-building operations (for a review, see Friederici, 2002). Friederici, Steinhauer, and Pfeifer (2002) found the characteristic ERP signatures of syntactic processing in learners of Brocanto, an artificial language that was learned under intentional induction in the context of a board game. It is perhaps surprising to find such native-like processing in an artificial-language-learning experiment involving relatively little exposure when studies on naturalistic learners fail to find such effects, particularly with regard to the ELAN (Hahne & Friederici, 2001). Friederici et al. (2002) argue that their participants can simply be regarded as having achieved a very high level of proficiency in a very small language. Morgan-Short (2007) replicated these results, but also found that there were no ERP effects for participants who were given explicit instruction in the rules of Brocanto prior to playing the board game, even though final GJT performance was similarly high for the instructed and uninstructed groups. Thus, ERPs can reveal differences in underlying processing that are not reflected in behaviour. Also there appears to be a big difference between being told rules and working them out for oneself, with (intentional) induction being more likely to lead to native-like processing than instruction, even when the amount of practice is held constant. As in the case of Rebuschat (2008), we see that intentional induction can lead to implicit knowledge.

There is also evidence for native-like brain responses after very little exposure in classroom settings (Osterhout, McLaughlin, Pitkanen, Frenck-Mestre, & Molinaro, 2006), suggesting rapid assimilation of the L2 into the learners' comprehension system. Particularly impressive are cases that show dissociations between brain responses and more 'direct' behavioural tests. McLaughlin, Osterhout, and Kim (2004) found that after only 14 h of instruction, learners of French showed different ERP responses to words and nonwords, yet they were unable to distinguish them in a lexical decision task. Tokowicz and MacWhinney (2005) studied beginner learners of Spanish and found strong ERP responses to gender violations in online processing despite low sensitivity to gender violations in an offline task. As we will see below, we must be cautious in interpreting these task dissociations as evidence for different implicit and explicit knowledge systems. But these experiments do appear to show very rapid assimilation of some aspects of second languages, resulting in automatic and native-like brain responses.

C. Different Brain Systems

Another way of distinguishing implicit and explicit knowledge may be in terms of the brain regions that support them. Amnesics show dissociations between implicit and explicit memory in that they perform normally on 'indirect' tests of implicit memory, but relatively poorly on 'direct' tests of explicit memory (see Gabrieli, 1998, for a review). After being exposed to a list of words, they will show poor recognition memory, but intact priming (Haist, Musen, & Squire, 1991). Or in AG learning experiments, they will show normal levels of performance on GJT, but impaired recognition for the training items compared to controls (Knowlton & Squire, 1996). An obvious conclusion from such dissociations is that implicit and explicit memories are subserved by different brain regions (N. C. Ellis, 1994; Gabrieli, 1998; Squire, 1992).

But despite what appears to be compelling evidence for dissociations between implicit and explicit knowledge, it is still possible to defend a single system view in which performance on different tasks, such as recognition and priming, simply reflects the differential accessibility of the same knowledge. What determines conscious accessibility is not where the knowledge is stored in the brain, but its level of 'analysis' (Bialystok, 1982) or more generally its 'quality' as defined in terms of stability, strength and distinctiveness (Cleeremans & Jiménez, 2002). The further the knowledge has progressed along these dimensions, the more likely it is to become amenable to conscious control and to enter into the 'global workspace' where it becomes available to other cognitive systems (Dehaene & Naccache, 2001). It has also been shown through computational modelling that dissociations between tasks can even be obtained when the underlying knowledge is of the same level of analysis or quality (Kinder & Shanks, 2003). Retrieving a specific learning episode in, say, a recognition memory task is relatively difficult because it involves a fine discrimination between memory traces, whereas making an intuitive judgement about well-formedness is relatively easy because it can utilise information that is aggregated over all training items. We should bear in mind, therefore, that demonstrations that learners' brain responses show effects that are not evident in behavioural measures (McLaughlin et al., 2004; Tokowicz & MacWhinney, 2005) may reflect differential sensitivity of the tasks to the same underlying knowledge. The moral is that we should be very cautious in interpreting dissociations between tasks as evidence for dissociations between implicit and explicit knowledge systems.

Nevertheless, within SLA there is a strong preference for a multiple systems perspective. For example, Paradis (1994, 2004) distinguishes the kind of procedural knowledge acquired in learning a motor skill, or one's first language, from the kind of declarative knowledge acquired in a geography lesson or the metalinguistic knowledge acquired in a foreign language lesson. Obviously, given the radically different form of representation involved, we would expect these types of knowledge to be represented in different brain regions. Ullman (2001, 2004) also draws a distinction between different memory systems in his declarative–procedural (DP) model and argues specifically that the rule-governed aspects of language (across syntax and morphology) are supported by the procedural system (rooted in frontal/basal ganglia circuits) and item-based aspects are supported by the declarative system (rooted in medial and lateral temporal lobe structures).

From a single system perspective, the fact that different types of knowledge are represented in different brain regions is irrelevant to the issue of conscious availability. What is relevant is the level of analysis/quality of the knowledge (indeed, for Cleeremans & Jiménez, 2002, all knowledge is ultimately represented in the same subsymbolic form). In contrast, Paradis (1994, 2004) appears to equate declarative and procedural knowledge with explicit and implicit knowledge, respectively, such that any aspect of language that is known implicitly must be assumed to be represented in the procedural system. However, there is evidence that contradicts such a strong association between memory systems and conscious accessibility. For example, damage to the declarative system can impair certain forms of implicit learning (Chun & Phelps, 1999), and damage to the procedural system does not necessarily impair implicit AG learning (Reber & Squire, 1999; Witt, Nuhsman, & Deuschl, 2002). Such findings can be accommodated by the DP model since here the terms ‘procedural’ and ‘declarative’ are used primarily to refer to differing forms of knowledge (and associated brain systems). The model does not assume an isomorphic relation between declarative/procedural memory and explicit/implicit knowledge and assumes that declarative memory underlies implicit and explicit knowledge, while the procedural system is one of several brain systems underlying different types of implicit knowledge (Ullman, personal communication; see Ullman, 2005, for discussion). Just because some knowledge is known implicitly does not necessarily imply that it is represented in the procedural system, although if there is anatomical evidence that it is represented in the procedural system, the prediction would be that it should also bear the hallmarks of implicit knowledge.

D. Conclusion

Following the definition of implicit knowledge as knowledge that a person does not know that they know (Dienes & Perner, 1999), implicitness can be operationalised only through assessments of subjective mental states, that is, through measurements of awareness. When automaticity is used as a diagnostic, we must recognise that we are *assuming* that conscious knowledge could not have been used in the moment of generating the behaviour that we are measuring. The more converging lines of evidence there are, such as speed, independence from attention and native-like brain responses, the more convincing this assumption will be. An advantage of this approach is that we can accept that learners might have conscious knowledge, as assessed by unspeeded tasks and yet still be producing automatic behaviour using implicit knowledge.[3] Having said this, the most convincing evidence for implicit knowledge will always come from subjective measures.

To date, few language studies have attempted to establish that implicit knowledge was acquired according to any of the above criteria. The term *implicit learning* is often simply used to refer to a mode of learning that is incidental and inductive.

[3]Morgan-Short’s (2007) study provides an example. Given that GJT performance was around 80% correct, it seems likely that both instructed and uninstructed groups had conscious knowledge of the rules of Brocanto, yet only the uninstructed group showed evidence of implicit knowledge according to an automaticity criterion.

This is presumably a reflection of a concern with the effectiveness of particular modes of learning rather than the status of the resulting knowledge. So in the remainder of this review we shall be essentially concerned with the nature of incidental inductive learning, treating this as 'implicit learning', even if the implicitness of the resulting knowledge was not actually established. Ideally, though, the term *implicit learning* ought to be used to refer to situations where implicit knowledge was acquired, as established by the awareness criterion, or as assumed by virtue of automaticity.

III. THE NATURE OF WHAT CAN BE LEARNED IMPLICITLY

A. Chunking and Statistical Learning in Orthography, Phonology and Syntax

Research using the AG and SRT paradigms consistently shows evidence for chunk learning, that is, short sequences of letters or stimulus positions that frequently occur in the input. For example, the test item XXVXJ contains the bigrams XX, XV, VX and XJ and the trigrams XXV, XVX and VXJ. Participants may simply learn these bigrams and trigrams, perhaps also being sensitive to which occur at the beginnings and ends of strings. The more of these bigrams and trigrams a test item contains, the greater the likelihood that it would be classified as grammatical (Johnstone & Shanks, 1999; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990).

Chunking is just one instantiation of what, more recently, has come to be referred to as 'statistical learning', a strongly empiricist and emergentist approach that sees acquisition as the absorption of statistical regularities in the environment through implicit learning mechanisms. For example, it has been shown that when people are played what appears at first to be a random sequence of syllables such as *ba-bu-pu-du-ta-ba-bu-pa-da-du-ta-ba-tu-ti-bu-ba-bu-pu-tu-ti-bu*, they rapidly acquire a sense of recurring sequences that could be regarded as constituting lexical items, that is *babupu dutaba bupada dutaba tutibu babupu tutibu* (Saffran, Newport, & Aslin, 1996b). The dominant interpretation of this result is that people unconsciously 'tally' the transition probabilities (Aslin, Saffran, & Newport, 1998), or more precisely contingencies (N. C. Ellis, 2006a) between syllables. Because the predictability of syllables is higher within words than between them, dips in transitional probability signal lexical boundaries (but see Perruchet & Vinter, 1998, for an alternative account in terms of chunking effects in memory and perception). This effect has been demonstrated in adults after 20 min of exposure (Saffran et al., 1996b), in eight-month-old infants (Saffran, Newport, & Aslin, 1996a) and even in cotton-top tamarins (Hauser, Newport, & Aslin, 2001). Mirman, Magnuson, Graf Estes, and Dixon (2008) demonstrate that it actually feeds into vocabulary learning.

Implicit statistical learning effects have also been demonstrated in phonology and orthography. Dell, Reed, Adams, and Meyer (2000) investigated learning of artificial phonotactic constraints. For example, the fact that within the experimental materials [f] always occurred as an onset and [s] as a coda influenced the speech errors that were produced. The learning effect was assumed to be implicit because there was no difference between participants who were initially informed about the constraints and those who were not and because in any case, speech errors are produced unintentionally and automatically. Even more impressively, Dell et al. (2000) and Warker and Dell (2006) demonstrate implicit learning of second-order constraints; for example, if the vowel is [ae], [g] must occur as an onset and [k] as a coda, but if the vowel is [I], [k] must occur as an onset and [g] as a coda.

Following from research showing that statistical information can be used to break syllable streams into words, a number of experiments have explored whether it can help learners break streams of words into phrases. These experiments involve presenting meaningless strings of words that are generated by a phrase structure grammar and seeing if participants incidentally acquire sensitivity to the underlying phrasal groupings. One statistical cue to phrase structure is what is referred to as 'predictive dependencies'. For example, in English, an article requires a noun to be present, but a noun can occur with or without an article. Artificial languages that respect this structure are learned better than ones that do not (Saffran, 2001; Saffran et al., 2008). In addition, transition probabilities between words are high within phrases and low at phrasal boundaries, and this can be made more evident by including optional, repeated and moved phrases. Even though such features increase the complexity of the language, they do improve sensitivity to underlying phrasal structure (Morgan, Meier, & Newport, 1989; Thompson & Newport, 2007). However, one problem with these studies is that they do not convincingly show that test performance reflects abstract grammatical categories as opposed to surface similarity to training items, so the generalisability of what is learned is not clear. Nevertheless, they do suggest that low-level statistical information could feed into the process of learning phrase structure.

If one is to develop a theory of statistical learning, one needs a theory of what is learned. Currently, the dominant approach is to regard statistical learning as contingency learning, which, broadly speaking, refers to associative learning of the predictability of outcomes given cues (Shanks, 1995). N. C. Ellis (2006a, 2006b) provides an extensive discussion of the possible role of contingency learning in SLA. Although probability theory provides formal methods for calculating contingency, this does not tell us how this is achieved in human brains. Connectionist models provide one indication of how contingency could be computed in a psychologically, if not neurally, plausible way (Shanks, 1995). An especially interesting type of connectionist model is the 'simple recurrent network' (SRN), which is specialised for the kind of sequence learning that is assumed to occur in the procedural system. The details of such models need not concern us here (see Elman, 1990, for examples); suffice it to say that such models treat sequence learning as a prediction task. For any particular training item, say the string ABCD in an AG experiment, the network is taught to predict the next element in the string, taking into account not only the current element, but also its context (e.g. it is trained to predict B from A, and C from B in the context of A). The network essentially learns the context-dependent contingencies between elements in training strings. SRNs have been used to successfully model implicit learning in lexical segmentation (Christiansen, Allen, & Seidenberg, 1998; Elman, 1990), AG and SRT experiments (Cleeremans & McClelland, 1991; Kinder & Shanks, 2001) and, using a somewhat different architecture, phonotactic constraints

(Warker & Dell, 2006). See Cleeremans and Dienes (2008) for a review of computational models of implicit learning.

B. Abstraction and Transfer in Statistical Learning

One obvious limitation of chunking and connectionist approaches is that responses to test items are a function of what appears to be surface similarity to training items. The response to novel test items is determined by how well they reflect the probabilistic structure of the training set. But in the case of natural language learning, we assume that learners internalise abstract grammatical representations that can be applied to word combinations that bear no similarity to previous utterances (e.g. we can appreciate the sense in which the sentence *Green ideas sleep furiously* is syntactically well-formed). Is there evidence that representations of sufficient abstraction to support this kind of generalisation can be learned implicitly by humans or connectionist networks?

Research within the AG tradition has examined the abstraction issue by simply changing the letter set between training and test (e.g. the string AABCAB in training would correspond to the string DDEFDE in test). Typically, GJT performance for changed letter sets is lower than that for same letter sets, although still significantly above chance (Knowlton & Squire, 1996; Matthews et al., 1989). More impressive are demonstrations of transfer to different modalities, as for example when the grammar used to generate letter sequences in training is used to generate tone sequences at test (Altmann, Dienes, & Goode, 1995), although here too performance is lower than for same modality test items. The question then is what kind of knowledge supports this limited generalisation? Is it knowledge of the abstract structure of the grammar, as argued originally by Reber (1967)? The current consensus appears to be that this is not strictly the case; rather people pick up 'some rules about permissible locations of letter repetitions, alternations, or dependencies between different parts of the letter strings' (Knowlton & Squire, 1996, p. 179). Sensitivity to abstract patterns of alternation and doubling in syllable strings has been demonstrated in seven-month-old infants (Marcus, Vijayan, Bandi Rao, & Vishton, 1999) and even in tamarin monkeys (Hauser, Weiss, & Marcus, 2002). It appears that human and primate perceptual mechanisms code events in terms of change. It would not be unreasonable to assume that such codings could form part of the input to an associative learning system. When this is done, transfer problems of the type explored by Marcus become trivial for connectionist networks and hence do not pose a challenge to connectionist and other associative learning approaches (Dominey & Ramus, 2000; McClelland & Plaut, 1999). In a natural language context, Pacton, Perruchet, Fayol, and Cleeremans (2001) show implicit learning of constraints on consonant doubling in French and that the results can be modelled by an SRN.

What about transfer to sentences with new lexis in natural language grammar? Robinson (2005) examined incidental learning of Samoan by Japanese participants, targeting rules for ergative marking in transitive sentences (*ave e le tama le taavale* = drive ERG the boy the car), an incorporation rule (*inu-pia le tama* = drink-beer the boy) and locative (*taalo le tama i le paka* = play the boy IN the park). Performance on an immediate GJT showed high accuracy on old grammatical sentences, but poor performance on new grammatical and ungrammatical sentences (except for the locative, which corresponds to an English structure). Thus, there was an almost complete failure to transfer the knowledge of the novel trained structures to new sentences. However, it must be noted that there were only nine different training sentences, each of which was repeated 50 times. Each verb occurred in only one context during training, and so individual verbs were strongly associated with specific word order patterns. These results are therefore best explained in terms of learning of the kinds of item-based constructions that are characteristic of the early stages of first language acquisition (FLA) (Lieven & Tomasello, 2008; Tomasello, 2000), which here is encouraged by overlearning of a very small training set.

In an earlier study, Robinson (1996) did show a degree of transfer of a rule for forming pseudo-clefts of location (e.g. *Where my parents vacation is in Europe*, *Where LA is is in California*). After exposure to sentences illustrating this structure in an implicit (memory) task, learners of English were above chance on a GJT using sentences that contained different content words from the training sentences, despite being unable to state the rule in a debriefing. Whilst this might appear to show implicit learning of abstract structure, Robinson points out that it could also be based on memory for doubling of the verb *to be* (which was repeated from training as either an *is is* or *are is* pattern) along with patterns of plural and tense marking. One could go further and argue that the participants learned word order templates or 'constructions' (Goldberg, 1995, 2006) using a combination of specific lexis and abstract categories, such as *Where* N-PL V *is* PP and *Where* N *is is* PP, which then transfer to sentences with new lexis.

More direct evidence for learning of word order templates that generalise to new lexis comes from the studies of incidental learning of German word order by Rebuschat (2008, Experiment 3) that were mentioned earlier (see also Rebuschat & Williams, 2006). A unique aspect of these studies was that the materials used English lexis but German word order. Participants performed a semantic plausibility judgement task on 120 training sentences, 40 for each of three German structures (examples of each structure: V2, *In the evening ate Rose excellent dessert at a restaurant*; VF-V1, *Since his teacher criticism voiced, put Chris more effort into his homework*; V2-VF, *George repeated today that the movers his furniture scratched*). Participants then received a surprise GJT on sentences containing new lexis. Sentences that repeated grammatical patterns encountered in training were accepted at levels well above chance (and better than a control group who had received no training), whilst performance on ungrammatical sentences was at chance. Thus, there was rapid incidental learning of abstract word order patterns, but judging by performance on ungrammatical items, no learning of the actual verb placement rules (see below). Acceptance of grammatical items was likely to reflect template representations using categories that were sufficiently abstract to support transfer to new lexis (e.g. categories such as subject, verb and time adverbial). Williams and Kuribara (2008) obtained similar results in a study of incidental learning of Japanese word order. We also performed connectionist (SRN) simulations in which the input was coded as sequences of grammatical categories rather than words. The simulations accounted well for the relative difficulty of most of the test items, suggesting that the participants were learning sequences of abstract categories in much the same way as they learn the sequences of letters in AG experiments. These studies demonstrate incidental learning of word order patterns represented at a sufficient level of abstraction to support transfer to sentences with new lexis.

C. Implicit Learning of Grammatical Form–Meaning Connections

All of the demonstrations of implicit learning effects that have been mentioned so far essentially involve learning contingencies between representations within the same domain—be they letters, phonemes, syllables or grammatical categories. But what about learning associations between forms and meanings? After all, from functionalist and usage-based perspectives, form–meaning mappings lie at the heart of language processing and learning (Bates & MacWhinney, 1989; Goldberg, 1995; Tomasello, 2003). According to the Competition Model (Bates & MacWhinney, 1989), learners track the probabilities with which input cues in the domains of word order, morphology and meaning are associated with specific interpretations. Basic principles of associative learning such as cue competition, salience, interference, overshadowing and blocking can be used to explain first and second language learning phenomena such as morpheme acquisition orders, fossilisation, transfer and interference (N. C. Ellis, 2006b). Clearly, knowledge of the cue-interpretation contingencies underlying language is implicit. We have no awareness of these contingencies or of the process by which they are constantly updated through usage.

However, it has been argued that whilst the tuning of existing form–meaning connections may proceed implicitly, establishing new connections requires explicit learning processes. This is because of the requirement to integrate information across different cognitive systems, and such 'relational encoding' (Eichenbaum, Otto, & Cohen, 1994) requires declarative memory systems such as the hippocampus (N. C. Ellis, 1994; N. C. Ellis, 2005). The main line of evidence for this argument is that vocabulary acquisition is impaired in amnesia (Gabrieli, Cohen, & Corkin, 1988). However, the kind of hippocampus-dependent relational encoding that is assumed to be required for learning form–meaning connections does not appear to be confined to explicit learning. It is important for certain types of implicit learning as well, as suggested by the research on 'contextual cuing' that will be described later (Chun & Phelps, 1999; Park, Quinlan, Thornton, & Reder, 2004). The fact that amnesics cannot learn form–meaning connections does not mean that explicit memory is necessary, but only that this kind of learning depends on an intact hippocampus. Thus, it may be possible to obtain implicit learning of form–meaning connections in the normal population.

There have been few empirical investigations of implicit learning of novel form–meaning connections. DeKeyser (1995) employed a miniature artificial language with rich inflectional morphology for marking biological gender, number and object role. Some sample sentences are *Bep-on warufk-at rip-us* (Worker-PL build-PL house-OBJ; 'The workers are building a house') and *Hadeks-on wulas-in-it melaks-is-on* (Queen-PL peel-FEM-PL apple-OBJ-PL; 'The queens are peeling apples'). During training participants had to indicate whether a given sentence correctly described a picture, and in the test phase, they were required to describe pictures using the artificial language. When it was possible to use stem–inflection combinations that had occurred in training, performance was very good, but when tested on items that required novel stem–inflection combinations, performance was at chance, indicating no learning of the semantic correlates of the inflectional morphemes. This was despite extensive training of 20 learning sessions of 25 min.

More positive evidence has come from a series of studies by myself (Williams, 2005) and Janny Leung (Leung, 2007; Leung & Williams, 2006; Leung & Williams, in preparation). These studies all had a much narrower focus than DeKeyser's, involving fewer novel forms and fewer meaning distinctions. The training tasks also involved greater attention to the relevant forms and meanings and test procedures that were potentially more sensitive to implicit knowledge than the production task used by DeKeyser. In all cases, the participants were taught just four novel grammatical morphemes (*gi*, *ro*, *ul* and *ne*, which might be introduced as determiners) and told that they encoded a certain meaning dimensions (e.g. *gi* and *ro* occurred with near objects, *ul* and *ne* with far objects). The aim was to see if the participants would spontaneously induce a correlation with another, hidden, meaning dimension (e.g. that *gi* and *ul* were used with animate nouns and *ro* and *ne* with inanimate nouns). The novel forms were embedded in English carrier phrases or sentences (e.g. *I was terrified when I turned around and saw gi lion right behind me*) upon which the participants had to perform tasks that forced them to process the novel determiners in relation to the meaning dimension they had been taught. After training, Williams (2005) found significantly above-chance selection of determiners according to the non-instructed meaning dimension (in this case animacy) in entirely novel sentences even for participants who reported no awareness of the relevance of that dimension to determiner usage.

Extending this work, Leung (2007) developed a novel reaction time methodology that hinged on the use of form–meaning connections to direct attention. For example, suppose a person knows that the determiner *gi* always occurs with animate objects. If presented with a display containing a picture of a lion and a clock, on hearing the phrase 'gi lion', they would be able to orient their attention to the lion on hearing *gi*, that is, even before hearing the word *lion*. Their time to respond to this object in a reaction time task would therefore be facilitated. If the knowledge of the animacy correlation were implicit, then this orienting effect would occur outside of awareness, providing an online measure of the automatic use of implicit knowledge in comprehension. The experiments provided evidence for such effects across a range of form–meaning correlations: animacy, thematic role (agent/patient) and, in a case where the novel forms acted as reflexive pronouns, reflexivity.

Constructionist approaches stress the acquisition of linking rules between word order and a verb's argument structure. Can these be learned implicitly? Casenhiser and Goldberg (2005) provide evidence for rapid acquisition of the mapping between SOV word order and novel verbs encoding appearance (e.g. *The spot the king mooped* was paired with a video of a spot appearing on a king's nose). English-speaking children between age 5 and 7 simply observed pairings of videos of sentences over a 3 min training period. In the test phase, a sentence containing a new verb had to be matched with either of two videos. If the sentence had SOV order the children tended to choose the video depicting a scene of appearance, whereas if it had the familiar SVO order they tended to choose a scene depicting a transitive action. This learning effect was claimed to be implicit because the children were unable to articulate the meanings of the novel verbs. Whilst it would be interesting to see whether above-chance responding would still be obtained using more sensitive measures of awareness (such as subjective ratings of guessing and intuition), this experiment does provide impressive evidence of rapid inductive learning of linking rules involving both a novel argument structure and a novel word order.

When considering what is implicitly learnable we must obviously bear in mind the possibility of interactions with prior linguistic knowledge. In the case of learning grammatical form–meaning connections there may be involvement of grammatical processes (e.g. in searching for a basis for agreement) or the search space for possible meanings may be constrained by biases towards the kinds of distinctions that are likely to be encoded in natural language grammars (Bickerton, 1999). In fact, Williams (2004, 2005) found that implicit learning effects were greater in participants who knew languages with grammatical gender systems, suggesting that prior linguistic knowledge facilitated learning. On the other hand, the novel appearance meaning that was learned in Casenhiser and Goldberg (2005) is not encoded in English and may even fall outside the scope of whatever universal linking rules have been proposed (see Goldberg, 2006, p. 83). The way in which prior knowledge influences implicit learning is clearly an important issue in SLA and arises again in relation to determining what is learnable, as we shall see below.

D. Limitations of Implicit Learning

There is a tendency to believe that a statistical learning approach implies that any regularity in the environment can be acquired. For example, Hayes and Broadbent (1988) characterise implicit learning as involving 'the unselective and passive aggregation of information about the co-occurrence of environmental events and features' (*ibid*., p. 251), and Cleeremans and Jiménez (2002), following O'Reilly and Munakata (2000), characterise it as 'model learning', the goal of which is to 'enable the cognitive system to develop useful, informative models of the world by capturing its correlational structure' (*ibid*., p. 18). However, it is becoming increasingly apparent that not all environmental regularities are equally learnable by implicit means. The question is though, do these limitations reduce the significance of implicit learning to language, or, as some believe, do they help us understand how the form of language might be constrained by our human cognitive capacities?

Long-Distance Dependencies

An important feature of natural languages is that they contain long-distance, or non-adjacent, dependencies, both in phonology and in syntax. Saffran et al.'s (1996b) work on segmentation of syllable sequences suggested that people rapidly learn associations between adjacent syllables. Newport and Aslin (2004) went on to examine learning non-adjacent dependencies in syllable sequences; for example, the frame ba_te recurred in the sequence, but with random intervening syllables. Over a series of seven experiments manipulating exposure, language size and task, they were unable to obtain any learning effects. These null results reveal a surprising limitation on statistical learning. However, learning effects were obtained when the dependencies concerned individual consonants and the intervening segment was a vowel; for example, the consonant frame p_t could be learned even though it occurred with random intervening vowels. This latter situation is more like that found in Semitic languages where words are formed from consonant frames. The learnability of the more natural system may derive from Gestalt principles of perceptual organisation, which group elements of a common type together (where 'type' is defined here in terms of different phonological tiers for consonants and vowels). Thus, even if statistical learning is limited to adjacent elements, this constraint can be overcome by bringing non-adjacent elements into adjacency at a common level of linguistic representation. Interestingly, tamarin monkeys show the converse pattern, being able to learn non-adjacent dependencies between syllables but not consonants (Newport, Hauser, Spaepen, & Aslin, 2004). Whilst it is not clear what kind of representation the monkeys impose on the input that makes this possible, it appears that the form of the coding provided by prior knowledge determines what is learnable (just as in the case of learning patterns of alternation or doubling mentioned earlier).

AG research provides another example of the problem of learning long-distance dependencies. In 'biconditional grammars', letter sequences such as TPPV.XCCS and TPVP.XCSC are formed by substituting letters (in this example any T on the left is substituted by an X on the right, any P with a C and any V with an S). This kind of grammar is not learnable under incidental training conditions (Johnstone & Shanks, 2001; Matthews et al., 1989). Given Newport & Aslin's (2004) failure to find learning of associations between non-adjacent syllables, this is hardly surprising.

Embedding in natural language syntax is another domain where long-distance dependencies are critical. The ability to understand and produce such structures depends on a grasp of the principle of recursion, which it has been claimed is uniquely human (Hauser, Chomsky, & Fitch, 2002). Fitch and Hauser (2004) compared a context-free grammar with 'centre embedding' of the form A^nB^n (e.g. AAABBB) with a finite state grammar of the form $(AB)^n$ (e.g. ABABAB), where A and B stand for different categories of syllables, spoken by male and female speakers, respectively. Both grammars were learnable by humans (under incidental conditions), but only the finite state grammar was learnable by tamarin monkeys. It was subsequently found that only the A^nB^n grammar activates Broca's area (Friederici, Bahlmann, Heim, Schubotz, & Anwander, 2004). However, it is important to note that the ability to distinguish, say, an AAABBB string from an ungrammatical AAABBA string does not necessarily entail sensitivity to centre embedding. It could just reflect an understanding that there has to be an equal number of A items followed by an equal number of B items. Indeed, it has been found that the ability to reject ungrammatical AAABB items after training on the A^nB^n grammar is confined to subjects who reported using a counting strategy (Hochmann, Azadpour, & Mehler, 2008). Starlings have also been shown to perform this kind of discrimination (Gentner, Fenn, Margoliash, & Nusbaum, 2006), but again this is more likely to be due to their counting (or more probably subtising) abilities than an appreciation of recursion (Corballis, 2007). None of these experiments test the essential characteristic of embedding, which is that there are long-distance dependencies between specific elements. That is, for the case of English the relevant structure is not AAABBB but $A_1A_2A_3B_3B_2B_1$. When an AG is constructed along these lines, it turns out to be unlearnable even by humans under incidental training conditions (Perruchet & Rey, 2005), again pointing to problems learning long-distance dependencies.

Does this line of research necessarily pose problems for an implicit, associative learning account of natural language? Not necessarily. First, the unlearnability of non-adjacent dependencies has generally been demonstrated when the intervening stimuli are randomly generated. In contrast, SRT research shows that people are sensitive to the predictiveness of elements up to three stimuli back in the sequence (Cleeremans & McClelland, 1991). The difference here is that the intervening material is structured according to a finite state grammar, a situation that is more like long-distance dependencies in natural language syntax than the systems studied by Newport et al. (2004). Second, whether effects are obtained probably depends on whether the separated items can be brought into adjacency through Gestalt principles of perceptual organisation, that is, through attentional processes (Pacton & Perruchet, 2008, see below). The existence of a common underlying representation provides an underlying motivating force towards perceptual grouping, as shown by Newport et al.'s (2004) contrasting results for syllable and consonant frames. In the case of long-distance dependencies in syntax, meaning provides a level of representation at which disparate forms can be related to each other. For sentences such as *The mouse the cat chased escaped*, the knowledge that the cat is likely to have done the chasing and the mouse the escaping would surely aid the learner in bringing the relevant words into adjacency at the level of meaning. Thus, just because statistical learning at the level of form cannot solve this problem does not mean that such structures are unlearnable.

Learning Grammatical Categories by Distributional Analysis

It has been argued that abstract lexical and grammatical categories can be learned by distributional analysis of forms, that is, by analysing patterns of lexical co-occurrence (Maratsos, 1982; Redington & Chater, 1998). The idea is that forms that show similar patterns of co-occurrence with other forms come to constitute a category. Grammatical gender classes provide a simple example, where nouns of different classes might occur with different sets of articles. However, there is little evidence that abstract grammatical categories can be formed through implicit/incidental learning when those categories are 'arbitrary', that is, when they are not also correlated with semantic or phonological properties of the words. Even when people have good memory for the specific items they have been trained on, their behaviour on tests of generalisation shows no learning of the underlying noun class distinction (Braine, 1987; Braine et al., 1990; Brooks, Braine, Catalano, & Brody, 1993; Frigo & McDonald, 1998; Gerken, Wilson, & Lewis, 2005).[4] Frigo and McDonald (1998) argue

[4]Thompson and Newport (2007) claim to demonstrate learning of phrase structure based on word classes that are defined purely distributionally. But, as they themselves admit, their experiment did not permit a proper test of generalisation. Mintz (2002) found learning of a word class that was signalled by two surrounding markers (rather than the one marker used in previous studies) and when there was only one withheld example of the paradigm in training. These factors may have provided sufficiently strong distributional cues to permit construction of the class, but one must wonder whether they are representative of natural language.

that models of noun class learning that depend on pure distributional analysis are 'too powerful' (*ibid*., p. 237). However, the power of connectionist networks to learn arbitrary noun classes may depend on the specific architecture that is adopted. A network that modelled the kind of passive, and unsupervised, model learning occurring in incidental learning situations was indeed unsuccessful at learning the same system that had been shown to be unlearnable by humans (Williams, 2003). Connectionist models are not necessarily too powerful in this respect.

Of course, in natural languages, distributional information is just one of the possible cues to grammatical classes since there are also phonological and semantic correlates of varying predictiveness in different languages (Kelly, 1992). The same studies that showed no learning on the basis of distributional information also showed high levels of learning when such correlates were included, even if they were present only for a subset of category members (Braine, 1987; Brooks et al., 1993; Frigo & McDonald, 1998; Gerken et al., 2005). However, even here the picture is not so clear because in some studies generalisation to nouns without the relevant cues is at best marginally significant (Brooks et al., 1993) or not significant at all (Frigo & McDonald, 1998), suggesting that in these cases the nouns were not actually represented as belonging to different abstract grammatical classes (see Williams, 2003, for discussion).

Do these limitations on implicit (or rather incidental) learning limit its relevance to SLA? Not necessarily. Gender classes are notoriously difficult for learners to master (Carroll, 1999; Holmes & Dejean de la Batie, 1999), and processing of gender agreement is impaired in the second language (Guillelmon & Grosjean, 2001; Sabourin & Stowe, 2008). Furthermore, in contrast to first language learners, second language learners are overly sensitive to phonological cues to gender in languages such as French (Holmes & Dejean de la Batie, 1999) and Russian (Taraban & Kempe, 1999), in line with the above research on artificial languages. What we see in these experiments, therefore, could be just a reflection of a limitation of the adult implicit learning mechanism. See Blom and Polisenská (2008) and accompanying articles for recent research on gender in SLA and FLA.

Other Grammatical Rules

Studies that have examined other kinds of natural language regularities have failed to find implicit learning effects. Ellis (1993) examined the soft mutation rule in Welsh. A word like *trwyn* (nose) would appear in its citation form in isolation or in a context such as *blae mae trwyn* (where is a nose), but the initial /t/ mutated to /d/ in contexts such as *ei drywn o* (his nose). Not all initial consonants displayed mutation, however. After receiving examples in an implicit (Welsh-to-English translation) task, there was no evidence that the soft mutation rules had been learned. Well-formedness decision on items that had been received in training was 82% (in the 'yoked random' condition), whereas on incorrect items such as *ei trwyn o*, it was 50%. Subjects clearly knew that *ei trwyn o* was different from any training items; what they did not feel confident about was whether it was well-formed or not. Given that each of the five different mutation patterns was only exemplified by two different examples in training, it is quite possible that the quantity of input fell short of the 'critical mass' required to move beyond individual items to generalisations (Marchman & Bates, 1994). But the system itself may be beyond the scope of implicit learning. What appears to be required is that certain contexts be identified as triggering an abstract notion of ‘mutation’ that then selects an alternate form of the noun which is derived according to phonological rules (such as addition of voicing). As in the case of noun classes, it is not obvious how a connectionist network simulating unsupervised learning could acquire this kind of system. Also bear in mind that we do not know whether soft mutation rules are acquired (as opposed to learned) in SLA.

Robinson (1997) examined learning of a phonological constraint on the dative alternation, such that monosyllabic verbs could take both prepositional (PO) and double-object (DO) datives, whereas disyllabic verbs took only the PO dative. After implicit (memory for form) and incidental (focus on meaning) training, there was no evidence for generalisation of this rule to new verbs. Once again, this is a difficult learning problem. Learners would presumably have to unconsciously ‘realise’ that certain verbs that would have otherwise been expected to appear with both PO and DO in fact only occur with PO (Goldberg, 2006, refers to this as ‘statistical pre-emption’). The next step would be to make a form-level generalisation across these verbs. But each verb in training was associated with only one structural alternative, making statistical pre-emption impossible. It should also be noted that most analyses assume that it is the semantics of verbs, rather than their form, that determines their argument structure possibilities (Goldberg, 1995; Pinker, 1989), and so it is not even clear that this system would be learnable under naturalistic conditions. Thus, as in the case of Ellis (1993), it is rather difficult to draw conclusions about what may or may not be learnable because it is not clear whether the input that was provided licensed the generalisations that were tested.

With regard to learning syntactic rules, as described earlier, Rebuschat (2008) examined incidental learning of German word order rules using materials that combined English lexis with German word order patterns (e.g. *Yesterday scribbled David a long letter to his family*). Whilst patterns that had been received in training were accepted in a GJT, performance on ungrammatical items showed that the underlying word order rules had not been learned. For example, there was high endorsement of single-clause verb-final structures (e.g. **After dinner Susan an old car with her savings bought*). It appeared that whilst participants had learned possible verb positions at clause level (verb-second, verb-first, verb-final), they had not learned how verb position was determined by clause type and clause sequence.

Williams and Kuribara (2008) adopted a similar methodology in order to examine the acquisition of Japanese scrambling. From a generative perspective, scrambling is an optional syntactic operation that moves a phrase in the direction opposite to the head direction (Saito & Fukui, 1998). So in a right-headed language like Japanese, scrambling takes place to the left. The materials employed English lexis combined with Japanese word order and case markers (e.g. *John-ga pizza-o ate*). The training set contained a majority of simple and embedded canonical SOV structures and a minority of scrambled structures, but only scrambled structures involving movement of the direct object occurred in training (e.g. OSV). The training task was to perform plausibility judgements on a total of 194 sentences, and learning was assessed by a surprise GJT on sentences with new lexis. There was evidence of learning canonical structures, confirming the incidental acquisition of abstract grammatical patterns.

With regard to scrambling, 44% of participants showed a general preference for canonical structures and did not reliably endorse even the scrambled structures they had been trained on. The remaining 56% of participants accepted these structures and even generalised to certain, but not all, scrambled structures that they had not been trained on (involving fronting of an indirect object). However, they also failed to reliably reject structures that manifested the word orders reflecting a head-initial (i.e. English) parameter setting (e.g. SVO). We concluded that even amongst these participants there was no learning of scrambling defined in terms of optional movement constrained by head direction. On the other hand, a connectionist simulation (using an SRN trained on sentences coded as sequences of grammatical categories) provided a good fit to the GJT data, taking into account that performance was also affected by a general preference for canonical structures in some participants, as well as processing difficulties involved in embedded structures. It appeared that GJT performance on both grammatical and ungrammatical items was strongly influenced by their similarity to training sentences (as determined by context-dependent contingencies in the sequences of grammatical categories). Note that the failure to obtain evidence for incidental learning of scrambling is consistent with reported problems acquiring scrambling in adult SLA (Iwasaki, 2003) and contrasts with the apparent ease with which scrambling is acquired in FLA, despite its rarity in the input (Murasugi & Kawamura, 2005).

E. Conclusions

The studies reviewed in the first part of this section show that it is possible to obtain implicit learning of linguistically relevant regularities. Humans possess a powerful learning mechanism that can absorb the statistical structure of the environment, defined as the contingencies between events. This type of learning is successful in the areas of lexical segmentation, phonological and orthographic structure, phrase structure and grammatical form–meaning connections. It may also support the rapid absorption of word order patterns (templates, schemas or constructions), represented at a sufficient level of abstraction to be independent of lexical content. But there appear to be limits to what can be learned in this way. There is evidence that implicit learning is temporally constrained, so that associations between events are only learned if they are adjacent or brought into adjacency through some other means (by attention or by virtue of the way they are represented). Whether this causes a problem for learning long-distance dependencies in language is debatable. But there also seem to be problems in going beyond the statistical properties of the input to deeper regularities that depend on abstract notions, as exemplified by the above studies on word classes, scrambling and possibly soft mutation. In the case of word classes and scrambling, there is evidence for similar difficulties in naturalistic SLA.

A common feature of many studies is that whilst there is good learning of trained items, and even transfer to new sentences with familiar underlying structures, there is poor rejection of ill-formed items in GJTs. In fact the simulations in Williams and Kuribara (2008) showed that performance on ungrammatical items could be explained largely by their similarity to trained items. This is not a phenomenon confined to laboratory studies. R. Ellis (2005) found that whereas performance on grammatical items in a speeded GJT loaded on the same factor as other speeded tasks assumed to tap implicit knowledge, performance on ungrammatical items in an unspeeded GJT loaded on the same factor as metalinguistic knowledge. Ellis concluded that whereas acceptance of grammatical items can be driven by implicit knowledge, reliable rejection of ungrammatical items is dependant on explicit knowledge. Similarly, Roehr (2008) suggests that implicit knowledge of a second language is exemplar –based, leading to prototype and similarity effects, whereas categorical, and context-independent, performance can be achieved only by using explicit metalinguistic knowledge.

Yet in the case of FLA, reliable rejection of ungrammatical sentences appears to be possible using implicit knowledge, and grammatical gender and Japanese scrambling are acquired with ease. Whether such divergences between FLA and SLA can be explained purely within an associative learning framework is at present unclear. For the moment it appears that what is currently known about the limitations of associative learning makes it a more promising approach to explaining SLA than FLA.

IV. THE ROLE OF ATTENTION IN IMPLICIT LEARNING

Implicit learning was characterised earlier as a form of incidental learning, that is, an automatic form of learning that occurs without intention. Given that a characteristic of automaticity is that it makes relatively few, if any, demands on attentional resources, the implication is that implicit learning can occur without attention. This is important in the context of SLA because it would mean that, for example, acquisition of one aspect of form could occur even if the learner's attention is focused on some other aspect of form or on meaning.

A common way to address this question is by using dual-task paradigms. For example, whilst performing an SRT task, participants might also be required to indicate whether tones are of low or high pitch. It has been found that learning is still obtained, sometimes being equivalent to that obtained under single-task conditions (Jiménez & Méndez, 1999) and sometimes reduced, but still significant (Shanks & Channon, 2002). Clearly, the demands of the secondary tasks prevent participants from actively trying to work out the underlying regularities of the system, and yet learning effects are still obtained. This is enough to suggest that learning is largely independent of the kinds of attention-demanding processes assumed to underlie explicit learning (for reviews, see Goschke, 1997; Shanks, 2005).

However, even granted that learning is incidental, we can still ask whether attention needs to be paid to the relevant stimuli for learning to occur. In dual-task situations, responses are required to stimuli in both tasks, and so it is obvious that participants are attending to the relevant stimuli, even if they do not have the resources to engage in additional explicit learning. What is the evidence for learning from unattended stimuli? Within SLA we may wonder whether learning about a particular form can occur when attention is directed to meaning or other aspects of form.

It is widely assumed that learning is dependent on focal attention (Cowan, 1999; Logan & Etherton, 1994; Perruchet & Gallego, 1997). Only attended content, or more specifically, content that is in 'access' (as opposed to 'phenomenal') consciousness (Block, 1990) is remembered;[5] an assumption that within SLA is encapsulated by the 'noticing hypothesis' (Schmidt, 2001). For example, Leow (2000) found that only learners whose think-aloud protocols suggested that they had noticed certain verb forms during the training task showed learning of those forms in a post-test. Pacton and Perruchet (2008) showed that non-adjacent dependencies can be learned only if the subjects actively maintain the to-be-associated items in focal attention as part of the task they are set. As Pacton and Perruchet (2008) put it, 'associative learning is an automatic process that links together all the components that are present in the attentional focus at a given point' (*ibid*., p. 82).

Whilst noticing may be necessary for encoding instances of language use in memory, extraction of regularities across instances might still occur unconsciously (Robinson, 1995). We must separate awareness at the level of noticing instances of language from awareness at the level of understanding generalisations across them (Schmidt, 2001). For example, Rosa and Leow (2004) found evidence for learning a generalised rule in participants whose think-aloud protocols revealed awareness at the level of noticing but not at the level of understanding. Thus, attention facilitates memory encoding, but learning of generalisations may still be implicit.

However, we should not be too hasty in assuming that focal attention is always a necessary condition for implicit learning. There is evidence that associations can be learned between attended stimuli and ambient stimuli that are not focally attended. Such a mechanism might be relevant to implicit learning of associations between attended words and non-attended contextual information.

Vision research suggests that a certain amount of semantic processing occurs for even complex stimuli, such as natural scenes and faces, that are presented outside of the focus of attention (Koch & Tsuchiya, 2007). The phenomenon of 'contextual cuing' demonstrates implicit learning of the association between such stimuli and a focally attended target (Chun, 2000). In a visual search task, participants might be asked to locate a rotated T amongst a number of distracting rotated L's. What they do not know is that displays are repeated, such that certain spatial configurations of distracters are always paired with certain target positions. It is found that targets are located more quickly on these repeated trials than on trials where the distracter positions are determined randomly. In a subsequent recognition task, participants are unable to distinguish repeated arrays from random ones, suggesting that the learning effect is implicit.[6] Similar effects have been obtained when the target location is predicted by the shapes of the distracters, as opposed to their position (Chun, 2000), or even by aspects of their meaning (Goujon, Didierjean, & Marmèche, 2007). Thus, an

[5]Content that is in access consciousness is available to other cognitive systems via the global workspace and so, for example, can be reported. Phenomenal consciousness refers to sensations, such as the difference between red and green. Content that is only phenomenally conscious is assumed to be rapidly forgotten (Lamme, 2003).

[6]Contextual cuing effects are absent in amnesics, suggesting that an intact hippocampus is involved in implicit as well as explicit learning (Chun & Phelps, 1999; Park et al., 2004).

attended stimulus can pick up correlations with ambient stimuli, even if those stimuli are not focally attended. Seitz and Watanabe (2005) refer to this as 'task irrelevant learning' and report studies which show that it can occur even when the task-irrelevant stimuli are presented subliminally. Here, then, we have cases of learning without noticing. The target item is noticed, but the stimuli with which it comes to be associated are not.

In the domain of verbal learning, Logan and Etherton (1994) studied similar effects in situations where pairs of words were presented for semantic categorisation. For example, in their Experiment 5, participants were asked to respond when the member of a word pair cued by an arrow was a metal. Certain specific word pairings were repeated (e.g. *Gold* might always be paired with *Sky*), and response times to those pairings became faster than non-repeated pairings. Although subjects were clearly learning these associations incidentally, there was no test of whether learning was implicit. But as in the contextual cuing situation, a non-attended stimulus that is part of the context when a target response is made becomes associated with the target.

Seitz and Watanabe (2005) suggest that task-irrelevant learning is due to the alerting function of attention. When a task-relevant stimulus is detected there is a general alerting response (a release of neurotransmitters) that allows currently processed ambient stimuli to be associated with that target. This view stresses simultaneity; the target and non-target stimuli have to simultaneously activate mental representations for learning to occur. The importance of temporal contiguity in implicit learning of associations is also evident from research on verbal learning in amnesia (Gabrieli, Keane, Zarella, & Poldrack, 1997; Goshen-Gottstein & Moscovitch, 1995).

However, there is an important constraint on this kind of learning apart from timing. In situations where there is prior orientation of attention to a target, there is no learning of task-irrelevant associations. In another condition of Logan and Etherton's (1994) Experiment 5, the arrow appeared half a second before the word pair. Now subjects did not learn the repeated word pairings. Similarly, Toro, Sinnett, and Soto-Faraco (2005) found that lexical segmentation of syllable sequences by statistical learning (Saffran et al., 1996b) is completely eliminated when the subjects performed a demanding distracting task, for example monitoring a rapid stream of line drawings for repetitions. These results should not be surprising when we consider that there appears to be a lack of perceptual processing for stimuli when sustained attention is directed elsewhere (Dupoux, Kouider, & Mehler, 2003; Naccache, Blandin, & Dehaene, 2002). Returning to the issue of learning form–meaning connections, what these studies suggest is that if the task encourages sustained attention to form alone, then contextual associations will not be learned.

When considering the influence of attention on learning, we must also consider the fact that attention not only is directed to discrete stimuli in space or time, but also can be directed to different dimensions of the same stimulus. Here too, sustained attention to one dimension will eliminate learning effects related to another dimension. Toro et al. (2005) showed that no learning of lexical segmentation occurred when subjects had to monitor the syllable stream for pitch changes. Attention was focused on the syllables, but not on the relevant dimension. Jiménez and Méndez (1999) used an SRT task in which the sequence of positions was generated in the usual way by a finite state grammar. However, they also built in a separate regularity such that the identity of the characters that were used as stimuli also predicted the position of the next stimulus (e.g. a * predicted position A, whereas a ? predicted position C). Using the standard SRT procedure in which subjects simply indicate at which position each stimulus occurred, only the position-based regularity was learned. But when they also had to keep a running count of how often certain characters (x or *) occurred, learning of the second, identity-based, regularity was obtained. Both regularities were learned implicitly. Thus, awareness at the level of noticing stimulus identity led to learning without understanding its predictiveness. Together, these experiments also illustrate the importance of attending to the appropriate stimulus dimensions for learning the regularities that relate to them, even when attention is always apparently directed to the same stimuli.

Work on task-irrelevant learning and contextual cuing suggests that an attended, and noticed, word might implicitly, and unselectively, acquire associations to contextual information that is outside the focus of attention. However, based on the above, we can hypothesise that this will occur only under specific conditions: the word and the contextual information have to be simultaneously active, and attention must not be oriented in advance to either the word or the context, or to some irrelevant dimension of either. If the word and the contextual information are not simultaneously active, then unitisation through joint attention will be necessary. Clearly, implicit learning is highly sensitive to attentional effects. Only by working through the microstructure of learning processes at this level of detail will we be able to understand the precise conditions under which implicit learning is likely to occur.

V. CONCLUSION

Even though there is a long tradition of research on implicit learning dating back to Reber's seminal 1967 publication, one senses that the study of implicit language learning is still in its infancy. On the positive side, there are now clearly developed ideas on how to measure subjective states and technologies that can provide indications of the automaticity and nativelikeness of brain responses, so we are in a good position to at least identify when implicit knowledge has been acquired. There are also clearly developed ideas about how attention is involved in learning processes, and these should give a good indication of the task conditions under which implicit learning is most likely to occur. Whilst of practical relevance in themselves, these advances also provide important methodological groundwork for the investigation of the crucial theoretical issues concerning the nature of the implicit learning process itself and the nature of what is learnable. Here there is much more that could be done. When investigating what is learnable we need to consider how characteristics of the hypothesised learning mechanism might determine learnability (perhaps following the example of research on non-adjacent dependencies). Ideally, computational modelling will be used to explore the learnability of different regularities, helping us to make explicit what exactly a wholly empiricist and associative view of implicit learning predicts.

Certain areas of potentially important implicit learning research have been curiously neglected. We need to know far more about the influence of prior knowledge and where that knowledge comes from (L1, L2 or UG?). The issue of the interface between implicit and explicit knowledge is remarkably under-researched in both psychology and SLA. N. C. Ellis (2005) provides a theoretical framework for thinking about this issue, but there is a need for hard experimental evidence. We must consider not only how explicit knowledge can influence implicit learning, but also how implicit knowledge can become explicit (see Haider & Frensch, 2005, for an intriguing suggestion). The issue of individual differences has been dominated by Reber's (1989) hypothesis that implicit learning should be relatively immune to factors such as IQ and age. Whilst this appears to be true (for recent evidence, see Don, Schellenberg, Reber, DiGirolamo, & Wang, 2003; Gebauer & Mackintosh, 2007), does this mean that implicit learning is completely independent of all dimensions of individual differences? Robinson (2005) found no relationship between the Modern Language Aptitude Test (MLAT) and incidental learning of a natural language, but suggests that the component skills measured by the MLAT are probably more relevant to explicit than implicit learning. Given that implicit learning is a memory-driven process, one would expect it to be related to memory ability. There is some evidence for this connection in AG learning (Karpicke & Pisoni, 2004), but as yet there is no evidence from implicit learning of natural language.

Clearly, much more needs to be known about these issues before the exact role of implicit learning in SLA can be specified. However, given that we are clearly endowed with a powerful associative learning mechanism for unintentionally picking up aspects of the statistical structure of the environment, it would surely be absurd to argue that it makes no contribution to language learning. The goal is to specify exactly what that contribution is.

ACKNOWLEDGMENTS

Thanks to Peter Robinson, Michael Ullman and Chieko Kuribara for their comments on portions of this chapter.

REFERENCES

Altmann, G. T. M., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning Memory, and Cognition*, *21*, 899–912.

Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*, 321–324.

Bates, E., & MacWhinney, B. (1989). Functionalism and the competition model. In B. MacWhinney & E. Bates (Eds.), *The crosslinguistic study of sentence processing* (pp. 3–73). Cambridge: Cambridge University Press.

Bialystok, E. (1982). On the relationship between knowing and using forms. *Applied Linguistics*, *3*, 181–206.

Bickerton, D. (1999). Creole languages, the language bioprogram hypothesis, and language acquisition. In W. C. Ritchie & T. K. Bhatia (Eds.), *Handbook of child language acquisition*. San Diego, London: Academic Press.

Block, N. (1990). Consciousness and accessibility. *Behavioral and Brain Sciences*, *13*, 596–598.

Blom, E., & Polisenská, D. (2008). The acquisition of grammatical gender in Dutch. *Second Language Research*, *24*, 259–265.

Braine, M. D. S. (1987). What is learned in acquiring word classes: A step towards an acquisition theory. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 65–87). Hillsdale, NJ: Lawrence Erlbaum Associates.

Braine, M. D. S., Brody, R. E., Brooks, P. D., Sudhalter, V., Ross, J. E., Catalano, L., & Fisch, S. M. (1990). Exploring language acquisition in children with a miniature artificial language: Effects of item and pattern frequency, arbitrary subclasses, and correction. *Journal of Memory and Language*, *29*, 591–610.

Brooks, P. J., Braine, M. D. S., Catalano, L., & Brody, R. (1993). Acquisition of gender-like noun classes in an artificial language: The contribution of phonological markers to learning. *Journal of Memory and Language*, *32*, 76–95.

Buchner, A. (1994). Indirect effects of synthetic grammar learning in an identification task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 550–565.

Carroll, S. E. (1999). Input and SLA: Adults' sensitivity to different sorts of cues to French gender. *Language Learning*, *49*, 37–92.

Casenhiser, D., & Goldberg, A. E. (2005). Fast mapping of a phrasal form and meaning. *Developmental Science*, *8*, 500–508.

Christiansen, M. H., Allen, J., & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: A connectionist model. *Language and Cognitive Processes*, *13*, 221–268.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, *4*, 170–178.

Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience*, *2*, 844–847.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*, 406–417.

Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. In R. Sun (Ed.), *Cambridge handbook of computational psychology* (pp. 396–421). Cambridge: Cambridge University Press.

Cleeremans, A., & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R. M. French & A. Cleeremans (Eds.), *Implicit learning and consciousness* (pp. 1–40). Hove: Psychology Press.

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

Corballis, M. C. (2007). Recursion, language, and starlings. *Cognitive Science*, *31*, 697–704.

Cowan, N. (1999). An embedded-processes model of working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 62–101). Cambridge: Cambridge University Press.

Curran, T., & Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 189–202.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

DeKeyser, R. (2003). Implicit and explicit learning. In C. Doughty & M. Long (Eds.), *Handbook of second language acquisition* (pp. 313–348). Oxford: Blackwell.

DeKeyser, R. M. (1995). Learning second language grammar rules: An experiment with a miniature linguistic system. *Studies in Second Language Acquisition*, *17*, 379–410.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning Memory, and Cognition*, *26*(6), 1355–1367.

Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, *8*, 343–350.

Destrebecqz, A., & Cleeremans, A. (2003). Temporal effects in sequence learning. In L. Jimenez (Ed.), *Attention and implicit learning* (pp. 181–213). Amsterdam: John Benjamins.

Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioural and Brain Sciences*, *22*, 735–808.

Dienes, Z., & Scott, R. (2005). Measuring unconscious knowledge: Distinguishing structural knowledge and judgment knowledge. *Psychological Research*, *69*, 338–351.

Dominey, P. F., & Ramus, F. (2000). Neural network processing of natural language: I. Sensitivity to serial, temporal and abstract structure of language in the infant. *Language and Cognitive Processes*, *15*, 87–127.

Don, A. J., Schellenberg, E. G., Reber, A. S., DiGirolamo, K. M., & Wang, P. P. (2003). Implicit learning in children and adults with Williams syndrome. *Developmental Neuropsychology*, *23*(1–2), 201–225.

Dupoux, E., Kouider, S., & Mehler, J. (2003). Lexical access without attention? Explorations using dichotic priming. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 172–184.

Eichenbaum, H., Otto, T., & Cohen, N. J. (1994). Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences*, *17*, 449–518.

Ellis, N. C. (1993). Rules and instances in foreign language learning: Interactions of explicit and implicit knowledge. *European Journal of Cognitive Psychology*, *5*, 289–318.

Ellis, N. C. (1994). Vocabulary acquisition: The explicit ins and outs of explicit cognitive mediation. In N. C. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 211–282). London: Academic Press.

Ellis, N. C. (2005). At the interface: Dynamic interactions of explicit and implicit language knowledge. *Studies in Second Language Acquisition*, *27*, 305–352.

Ellis, N. C. (2006a). Language acquisition as rational contingency learning. *Applied linguistics*, *27*, 1–24.

Ellis, N. C. (2006b). Selective attention and transfer phenomena in L2 acquisition: Contingency, cue competition, salience, interference, overshadowing, blocking, and perceptual learning. *Applied Linguistics*, *27*, 164–194.

Ellis, R. (2005). Measuring implicit and explicit knowledge of a second language: A psychometric study. *Studies in Second Language Acquisition*, *27*, 141–172.

Ellis, R. (2006). Modelling learning difficulty and second language proficiency: The differential contributions of implicit and explicit knowledge. *Applied Linguistics*, *27*, 431–463.

Ellis, R., & Loewen, S. (2007). Confirming the operational definitions of explicit and implicit knowledge in Ellis (2005): Responding to Isemonger. *Studies in Second Language Acquisition*, *29*, 119–126.

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

Elman, J. L., Bates, E. A., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Erlam, R. (2006). Elicited imitation as a measure of L2 implicit knowledge: An empirical validation study. *Applied Linguistics*, *27*, 464–491.

Fitch, W. T., & Hauser, M. D. (2004). Computational constraints on syntactic processing in a nonhuman primate. *Science*, *303*, 377–380.

Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences*, *6*, 78–84.

Friederici, A. D., Bahlmann, J., Heim, S., Schubotz, R. I., & Anwander, A. (2004). The brain differentiates human and non-human grammars: Functional localization and structural connectivity. *Proceedings of the National Academy of Sciences*, *103*, 2458–2463.

Friederici, A. D., Steinhauer, K., & Pfeifer, E. (2002). Brain signatures of artificial language processing: Evidence challenging the 'critical period' hypothesis. *Proceedings of the National Academy of Science*, *99*, 529–534.

Frigo, L., & McDonald, J. L. (1998). Properties of phonological markers that affect the acquisition of gender-like subclasses. *Journal of Memory and Language*, *39*, 218–245.

Gabrieli, J. D. E. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, *49*, 87–115.

Gabrieli, J. D. E., Cohen, N. J., & Corkin, S. (1988). The impaired learning of semantic knowledge following medial temporal lobe resection. *Brain and Cognition*, *7*, 157–177.

Gabrieli, J. D. E., Keane, M., Zarella, M. M., & Poldrack, R. A. (1997). Preservation of implicit memory for new associations in global amnesia. *Psychological Science*, *8*, 326–329.

Gebauer, G. F., & Mackintosh, N. J. (2007). Psychometric intelligence dissociates implicit and explicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 34–54.

Gentner, T. Q., Fenn, K. M., Margoliash, D., & Nusbaum, C. (2006). Recursive syntactic pattern learning by songbirds. *Nature*, *440*, 1204.

Gerken, L., Wilson, R., & Lewis, W. (2005). Infants can use distributional cues to form syntactic categories. *Journal of Child Language*, *32*, 249–268.

Goldberg, A. E. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago: Chicago University Press.

Goldberg, A. E. (2006). *Constructions at work*. Oxford: Oxford University Press.

Goschke, T. (1997). Implicit learning and conscious knowledge: Mental representation, computational mechanisms, and brain structures. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts, and categories*. Hove: Psychology Press.

Goshen-Gottstein, Y., & Moscovitch, M. (1995). Repetition priming effects for newly formed associations are perceptually based: Evidence from shallow encoding and format specificity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1249–1262.

Goujon, A., Didierjean, A., & Marmèche, E. (2007). Contextual cueing based on specific and categorical properties of the environment. *Visual Cognition*, *15*, 257–275.

Green, P., & Hecht, K. (1992). Implicit and explicit grammar: An empirical study. *Applied Linguistics*, *13*, 168–184.

Guillelmon, D., & Grosjean, F. (2001). The gender marking effect in spoken word recognition: The case of bilinguals. *Memory and Cognition*, *29*, 503–511.

Hahne, A., & Friederici, A. D. (2001). Processing a second language: Late learners' comprehension mechanisms as revealed by event-related brain potentials. *Bilingualism: Language and Cognition*, *4*, 123–141.

Haider, H., & Frensch, P. A. (2005). The generation of conscious awareness in an incidental learning situation. *Psychological Research*, *69*, 399–411.

Haist, F., Musen, G., & Squire, L. (1991). Intact priming of words and nonwords in amnesia. *Psychobiology*, *19*, 275–285.

Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, *298*, 1569–1579.

Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the speech stream in a non-human primate: Statistical learning in cotton-top tamarins. *Cognition*, *78*, B53–B64.

Hauser, M. D., Weiss, D. J., & Marcus, G. (2002). Rule learning by cotton-top tamarins. *Cognition*, *86*, B15–B22.

Hayes, N. A., & Broadbent, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, *28*, 249–276.

Hochmann, J. R., Azadpour, M., & Mehler, J. (2008). Do humans really learn A^nB^n artificial grammars from exemplars? *Cognitive Science*, *32*, 1021–1036.

Holmes, V. M., & Dejean de la Batie, B. (1999). Assignment of grammatical gender by native speakers and foreign language learners. *Applied Psycholinguistics*, *20*, 479–506.

Hulstijn, J. H. (2003). Incidental and intentional learning. In C. Doughty & M. Long (Eds.), *Handbook of second language acquisition* (pp. 349–381). Oxford: Blackwell.

Hulstijn, J. H., & Hulstijn, W. (1984). Grammatical errors as a function of processing constraints and explicit knowledge. *Language Learning*, *34*, 23–43.

Isemonger, I. M. (2007). Operational definitions of explicit and implicit knowledge: Response to R. Ellis (2005) and some recommendations for future research in this area. *Studies in Second Language Acquisition*, *29*, 101–118.

Iwasaki, N. (2003). L2 acquisition of Japanese: Knowledge and use of case particles in SOV and OSV sentences. In S. Karimi (Ed.), *Word order and scrambling* (pp. 273–300). Oxford: Blackwell.

Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 236–259.

Johnstone, T., & Shanks, D. R. (1999). Two mechanisms in implicit artificial grammar learning? Comment on Meulemans and Van der Linden (1997). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 524–531.

Johnstone, T., & Shanks, D. R. (2001). Abstractionist and processing accounts of implicit learning. *Cognitive Psychology*, *42*, 61–112.

Karpicke, J. D., & Pisoni, D. B. (2004). Using immediate memory span to measure implicit learning. *Memory and Cognition*, *32*, 956–964.

Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, *99*, 349–364.

Kinder, A., & Shanks, D. R. (2001). Amnesia and the declarative/nondeclarative distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience*, *13*, 648–669.

Kinder, A., & Shanks, D. R. (2003). Neuropsychological dissociations between priming and recognition: A single-system connectionist account. *Psychological Review*, *110*, 728–744.

Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 169–181.

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, *11*, 16–22.

Krashen, S. (1981). *Second language acquisition and second language learning*. London: Pergamon.

Krashen, S. D. (1994). The input hypothesis and its rivals. In N. C. Ellis (Ed.), *Implicit and explicit learning of languages*. London: Academic Press.

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*, 12–18.

Leow, R. P. (2000). A study of the role of awareness in foreign language behavior. *Studies in Second Language Acquisition*, *22*, 557–584.

Leung, J. (2007). *Implicit learning of form–meaning connections*. Ph.D. thesis, University of Cambridge, Cambridge.

Leung, J., & Williams, J. (2006). Implicit learning of form–meaning connections. In R. Sun & N. Miyake (Eds.), *Proceedings of the annual meeting of the cognitive science society* (pp. 465–470). Mahwah, NJ: Lawrence Erlbaum Associates.

Leung, J., & Williams, J. N. (in preparation). *Implicit learning of grammatical form–meaning connections in semi-artificial languages*.

Lieven, E., & Tomasello, M. (2008). Children's first language acquisition from a usage-based perspective. In P. Robinson & N. C. Ellis (Eds.), *Handbook of cognitive linguistics and second language acquisition* (pp. 168–196). New York: Routledge.

Logan, G. D., & Etherton, J. L. (1994). What is learned during automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1022–1050.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 3–26.

Maratsos, M. (1982). The child's construction of grammatical categories. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 240–266). Cambridge: Cambridge University Press.

Marchman, V. A., & Bates, E. (1994). Continuity in lexical and morphological development: A test of the critical mass hypothesis. *Journal of Child Language*, *21*, 339–366.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning in 7-month-old infants. *Science*, *283*, 77–80.

Marslen-Wilson, W., & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, *10*, 29–63.

Matthews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J.-R., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1083–1100.

McClelland, J. L., & Plaut, D. C. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Sciences*, *3*, 166–168.

McLaughlin, J., Osterhout, L., & Kim, A. (2004). Neural correlates of second language word learning: Minimal instruction produces rapid change. *Nature Neuroscience*, *7*, 703–704.

Mintz, T. H. (2002). Category induction from distributional cues in an artificial language. *Memory and Cognition*, *30*, 678–686.

Mirman, D., Magnuson, J. S., Graf Estes, K., & Dixon, J. A. (2008). The link between statistical segmentation and word learning in adults. *Cognition*, *108*, 271–280.

Morgan, J. L., Meier, R. P., & Newport, E. L. (1989). Facilitating the acquisition of syntax with cross-sentential cues to phrase structure. *Journal of Memory and Language*, *28*, 67–85.

Morgan-Short, K. (2007). A *neurolinguistic investigation of late-learned second language knowledge: The effects of explicit and implicit conditions*. Ph.D. thesis, Georgetown University, Washington.

Murasugi, K., & Kawamura, T. (2005). On the acquisition of scrambling in Japanese. In J. Sabel & M. Saito (Eds.), *The free word order phenomenon: Its syntactic sources and diversity* (Vol. 69). Berlin/New York: Mouton de Gruyter.

Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, *13*, 416–424.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, *48*, 127–162.

Newport, E. L., Hauser, M. D., Spaepen, G., & Aslin, R. N. (2004). Learning at a distance II. Statistical learning of non-adjacent dependencies in a non-human primate. *Cognitive Psychology*, *49*, 85–117.

Norman, E., Price, M. C., Duff, S. C., & Mentzoni, R. A. (2007). Gradations of awareness in a modified sequence learning task. *Consciousness and Cognition*, *16*, 809–837.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Osterhout, L., McLaughlin, J., Pitkanen, I., Frenck-Mestre, C., & Molinaro, N. (2006). Novice learners, longitudinal designs, and event-related potentials: A means for exploring the neurocognition of second language processing. *Language Learning*, *56*, 199–230.

Pacton, S., & Perruchet, P. (2008). An attention-based associative account of adjacent and nonadjacent dependency learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 80–96.

Pacton, S., Perruchet, P., Fayol, M., & Cleeremans, A. (2001). Implicit learning out of the lab: The case of orthographic regularities. *Journal of Experimental Psychology: General*, *130*, 401–426.

Paradis, M. (1994). Neurolinguistic aspects of implicit and explicit memory: Implications for bilingualism and SLA. In N. C. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 393–419). San Diego: Academic Press.

Paradis, M. (2004). *A neurolinguistic theory of bilingualism*. Amsterdam/Philadelphia: John Benjamins.

Park, H., Quinlan, J., Thornton, E., & Reder, L. M. (2004). The effect of midazolam on visual search: Implications for understanding amnesia. *Proceedings of the National Academy of Sciences*, *101*, 17879–17883.

Perruchet, P., & Gallego, J. (1997). A subjective unit formation account of implicit learning. In D. C. Berry (Ed.), *How implicit is implicit learning?* (pp. 124–161). Oxford: Oxford University Press.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, *119*, 264–275.

Perruchet, P., & Rey, A. (2005). Does the mastery of center-embedded linguistic structures distinguish humans from nonhuman primates? *Psychonomic Bulletin & Review*, *12*, 307–313.

Perruchet, P., & Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language*, *39*, 246–263.

Pinker, S. (1989). *Learnability and cognition*. Cambridge, MA: MIT Press.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 855–863.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219–235.

Reber, A. S., & Allen, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, *6*, 189–221.

Reber, P. J., & Squire, L. R. (1999). Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behavioral Neuroscience*, *113*, 235–242.

Rebuschat, P. (2008). *Implicit learning of natural language syntax*. Ph.D. thesis, University of Cambridge, Cambridge.

Rebuschat, P., & Williams, J. N. (2006). Dissociating implicit and explicit learning of natural language syntax. In R. Sun & N. Miyake (Eds.), *Proceedings of the annual meeting of the cognitive science society* (p. 2594). Mahwah, NJ: Lawrence Erlbaum.

Redington, M., & Chater, N. (1998). Connectionist and statistical approaches to language acquisition: A distributional perspective. *Language and Cognitive Processes*, *13*, 129–191.

Robinson, P. (1995). Attention, memory, and the "noticing" hypothesis. *Language Learning*, *45*, 283–331.

Robinson, P. (1996). Learning simple and complex second language rules under implicit, incidental, rule-search, and instructed conditions. *Studies in Second Language Acquisition*, *18*, 27–67.

Robinson, P. (1997). Generalizability and automaticity of second language learning under implicit, incidental, enhanced, and instructed conditions. *Studies in Second Language Acquisition*, *19*, 223–247.

Robinson, P. (2005). Cognitive abilities, chunk-strength, and frequency effects in implicit artificial grammar and incidental L2 learning: Replications of Reber, Walkenfeld, and Hernstadt (1991) and Knowlton and Squire (1996) and their relevance for SLA. *Studies in Second Language Acquisition*, *27*, 235–268.

Roehr, K. (2008). Linguistic and metalinguistic categories in second language learning. *Cognitive Linguistics*, *19*, 67–106.

Rosa, E., & Leow, R. P. (2004). Awareness, different learning conditions, and second language development. *Applied Psycholinguistics*, *25*, 269–292.

Sabourin, L., & Stowe, L. (2008). Second language processing: When are first and second languages processed similarly? *Second Language Research*, *24*, 397–430.

Saffran, J. R. (2001). The use of predictive dependencies in language learning. *Journal of Memory and Language*, *44*, 493–515.

Saffran, J. R., Hauser, M., Seibel, R., Kapfhamer, J., Tsao, F., & Cushman, F. (2008). Grammatical pattern learning by human infants and cotton-top tamarin monkeys. *Cognition*, *107*, 479–500.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996a). Statistical learning in 8-month-old infants. *Science*, *274*, 1926–1928.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996b). Word segmentation: The role of distributional cues. *Journal of Memory and Language*, *35*, 606–621.

Saito, M., & Fukui, N. (1998). Order in phrase structure and movement. *Linguistic Inquiry*, *29*, 439–474.

Schmidt, R. (2001). Attention. In P. Robinson (Ed.), *Cognition and second language instruction* (pp. 3–32). Cambridge: Cambridge University Press.

Segalowitz, N. (2003). Automaticity and second languages. In C. Doughty & M. Long (Eds.), *Handbook of second language acquisition* (pp. 382–408). Oxford: Blackwell.

Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences*, *9*, 229–234.

Servan-Schreiber, D., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 592–608.

Shanks, D. R. (1995). *The psychology of associative learning*. Cambridge: Cambridge University Press.

Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 202–220). London: Sage.

Shanks, D. R., & Channon, S. (2002). Effects of a secondary task on 'implicit' sequence learning: learning or performance? *Psychological Research*, *66*, 99–109.

Shanks, D. R., & St. John, M. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367–447.

Squire, L. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*, 195–231.

Taraban, R., & Kempe, V. (1999). Gender processing in native and nonnative Russian speakers. *Applied Psycholinguistics*, *20*, 119–148.

Thompson, S. P., & Newport, E. L. (2007). Statistical learning of syntax: The role of transitional probability. *Language Learning and Development*, *3*, 1–42.

Tokowicz, N., & MacWhinney, B. (2005). Implicit and explicit measures of sensitivity to violations in second language grammar: An event-related potential investigation. *Studies in Second Language Acquisition*, *27*, 173–204.

Tomasello, M. (2000). The item-based nature of children's early syntactic development. *Trends in Cognitive Sciences*, *4*, 156–163.

Tomasello, M. (2003). *Constructing a language*. Cambridge, MA: Harvard University Press.

Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, *97*, B25–B34.

Ullman, M. T. (2001). A neurocognitive perspective on language: The declarative/procedural model. *Nature Reviews: Neuroscience*, *2*, 717–726.

Ullman, M. T. (2004). Contributions of memory circuits to language: The declarative/procedural model. *Cognition*, *92*(1–2), 231–270.

Ullman, M. T. (2005). A cognitive neuroscience perspective on second language acquisition: The declarative/procedural model. In C. Sanz (Ed.), *Mind and context in adult second language acquisition: Methods, theory, and practice* (pp. 141–178). Washington, DC: Georgetown University Press.

Warker, J. A., & Dell, G. S. (2006). Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 387–398.

Williams, J. N. (2003). Inducing abstract linguistic representations: Human and connectionist learning of noun classes. In R. van Hout, A. Hulk, F. Kuiken, & R. Towell (Eds.), *The interface between syntax and the lexicon in second language acquisition* (pp. 151–174). Amsterdam: John Benjamins.

Williams, J. N. (2004). Implicit learning of form-meaning connections. In B. VanPatten, J. Williams, S. Rott, & M. Overstreet (Eds.), *Form meaning connections in second language acquisition* (pp. 203–218). Mahwah, NJ: Lawrence Erlbaum Associates.

Williams, J. N. (2005). Learning without awareness. *Studies in Second Language Acquisition*, *27*, 269–304.

Williams, J. N., & Kuribara, C. (2008). Comparing a nativist and emergentist approach to the initial stage of SLA: An investigation of Japanese scrambling. *Lingua*, *118*, 522–553.

Witt, K., Nuhsman, A., & Deuschl, G. (2002). Intact artificial grammar learning in patients with cerebellar degeneration and advanced Parkinson's disease. *Neuropsychologia*, *40*, 1534–1540.