

IMPLICIT LEARNING

Implicit learning: Below the subjective threshold

ZOLTÁN DIENES

University of Sussex, Brighton, England

and

DIANNE BERRY

University of Reading, Reading, England

In this review, we consider three possible criteria by which knowledge might be regarded as implicit or inaccessible: It might be implicit only in the sense that it is difficult to articulate freely, or it might be implicit according to either an objective threshold or a subjective threshold. We evaluate evidence for these criteria in relation to artificial grammar learning, the control of complex systems, and sequence learning, respectively. We argue that the convincing evidence is not yet in, but construing the implicit nature of implicit learning in terms of a subjective threshold is most likely to prove fruitful for future research. Furthermore, the subjective threshold criterion may demarcate qualitatively different types of knowledge. We argue that (1) implicit, rather than explicit, knowledge is often relatively inflexible in transfer to different domains, (2) implicit, rather than explicit, learning occurs when attention is focused on specific items and not underlying rules, and (3) implicit learning and the resulting knowledge are often relatively robust.

In everyday life, we often learn to respond appropriately according to criteria that we can readily state. For example, when we have solved an algebraic problem, we can usually describe the steps that we have taken and the rules that we have used. However, not all our abilities depend on our explicitly knowing how to use them. There appear to be many examples of our learning to respond in some rule-like way without being able to state the rules that govern our behavior. For example, most of us learn to recognize and produce grammatical utterances in our native language without ever being able to say what the rules of the grammar are.

This type of learning is the focus of interest of this review. Implicit learning, as it is known, has aroused increasing interest in recent years (see Berry & Dienes, 1993; Reber, 1993; Seger, 1994; Shanks & St. John, 1994). In fact, understanding the processes involved in implicit learning, and its relationship to explicit learning, have become central goals in current cognitive psychology.

Implicit learning has been investigated in a wide range of experimental paradigms, including artificial grammar learning (e.g., Mathews et al., 1989; Reber, 1967, 1989), probability learning (e.g., Reber & Millward, 1968), con-

trol of complex systems (e.g., Berry & Broadbent, 1984; Stanley, Mathews, Buss, & Kotler-Cope, 1989), serial reaction time (SRT; e.g., Lewicki, Czyzewska, & Hoffman, 1987; Nissen & Bullemer, 1987), learning of conditional responses (e.g., Shanks, Green, & Kolodny, 1994), acquisition of invariant characteristics (e.g., Bright & Burton, 1984; Cock, Berry, & Gaffan, 1994; McGeorge & Burton, 1990), perceptual learning (e.g., Kolars & Roediger, 1984), learning of perceptual categories (e.g., Jacoby & Brooks, 1984), and second-language acquisition (e.g., Ellis, 1993; Michas & Berry, 1994). In this review, we will focus on the three paradigms that have generated the most research—namely, artificial grammar learning, the control of complex systems, and sequence learning. What these (and other implicit learning) situations have in common is that the person typically learns about the structure of a fairly complex stimulus environment, without necessarily intending to do so, and in such a way that the resulting knowledge is difficult to express. This is what we mean by implicit learning. In terms of controlling complex systems, for example, people can learn to reach and maintain specified levels of target variables without being able to freely describe to others how the decisions should be made (e.g., Berry & Broadbent, 1984; Stanley et al., 1989). Similarly, people can learn to classify exemplars of an artificial grammar and can acquire knowledge about the sequential structure of stimuli without adopting explicit code-breaking strategies and without being able to articulate fully any rules they might be using or the basis on which they are responding

The authors thank Axel Cleeremans, Pierre Perruchet, Arthur Reber, Henry L. Roediger III, and Michael Stadler for valuable comments on a previous version of the manuscript. Correspondence should be addressed to Z. Dienes, Laboratory of Experimental Psychology, University of Sussex, Brighton BN1 9QG, England (e-mail: dienes@epunix.biols.susx.ac.uk, or d.c.berry@reading.ac.uk).

(e.g., Nissen & Bullemer, 1987; Reber, 1989). In this review, we will consider the ways in which the knowledge underlying performance on these tasks might be considered implicit and whether implicit knowledge is qualitatively different from explicit knowledge.

CRITERIA OF IMPLICITNESS

Some researchers appear to assume as a default position that a piece of knowledge is conscious until conclusively proven otherwise (Dulany, Carlson, & Dewey, 1984; Shanks & St. John, 1994). On the other hand, Reber (1993) argued that, for evolutionary reasons, we should take the unconscious as primary, and it is ultimately irrelevant whether we can demonstrate in the laboratory that there is a piece of knowledge of which the subject is completely unconscious. Reber pointed out that reflexive consciousness has arrived late on the evolutionary scene, and perhaps uniquely only in *Homo sapiens*. Thus, there is a type of unconscious learning engaged in by other animals and to which we have tagged on a conscious learning system. But, according to Reber, the primitive unconscious system must still be present in *Homo sapiens* (and it must have, according to the principles of evolutionary biology, various characteristics, such as greater robustness and less dependence on age than the conscious learning system). In contrast to Reber, we take it that the relationship between consciousness and the primitive learning system is not in fact obvious: Evolution could well have arranged it such that the knowledge acquired by the primitive system is conscious in *Homo sapiens*. In order to argue either that there are separate conscious and unconscious learning systems or that learning systems do not differ in terms of consciousness, appropriate criteria for establishing unconscious or implicit knowledge and conscious or explicit knowledge need to be considered. Although the adequacy of any one criterion of consciousness can always be disputed as a necessary and sufficient criterion of consciousness (e.g., Erdelyi, 1986), we take the stand that if a criterion inspired by our everyday notions of consciousness separates qualitatively different types of knowledge, then the criterion is a psychologically real one. It tells us the sense in which knowledge in one system is unconscious, whether or not that fully maps onto our everyday notion of unconscious (contrast Neal & Hesketh, 1997).

The implicitness of some piece of knowledge refers to the fact that the knowledge is relatively inaccessible in some way. But the notion of inaccessibility is not an unambiguous concept. Knowledge may be inaccessible simply in the sense that it is difficult to elicit with free report. In the artificial grammar learning, dynamic control, and sequential reaction time paradigms, people seem to have little ability to describe what they know or to answer general questions about it. Indeed, Broadbent (1991) suggested that the typical experiment with which implicit learning is investigated is “one in which the person is able to choose the correct reaction while in a task, but

is later unable to recall the key characteristics that controlled the behavior” (p. 128). However, free report is a relatively insensitive and incomplete measure. Knowledge not free recalled in a given period of time may be recalled when the subject is given another attempt (Erdelyi & Becker, 1974). Also failure of free report could simply reflect the problem of having to retrieve large amounts of low-confidence knowledge, rather than reflecting a deeper incompatibility between the mechanisms employed in free report and the type of knowledge stored. Such problems prompted Shanks and St. John (1994) to suggest that we should regard knowledge as unconscious only if it cannot be elicited by tests satisfying information and sensitivity criteria—that is, the tests should be tapping exactly the knowledge responsible for the performance changes (their information criterion), and the tests should be sensitive to all relevant conscious knowledge (their sensitivity criterion). Cued report and forced choice tests, rather than free report, are more likely to satisfy these criteria, and, so, many investigators have argued that knowledge is implicit or unconscious only if it cannot be elicited by relevant cued report tests. The results of experiments using a number of different tests have shown that much knowledge can be elicited using cued report tests (we could say that the knowledge is above an objective threshold of consciousness; see below). Indeed, in their review of implicit learning, Shanks and St. John (1994) concluded that there has been no satisfactory demonstration of unconscious learning. However, we will argue that there is another potentially useful criterion of implicitness—that of a subjective threshold—that may be closer to our everyday intuitions about what unconscious means and that, when applied to the implicit learning literature, leads to a different conclusion from that of Shanks and St. John.

The terms *objective threshold* and *subjective threshold* were originally used in the subliminal perception literature (Cheesman & Merikle, 1984). In a typical subliminal perception experiment, subjects are given a sequence of trials in which a stimulus is either presented or not. For each trial, subjects say whether a stimulus was presented on that trial. Subjective threshold occurs at the level of discriminative responding for which subjects claim not to be able to detect perceptual information (i.e., they claim to be literally guessing), whereas objective threshold occurs at the level of discriminative responding corresponding to actual chance performance (i.e., subjects claim a stimulus was present no more frequently when it was present than when it was not). Thus, the subjective threshold is the point at which subjects do not know that they know that a stimulus was presented; the objective threshold is the point at which subjects do not know that a stimulus was presented. Cheesman and Merikle found that subliminal perception occurred in the sense of a stimulus being below subjective but not objective threshold. Cheesman and Merikle’s results imply that subliminal perception exists in one sense meant by the layperson: People might not believe that they saw anything and yet still show, by their above-chance guesses and by reliable

priming, that the information nonetheless affected their behavior (see Bornstein & Pittman, 1992, for further discussions of subliminal perception).

In terms of the implicit learning literature, people's knowledge could be said to be below an objective threshold if there is chance performance on a direct test (a cued report test) that directly measures some knowledge that we infer that the subject must have (because of its indirect effects). People's knowledge could be said to be below a subjective threshold if they lack metaknowledge about their knowledge. Dienes, Altmann, Kwan, and Goode (1995) pointed out that subjects could lack metaknowledge in at least two ways. First, they might believe they are literally guessing when in fact they are above chance, the criterion for subjective threshold used by Cheesman and Merikle (1984) and called the *guessing* criterion by Dienes et al. Second, they might lack metaknowledge because their confidence is unrelated to their accuracy, a criterion introduced by Chan (1992) and called the *zero-correlation* criterion by Dienes et al. We will take both ways to be possibilities for defining a psychologically real subjective threshold. Lacking metaknowledge might be empirically associated with inability to report freely how the task was performed, because subjects would not know what questions to ask themselves to elicit their knowledge. These characteristics might correspond to an important aspect of what the layperson means by unconscious learning.

Both the guessing and zero-correlation criteria tap only one of two possible forms of metaknowledge. Dienes and Perner (1996) argued that full metaknowledge requires both representing oneself as being in the possession of certain propositional content (they called this type of representation *content explicit*) and representing an appropriate propositional attitude toward that content—namely, representing the content as knowledge and not, for example, just as guessing (they called the second type of representation *attitude explicit*). Knowledge could be unconscious according to the metaknowledge criterion either because content explicitness was lacking or because attitude explicitness was lacking. Finding that confidence is unrelated to accuracy, or that subjects are accurate when they believe that they are guessing, shows that subjects' representations lack attitude explicitness. But they may have content explicitness, as in the possible case of accurate verbal reports offered as guesses. Reber (1993, p. 136) suggested that subjects in artificial grammar learning experiments can show the converse; that is, they may know that they know something, even though they may not know what it was that they know. Similarly, in natural language, we may be very confident in our judgments of grammaticality, and so the judgments per se would be attitude explicit, but we may be unable to describe the bases of those judgments, and so the bases would be content implicit. Both types of lack of metaknowledge—content implicitness and attitude implicitness—are possible and even plausible in different circumstances. We will be reviewing mainly whether subjects' representations can lack attitude explicitness.

FEATURES OF IMPLICIT LEARNING

Theoretically, it can be argued that the finding of knowledge below a certain threshold is, by itself, of little value unless knowledge below and above the threshold is qualitatively different. Merikle (1992), for example, summarized evidence that the subjective threshold is a psychologically real one in the case of subliminal perception, because perception above and below the threshold is qualitatively different. We argue that implicit learning and the resulting knowledge are distinguished from explicit learning and knowledge by the following features: (1) limited transfer of the knowledge to related tasks; (2) learning tends to be associated with a focus on particular items rather than on the underlying rules; and (3) robustness of learning.

Limited Transfer of the Knowledge to Related Tasks

There is some evidence that there is only limited transfer to tasks that are based on the same underlying structure. Some studies show no transfer of learning at all, whereas others show reduced transfer. Only a few studies show equivalent performance levels on the structurally similar transfer task.

Learning Tends to be Associated With a Focus on Particular Items Rather Than on the Underlying Rules

Implicit learning tends to be associated with observation and memorization conditions, rather than with deliberate hypothesis testing. Many studies have shown that people who approach complex tasks in a relatively passive manner perform at least as well (and sometimes better) than do people who try to work out the underlying structure of the task explicitly.

Robustness of Learning and Knowledge

It has been suggested that implicit learning should be more robust than explicit learning (Reber, 1989). Reber argued that the robustness of the implicit over the explicit comes from the relative recency, in evolutionary terms, of the conscious, relative to the unconscious. Allen and Reber (1980) argued that implicit knowledge was robust in the sense that it is more durable than explicit knowledge. They found that subjects performed above chance on a classification task 2 years after an initial 10- to 15-min exposure to an artificial grammar. This aspect of robustness has not been explored in other studies, and so, for the purposes of this review, we consider three other aspects to robustness: (1) psychological and organic disorder; (2) other individual difference variables; and (3) secondary tasks.

Psychological and organic disorder. There have also been claims for robustness in the face of psychological and organic disorder. That is, implicitly acquired knowledge remains more intact than does explicitly acquired knowledge. This finding fits in with the growing literature on implicit–explicit distinctions in neuropsy-

chological patients (Schacter, McAndrews, & Moscovitch, 1988).

Other individual difference variables. Reber (1989) argued that implicit knowledge should also be less affected by other individual difference variables such as age or IQ.

Secondary tasks. Finally, it has been suggested that the storage and retrieval of implicit knowledge should be less affected by secondary tasks than should explicit knowledge.

The review is structured by considering, first, the sense in which the knowledge might be implicit and, second, evidence for the distinguishing features of implicit learning for each of the artificial grammar learning, the control of complex systems, and sequence learning paradigms. Issues of the nature of the underlying representation and how implicit learning may be best modeled are not discussed in detail in this paper (see Neal & Hesketh, 1997, and Berry & Dienes, 1993, for a consideration of these issues).

ARTIFICIAL GRAMMAR LEARNING

Reber has long been exploring the learning of finite-state grammars, which he claims occurs in some implicit way (see Reber, 1989, for an overview). In a typical study, subjects first memorize grammatical strings of letters generated by a finite-state grammar. Then, they are informed of the existence of the complex set of rules that constrains letter order (but not what they are) and are asked to classify grammatical and nongrammatical strings. In an initial study, Reber (1967) found that the more strings that subjects had attempted to memorize, the easier it was to memorize novel grammatical strings, indicating that they had learned to utilize the structure of the grammar. Subjects could also classify novel strings significantly above chance (69%, where chance is 50%). This basic finding has now been replicated many times, as will be seen. So subjects clearly acquire some knowledge of the grammar under these incidental learning conditions. But is this knowledge implicit?

The Accessibility of Classification Knowledge

Inaccessibility with free report. Reber has often asked for introspections regarding classification (e.g., Abrams & Reber, 1988; Allen & Reber, 1980; Reber & Allen, 1978; Reber, Kassin, Lewis, & Cantor, 1980; Reber & Lewis, 1977). Reber (e.g., Reber, 1989) concluded that although subjects emerged with a small but solid body of articulated knowledge, they still could not tell all that they knew. Although this claim was not formally tested by Reber, a systematic procedure to investigate the validity of the subjects' freely stated rules was used by Mathews et al. (1989). Experimental subjects first briefly studied a set of letter strings (the training phase) and then classified new letter strings as grammatical or not for 800 trials with feedback (test phase). After each 10-trial block, subjects were asked to give complete instructions on how to classify (the free-report measure of their

knowledge). The validity of these instructions was assessed by the classification performance of a group of yoked subjects who were requested to follow the transcribed instructions but were not given a training phase. That is, the prime source of information for the yoked subjects was the instructions given by the experimental subjects. The yoked subjects performed significantly (and substantially) above chance, and their performance gradually increased across blocks in the same manner as did that of experimental subjects. Thus, at least some of the knowledge used for classification could also be given in free report. However, the yoked subjects always performed significantly worse than did the experimental subjects, suggesting that the experimental subjects could not access their knowledge base equally by the classification and free-report measures.

Dienes, Broadbent, and Berry (1991) found similar results with a technique that did not depend on the possible idiosyncrasies of yoked subjects. After 10-min exposure to grammatical strings, subjects made 100 classification decisions and then attempted to describe as fully as possible the rules or strategies they used to classify. When the rules elicited in free report were directly used to simulate classification performance, the simulated performance (53%) was still considerably less than the actual classification performance (63%).

To summarize, the evidence indicates that the knowledge elicited from subjects in free report is impoverished, relative to the knowledge they can actually apply to classify test strings.

Objective threshold. Dulany et al. (1984) employed a forced-choice measure that did elicit the knowledge underlying classification performance. They asked subjects during classification to score that part of a string that "made it right" if it was classified as grammatical or that part that violated the rules if it was classified as nongrammatical. For each feature thus scored by each subject, a validity was calculated over all test strings: $P(\text{string is in the correct category}/\text{Feature } i \text{ is in the string})$. When the proportion of correct classifications for each subject was regressed onto the mean validity of all the features scored by each subject, the regression slope was not significantly different from 1.00, and the intercept was not significantly different from 0.00. This provides strong evidence that the scoring and classification tasks tapped the same data base with about the same sensitivity. Another similar task to which the classification knowledge appears to transfer, at least partially, is the filling in of one or two blank spaces of otherwise grammatical strings (McAndrews & Moscovitch, 1985).

Perruchet and Pacteau (1990) investigated the hypothesis that classification knowledge consisted of explicitly knowing permissible bigrams. After training, when subjects were asked to rate isolated bigrams for their legitimacy, their ratings could predict classification performance. Thus, subjects could apply their knowledge in cued-report tests about isolated bigrams. However, subjects also know more than just bigrams. For example, Gomez and Schvaneveldt (1994) and Manza and Reber

(in press) found that subjects trained by being exposed to bigrams could not transfer their knowledge to classifying in a new perceptual domain, but subjects trained by exposure to whole letter strings could transfer their knowledge. Also, Dulany et al.'s (1984) scoring task predicted classification performance without significant error only when the feature scored was taken to include position. Furthermore, Dienes et al. (1991) tested subjects' knowledge of the positional dependence of bigrams in incomplete strings. After 10 min of exposure to grammatical strings, subjects were asked which letters could occur after different stems, varying in length from zero letters upward (the sequential letter dependency, or SLD, task). Considering questions concerning permissible bigrams, subjects only responded "grammatical" at a level above chance if the bigram was in a grammatical position. That is, subjects' knowledge of the positional dependence of bigrams could be elicited by cued-report tests out of the context of particular exemplars.

Chan (1992) argued that the knowledge base underlying performance on the artificial grammar learning task might be more purely implicit if graphics symbols were used instead of letters to construct exemplars. Chan chose graphics symbols that pilot subjects found highly non-verbalizable and indeed found evidence of implicit knowledge when he used these stimuli. He gave subjects the standard training procedure and classification task, as well as bigram and SLD tests. In contrast to the findings of Perruchet and Pacteau (1990) and Dienes et al. (1991), Chan (1992) found dissociations between the tests: The bigram and SLD tests revealed knowledge only of the symbols at the beginnings and ends of sequences; classification performance, on the other hand, was sensitive to symbols in the middle of strings as well as those at the end. With these stimuli, the subjects' full knowledge of the grammar does appear to require a whole exemplar to elicit it rather than isolated parts of exemplars. However, as Dulany (1962) and Shanks and St. John (1994) forcefully argue, the failure of a cued-report (or other) test to elicit knowledge only indicates unconscious knowledge if it is testing for the same knowledge that underlies classification performance. The question of what subjects have learned in order to classify these stimuli still needs addressing: Subjects may have not learned bigrams at all, but properties of the whole exemplar such as its symmetries.

In summary, a number of cued-report tests do elicit subjects' knowledge of the artificial grammar, at least when stimuli are constructed with letters. That is, there is no convincing evidence that knowledge of artificial grammars can occur below an objective threshold.

Subjective threshold. Although subjects can produce the right answer to a classification or other test, they might not know that they can. In this case, subjects may lack metaknowledge. The level of subjects' metaknowledge in artificial grammar learning was first investigated by Chan (1992). Chan trained some subjects on graphics symbols and some on letters; both groups of subjects were then given a classification test with the symbol set that they were trained on. After each classification decision, Chan

asked subjects to rate how confident they were that the decision was correct. He found that, when subjects were just asked to memorize exemplars in the training phase, they were just as confident in incorrect decisions as in correct decisions, and this was so despite the fact that, in general, subjects acquire almost completely veridical knowledge about the grammar (Reber, 1989). The correlation between confidence and accuracy was .20 (Chan, 1992, Experiments 6 and 7). When subjects were asked to look for rules in the training phase, the overall level of accuracy was the same, and the variances for accuracy and confidence were the same. However, they were now more confident in correct decisions than in incorrect decisions, and the correlation between confidence and accuracy was .54 (significantly greater than the value for memorization subjects). That is, memorization subjects lacked metaknowledge about their knowledge, but rule-search subjects did not. Chan (1992) also found that the magnitude of the correlation depended on the type of stimuli used. For example, if memorization subjects were asked to rate bigrams for grammaticality, then confidence was correlated with accuracy ($r = .47$). Chan concluded that subjects acquired explicit knowledge of bigrams.

Manza and Reber (in press), using stimuli different from Chan's, found that confidence was reliably higher for correct decisions than for incorrect decisions. On the other hand, Dienes et al. (1995) replicated the lack of relationship between confidence and accuracy for letter stimuli, but only under some conditions: Metaknowledge was low particularly when strings were longer than three letters and presented individually. It may be difficult to find conditions under which Chan's measure of metaknowledge is literally zero; however, as long as it varies in meaningful ways (when classification performance is constant), it may provide a useful measure of degree of metaknowledge and, hence, degree of explicitness.

Dienes et al. (1995) also found that subjects could lack metaknowledge according to the criterion of Cheesman and Merikle (1984): Even when subjects believed that they were literally guessing, they were still classifying substantially above chance. On the basis of this criterion, about one third of subjects' responses could be regarded as being based on purely implicit knowledge.

Finally, Dienes and Altmann (in press) found that, when subjects transferred their knowledge to a different domain, they were classifying substantially above chance even when they believed that they were literally guessing, and their confidence was not related to their accuracy.

In summary, when trained by simply observing or memorizing exemplars, subjects have low, even if sometimes statistically significant, levels of metaknowledge about their knowledge of the artificial grammar. The knowledge appears to be below a subjective threshold.

Features of Implicit Learning

Transfer across domains. Reber (1969) asked subjects to memorize strings of letters generated by a finite-state grammar. The more strings that subjects had previously studied, the easier they found it to memorize new strings

generated by the grammar. There was also a benefit when the new strings were constructed from a different letter set, but the same grammar. That is, subjects could apply their knowledge of the grammar to a completely different letter set, leading Reber to claim that the knowledge was abstract and not perceptually bound. Subsequent work by Mathews et al. (1989), Brooks and Vokey (1991), Whittlesea and Dorken (1993), Gomez and Schvaneveldt (1994), and Manza and Reber (in press) exposed subjects to strings constructed from one letter set and later tested subjects' ability to classify strings constructed from another letter set. As well as showing transfer performance significantly above chance, these studies have all demonstrated a transfer decrement—that is, an advantage of same-domain performance over transfer performance. Altmann, Dienes, and Goode (1995) extended these findings to transfer between different modalities. In each of four experiments, involving transfer between letters and music or between graphic symbols and nonsense syllables, significant transfer occurred across domains. Furthermore, subjects performed better when tested in the same domain in which they had been trained than when tested in the different domain.

In these studies, subjects were not told of the mapping between the domains, and they were not given accuracy feedback. So it is not clear whether the transfer decrement (i.e., the difference between different- and same-domain performance) reflects the problem of inducing a mapping across domains or reflects the possibility that at least some of the knowledge is bound to particular perceptual features. There are a number of other studies that bear directly on this question. Whittlesea and Dorken (1993) trained subjects on exemplars from two grammars. The exemplars were distinguished by the task that subjects performed on them: For one grammar, subjects pronounced the exemplars; for the other grammar, subjects spelled them. Subjects were later asked to pronounce or spell test items and then classify these strings as belonging to either the “pronouncing” grammar or the “spelling” grammar. Whittlesea and Dorken argued that discrimination between the grammars was achieved simply by feelings of familiarity induced by test experiences perceptually matching or mismatching representations of prior training experiences: When test strings were common to both grammars, subjects tended to classify them as belonging to the spelling grammar if they spelled them and as belonging to the pronouncing grammar if they pronounced them.

Manza and Reber (in press) found no significant transfer decrement between sequences of tones and lights when subjects were informed of the mapping. However, the same-domain performance was 58% and the cross-domain performance was only 54%, consistent with a substantial decrement [the upper limit of the 95% confidence interval on the decrement must be at least $2(58\% - 54\%) = 8\%$]. Wright (1993), using number stimuli, found significant decrements when format of the stimuli was changed from digits (e.g., 4836) during training to words (e.g., four–eight–three–six) at test. Finally, Dienes and

Altmann (in press) found a significant decrement when subjects were transferred between the domains of colors and the names of the same colors. Dienes and Altmann pointed out that the size of the decrement was of the same order as that found in the implicit memory literature when priming is compared across different modalities (e.g., see Roediger & Blaxton, 1987).

In summary, despite the surprising ability of subjects to transfer across domains, the knowledge is partly perceptually bound and transfer is not normally complete even when a simple mapping is known. Manza and Reber (in press) found that when subjects were trained on two domains rather than one, same-domain and transfer performance were more closely equalized. Future research could investigate whether the extent to which knowledge of artificial grammars is perceptually bound depends on the training regime.

Learning tends to be associated with a focus on particular items rather than on the underlying rules. In a typical artificial grammar learning experiment, subjects are simply asked to observe or memorize the strings of letters; they are not informed of the existence of a set of rules. Reber (1976) showed convincingly that, for some stimuli, subjects do not learn by actively searching for rules. He instructed one group of subjects to memorize a set of randomly ordered grammatical strings (memorization instructions) and instructed another group to search for rules to assist their memorization (rule-search instructions). The memorization subjects subsequently classified more test strings correctly than did the rule-search subjects. Conversely, Reber et al. (1980) showed an advantage of rule-search instructions over memorization instructions when the learning strings were presented in an order that highlighted the grammatical rules. The effect seems to depend on the exact grammar used; in many cases, when the learning strings are presented in a random order, subsequent classification performance is equally good after memorization instructions or rule-search instructions (e.g., Dienes et al., 1991; Durlany et al., 1984; Mathews et al., 1989, Experiments 1 and 2). Furthermore, Dienes et al. found no interaction of memorization instructions versus rule-search instructions with the effect of a secondary task. Mathews et al. (1989, Experiment 3) used a much stronger manipulation of how actively subjects might be induced to search for rules and still found no difference: For the passive task (the match task), subjects held a string in memory for a few seconds and then had to recognize it among one of several choices. For the active task (the edit task), subjects were informed of the presence of rules and were asked to underline possible incorrect letters in invalid strings. After each invalid string, the correct string was presented. Classification performance was equally good after the match task or edit task alone or in combination.

Mathews et al. (1989, Experiment 3) also compared normal artificial grammar learning with learning a “biconditional” rule that determined the mapping between corresponding letters in the first and second halves of a string. The first four letters of any string were taken ran-

domly from the set T, P, V. The next four letters were chosen such that T predicted an X in a corresponding position, P predicted a C, and V predicted an S. Thus, for example, TPPV.XCCS was a valid string. As described above, for the normal artificial grammar, the match and edit tasks produced identical classification performance. With the biconditional rule, on the other hand, classification performance was better after the edit task than after the match task; performance on the match task was next to chance. Performance was best of all for a group of subjects who received first the match task and then the edit task. Matthews et al. hypothesized that only the finite-state grammar involved a family-resemblance structure that could be learned purely implicitly, and they hypothesized that some more explicit tasks benefit from a synergy between implicit and explicit learning modes. In general, if stimuli have a structure suited for the implicit learning mechanism, memorizing randomly ordered grammatical learning strings leads to performance that is as good as, or better than, the performance resulting from searching for rules.

Robustness: Psychological and organic disorder. Knowlton, Ramus, and Squire (1992) found that, on the artificial grammar learning task, amnesics classified new exemplars at the same level as did control subjects (63% and 67%, respectively), but, on a recognition test of the exemplars that had been presented, they performed more poorly than did control subjects (62% and 72%, respectively), where chance is 50% in both cases. Presumably, normal controls were using some explicit memory of the exemplars to perform the recognition task but not the classification task. Knowlton and Squire (1994) replicated the normal classification performance of amnesics and found that amnesics were similar to normals in terms of which items they found easy and which items they found difficult to classify.

Abrams and Reber (1988) found that, after exposure to grammatical strings, psychiatric patients and normals classified grammatical and nongrammatical strings similarly; however, on a task that required determining a mapping between letters and numbers and that was regarded as requiring explicit knowledge, psychiatric patients were inferior to normals. In one failure to demonstrate robustness of artificial grammar learning, Reber (1967) found that undergraduates performed significantly better than did high school students.

Robustness: Other individual difference variables. Reber, Walkenfeld, and Hernstadt (1991) used the evolutionary argument to predict that implicit learning should show less between-subject variability and operate largely independently of IQ. They compared artificial grammar learning with a series-completion problem-solving task (an example question would be “Should a D or C come next in the sequence ABCBCDCDE_?”) of the same difficulty. The results showed that artificial grammar learning had a reliably smaller variance and smaller correlation with IQ than did the problem-solving task.

Robustness: Secondary tasks. Dienes et al. (1991), using the artificial grammar learning paradigm, investi-

gated the effect of a secondary task (random number generation [RNG]) on both classification performance and free report of the rules of the grammar. In the initial learning phase, subjects were asked to memorize grammatical strings of letters. At the same time, some subjects spoke a random digit every 2 sec. A control group of subjects tried to memorize the exemplars without performing the RNG task. All subjects then classified new exemplars without performing a secondary task. The dual-task manipulation interfered with classification performance. Analyses revealed that the disruption was specifically in subjects' ability to establish the positional dependence of bigrams—a finding that will emerge again when considering the SRT task below.

Dienes et al. (1995) employed the RNG task during the test phase of an artificial grammar learning task for some subjects; other subjects classified without generating random numbers. All subjects also gave confidence ratings during the test phase. Overall, the secondary task interfered with classification performance. Analyses indicated that this only occurred for cases in which subjects had some confidence in their responses. When subjects believed that they were literally guessing, the secondary task did not interfere. Thus, in cases where subjects had the least metaknowledge, their performance was most resistant to the effects of a secondary task.

In summary, incidentally acquired knowledge of artificial grammars is surprisingly robust to the effects of psychological disorder and age. Secondary tasks interfere with learning (but perhaps only with learning *n*-grams of higher order than bigrams) and with the application of the knowledge (but perhaps only of that part of the subjects' knowledge about which they have metaknowledge).

Conclusion

The knowledge underlying classification performance appears to be below a subjective threshold: Subjects classify above chance when they believe that they are guessing, and confidence is often unrelated to accuracy. Is this a psychologically interesting threshold? In terms of features that may distinguish implicit knowledge from explicit knowledge, learning is more likely to result in knowledge below rather than above the subjective threshold when subjects focus on items rather than underlying rules (Chan, 1992); the application of knowledge below rather than above the subjective threshold is resistant to the effects of a secondary task (Dienes et al., 1995). Future research needs to establish whether knowledge below the subjective threshold is more perceptually bound or more robust in terms of psychological disorder or other individual difference variables, but the existing results are promising: Subjects' knowledge of finite-state grammars, taken as a whole, is more perceptually bound and robust than is knowledge that is plausibly largely above subjective threshold.

THE CONTROL OF COMPLEX SYSTEMS

A substantial body of evidence for the existence of distinct implicit and explicit learning modes comes from

a series of studies on the “dynamic control tasks.” The dynamic control tasks involve subjects’ attempting to control the level of one or more output variables by deciding on the level of one or more input variables. Subjects have continuous feedback on the output variables while they manipulate the input variables (e.g., Berry & Broadbent, 1984, 1987, 1988; Broadbent, 1977; Broadbent & Aston, 1978; Broadbent, FitzGerald, & Broadbent, 1986). For example, Berry and Broadbent (1984) devised and used what are now known as the *sugar-production* and *person-interaction* tasks. In the case of the sugar-production task, subjects took on the role of manager of a simple sugar-production factory and were required to reach and maintain specified levels of sugar output by varying the number of workers employed. Each time the subject entered the number of workers employed, the sugar output would be updated, and the subject would then enter a new value for the number of workers. In the case of the person-interaction task, subjects were required to interact with a “computer person.” The communication was based on a fixed set of adjectives describing various degrees of intimacy of personal interaction. Subjects were told to adjust the behavior of the person to the “very friendly” level and to keep it there. The equation relating sugar output to work force was the same as that relating the computer person’s responses to those of the subject. There was not a unique output associated with any one input. In the case of sugar production, for example, the resulting sugar output depended on the previous sugar output as well as on the new work force figure. This made it more likely that subjects would exercise continuous control, rather than being able to hit target by chance and then remain there simply by typing in the same input value.

Berry and Broadbent (1987, 1988) suggested that nonsalient relationships might be learned implicitly and that salient relationships might be learned explicitly (cf. Reber et al., 1980). In their 1988 study, Berry and Broadbent devised salient and nonsalient versions of the person-interaction task. In both cases, the output variable was simply a constant added to the input variable (plus a small amount of noise). In the salient case, the value of the input variable used was the one just entered by the subject; in the nonsalient case, the value was the one entered on the previous trial (i.e., there was a lag).

We now consider how accessible the knowledge is that subjects acquire about controlling dynamic systems.

Accessibility of the Knowledge

Free report. Berry and Broadbent (1984) found that, with practice, subjects improved in their ability to control the sugar-production and person-interaction tasks. However, retrospectively asking subjects to write down how they went about attempting to reach and maintain target values elicited uninformative answers. Many subjects did not answer the question at all, protesting that they were “unable to put it into words.” Others wrote general statements, such as “increasing and decreasing the work force to change the sugar level.”

Stanley et al. (1989) investigated the knowledge elicited in free report in a more systematic way. They also used the sugar-production and person-interaction tasks. They asked subjects after every 10-trial block to give complete instructions on how to perform the task. The informativeness of these instructions was assessed by the performance of yoked subjects requested to follow the transcribed instructions. Stanley et al. demonstrated that sudden improvements in performance by the original learners were not associated with simultaneous increases in the informativeness of the instructions. In fact, instructions helped the performance of yoked subjects only if the instructions were taken about four blocks after the improvement in performance. Therefore, subjects seemed to operate on acquired knowledge well before they could explain the knowledge verbally.

McGeorge and Burton (1989) also measured the validity of instructions, but, this time, by computer simulations of performance based on the instructions. Where possible, heuristics elicited from subjects were translated directly into computational procedures. Using this method, McGeorge and Burton found that only about one third of their subjects reported heuristics that were at least as good as observed performance.

In summary, the evidence suggests that subjects do not freely report sufficient knowledge to account for their actual control performance. Good free-report knowledge comes only after control performance has improved.

Objective threshold. In the Berry and Broadbent (1984) study, subjects were required to complete written posttask questionnaires that asked about the relationships within the system. Specifically, the questions asked subjects to predict what the next value of the output would be (e.g., sugar production) given the previous values of the output and input variables. It was found that, although practice significantly improved ability to control the tasks, it had no effect on ability to answer the questionnaire. In contrast, detailed verbal instructions about the principles of the system improved subjects’ ability to answer questions but had no effect on control performance. Berry and Broadbent (1987, 1988) showed that subjects were able to answer questions about salient relationships correctly; only questions about nonsalient relationships caused difficulty.

The above questionnaire results appear to suggest that knowledge elicited by the continuous demands of controlling the system is not elicited by most cued-report tests. However, if control performance does not require the ability to predict the consequences of changes in the input variables (except insofar as they bring one closer to target), the lack of relationship between control and questionnaire performance does not imply that subjects unconsciously know how the system works or can unconsciously predict the consequences of changes in the input variables (Sanderson, 1989; Shanks & St. John, 1994).¹ For example, Broadbent et al. (1986) suggested that subjects could be learning (nonsalient) tasks by forming a lookup table—that is, by linking specific responses to specific situations as they experience them. Forming an accurate lookup

table does not require understanding any of the general properties of the system.

Three studies have tested how subjects performed on specific situations after an initial training phase. Berry (1984, 1991) worded questions such that subjects were given a series of values for all variables and were asked which value of the input variable should be entered next to bring the output variable to target. Although this is what is required while controlling the task, performance was at chance on these “estimated-input questions.” Two further studies, by Marescaux, Luc, and Karnas (1989) and Dienes and Fahey (1995), showed why there might be this discrepancy between task performance and estimated input questions. Marescaux et al. (1989), using the sugar-production task, demonstrated that subjects know more about situations that they have personally experienced, as would be predicted by a lookup-table approach. Subjects were trained on the sugar-production task, manipulating level of work force to control level of sugar production, for two sets of 30 trials. Subjects then answered a number of questions about specific situations. The questions were matched closely to the task, in that they required subjects to estimate how many workers would be needed to reach target sugar production in different situations. An example situation was, “If you had just employed 400 workers and if the sugar production was then 8,000 tons, what should you do next to bring the sugar production to target?” The questions varied along two basic dimensions. First, the target sugar output was either the same as or different from that experienced while controlling the task. Second, the minihistory given to subjects either was taken from their own interaction in the training phase or was randomly constructed. The results showed superior questionnaire performance when subjects had to reach the same target as they had experienced while interacting with the task and when the minihistories were taken from their own past experience. The subjects did not perform well when faced with randomly constructed situations—a result that replicates Berry’s (1984, 1991) findings with the estimated-input questions. Also consistent with the idea that subjects formed a lookup table, Marescaux et al. found that, when subjects were shown situations in which they had previously been correct, the responses given by subjects matched their previous responses 57% of the time.

Dienes and Fahey (1995) also tested subjects on specific situations after an initial training phase. Like Marescaux et al. (1989), they found that subjects performed best on old situations in which they had previously been successful (they were still at chance on new situations), and they found that subjects were consistent in their responding to a situation in which they had previously been correct. This consistency was greater than the consistency for incorrect situations. Dienes and Fahey argued that subjects had learned only specific responses for situations in which they had previously been correct and subjects had not learned any general rules. This may explain why previous studies had found a dissociation between performance and questions about the principles of the system.

Dienes and Fahey (1995) examined performance on their situations task following experience of interacting with the salient- or nonsalient-person control task. They found that, following interaction with the nonsalient task, subjects’ knowledge was limited to situations that they had previously seen. However, in the case of the salient task, subjects were able to respond correctly to new situations. Dienes and Fahey suggested that subjects might learn the salient task partly by a lookup table and partly by acquiring general knowledge that they could apply to any situation.

Buchner, Funke, and Berry (1995) provided a similar analysis of why control performance is not positively related to ability to predict the consequences of changes in the input variables, as measured by a posttask questionnaire. Most subjects with many trials on target (good control performance) experience only a small number of specific situations. In contrast, subjects who are poor at controlling the system experience a large number of different types of specific situations. The greater the number of different types of specific situations that subjects experience, the greater the probability that subjects can answer the questionnaire with reference to their specific experience. In support of this analysis, Buchner et al. found that the number of different situations experienced while controlling the system was negatively correlated with number of trials on target and was positively correlated with questionnaire performance.

In summary, when subjects are given cued-report tests on the knowledge actually used to control the system, they perform well. People initially learn nonsalient dynamic control tasks by memory for specific situations, but people tend to develop knowledge of how the system works as a result of task experience. When dealing with salient systems, people quickly acquire knowledge of how the system works, as revealed by cued-report tests. That is, the knowledge of both nonsalient and salient systems is not below an objective threshold.

Subjective threshold. The subjective threshold has not been directly tested using the dynamic control tasks. However, many subjects report that they relied on intuition or mainly guessed in making decisions. Research is needed on this point.

Features of Implicit Learning

Transfer across domains. Berry and Broadbent (1988) looked at transfer between control tasks that were either perceptually similar or perceptually dissimilar. There were four tasks in all: two different person-interaction tasks and two transport tasks (one bus and one train). In each case, the underlying equation (a nonsalient relationship) was the same. Berry and Broadbent found positive transfer between performance on the two person tasks and between performance on the two transport tasks. There was no transfer across two dissimilar tasks (person followed by transport, or vice versa). Furthermore, transfer was impeded between the two similar tasks when subjects were informed of the critical relationship between them. Berry and Broadbent suggested that informing

subjects of the connection prior to controlling the second task induced an explicit mode of learning that had a detrimental effect on subsequent performance. This negative effect of a “transfer hint” goes against a series of experiments in the problem-solving literature (e.g., Gick & Holyoak, 1980). In the latter case, however, the critical information necessary for task solution was generally explicitly available to subjects. Once the necessary connection had been pointed out to them, subjects could draw on this critical information and consciously apply it to the problem requiring solution. This was not the case in the Berry and Broadbent study. Squire and Frambach (1990) also reported no transfer between performance on sugar-production and person-interaction tasks. Transfer has not yet been examined using salient versions of the control tasks, but one would predict, in this case, transfer across dissimilar tasks.

A further demonstration of transfer specificity comes from a study by Berry (1991). An initial experiment showed that experience of watching another person controlling the sugar-production and person-interaction tasks had no effect on subsequent control performance. Berry went on to use a transfer paradigm to investigate what aspect of personal control leads to learning—whether it is the decision-making component or the physical interaction component that is important. In one condition, subjects made decisions for the experimenter to enter into the computer. In a second condition, the experimenter made the decisions and the subjects typed them in. In both cases, subjects then controlled the entire task themselves. It was found that subjects in neither condition performed well on the second set of trials, relative to a standard group who carried out both aspects of the task on both the first and the second set of trials.

An interesting finding was that subjects who made decisions for the experimenter to type in on the first set of trials performed as well on that set as did subjects who carried out the entire task themselves. Despite this, they still performed poorly on the second set of trials when they had to carry out the entire task. One suggestion was that subjects did not transfer knowledge acquired during the initial set of trials to the second set because of a mismatch between the conditions of learning and testing (see, e.g., Kolers & Roediger, 1984). Berry therefore carried out a further experiment in which subjects made decisions in the absence of action (i.e., someone else typed them in) for both sets of trials, made decisions alone on Set 1 and carried out the entire task on Set 2, carried out the entire task on Set 1 and made decisions alone on Set 2, or carried out the entire task for both sets of trials. It was found that subjects who made decisions alone for both sets showed a small improvement across the sets. This finding supports the type of specificity of transfer approach advocated by Kolers and Roediger. However, Berry also found that subjects who controlled the entire task on Set 1 also performed well on Set 2 when they had to make decisions alone. It may be that subjects performing the entire task always made a subvocal verbal response as well as a typing

response, but subjects making decisions alone made a verbal response but not an abbreviated typing response.

In contrast to these results for nonsalient tasks, Berry (1991) found that experience of watching another person interacting with the salient-person control task had a beneficial effect on subsequent control performance with the same task. Salient tasks do not show the specificity of transfer demonstrated by nonsalient tasks.

Dienes and Fahey (1995) found positive transfer between visual and auditory versions of the person-interaction task. Subjects who interacted with the task in one modality were able to answer questions about previously experienced situations just as successfully when presented in a different modality as when presented in the same modality. Again, perhaps a subvocal verbal response mediated transfer.

In summary, transfer is surprisingly limited in controlling nonsalient systems. When the system obeys the same rules, but the cover task has been changed (Berry & Broadbent, 1988) or the responses are typed instead of spoken (Berry, 1991), the appropriate knowledge is not elicited.

Learning tends to be associated with a focus on particular items rather than on the underlying rules.

The control tasks are different from the artificial grammar learning task, in that controlling the system is the subjects’ primary task. However, the evidence suggests that people can go about this primary task either in a relatively passive way, storing the particular situations they have come across, or in a more explicit hypothesis-testing way, looking for general rules. The effectiveness of the latter mode of learning depends on the nature of the underlying relationship (Berry & Broadbent, 1988), with a more passive mode being better for a nonsalient relationship. Berry and Broadbent looked at the differential effects on performance on the salient- and nonsalient-person tasks of presenting subjects with an explicit search instruction (see also Reber, 1976). Following the first set of trials, half of the subjects were told that the computer person’s responses were determined by their responses and it might help them control the person better if they tried to work out exactly how the computer’s responses were related to their own. Subjects who received the search instruction performed better on the salient task than did subjects who did not receive it. In contrast, the instruction had a detrimental effect on those interacting with the nonsalient person. Similarly, as described in the previous section, Berry and Broadbent (1988) also found that transfer was impeded if subjects were presented with two similar nonsalient tasks but were informed of the critical relationship between them.

Robustness: Psychological and organic impairment.

Squire and Frambach (1990) examined whether amnesics could learn to control the sugar-production task at the same level as do normal subjects. They found that the amnesic patients performed just as well as the normal controls in an initial training session, but they performed significantly worse in a second session. Squire and Frambach suggested that this was because, by this stage of

practice, the normal subjects were starting to build up explicit knowledge, which could be used to improve performance still further. The amnesics, in contrast, were not able to do this.

Robustness: Other individual difference variables. Myers and Connor (1992) showed that older subjects (30–59 years of age) were able to control the sugar-production task as well as did younger subjects (16–19 years of age), but they performed less well on the standard questionnaire asking subjects to predict the consequence of changes in the input variables.

Robustness: Secondary tasks. Hayes and Broadbent (1988) argued that explicit learning rather than implicit learning would be more affected by a secondary task. They found that performance on the nonsalient-person task was disturbed by an unexpected reversal of the relationships within the task (output went from being about two steps less than input to being about two steps greater than input), whereas performance on the salient task was not. They argued that, because subjects had an explicit understanding of the rule in the salient case alone, it was easy to modify the subjects' representation of the rule in this case alone. They also showed that when a concurrent memory-demanding task was performed, subjects learned the reversed relationship more easily for the nonsalient version than for the salient version. They argued that this was because the concurrent task interfered with the process of consciously reflecting on the rule and changing it. Unfortunately, these results have not been subsequently replicated. Green and Shanks (1993) found that performance on the salient task, relative to performance on the nonsalient task, was superior after an unexpected reversal, consistent with the results of Hayes and Broadbent. However, with a secondary task, they found that performance on the salient task, relative to performance on the nonsalient task, was still superior. Green and Shanks argued that the results were consistent with the claim that the nonsalient task was simply more difficult than the salient task. Sanderson (1990) also found that performance on the salient task, relative to performance on the nonsalient task, was (nonsignificantly) superior after reversal with the concurrent task. (See Berry & Broadbent, 1995, for a more detailed discussion of the effects of secondary tasks on control performance.)

In summary, secondary tasks appear to interfere with both salient and nonsalient tasks. Of course, there is no reason to presume that, in the absence of explicit learning, there is no implicit learning of the salient task. Even an implicit learning mechanism may find the salient task easier than the nonsalient task (as is true of the computational model of Dienes & Fahey, 1995), and this might be complicating the above experimental findings. Future research could usefully determine whether a secondary task affects rule learning more than it affects lookup-table-based learning (using the methodology of Dienes & Fahey, 1995) or whether it affects confident responses more than it affects guessing responses (as in the case of artificial grammars; Dienes et al., 1995).

Conclusion

In terms of the distinction between subjective and objective thresholds, the evidence suggests that the knowledge relevant to performance lies above an objective threshold. When relevant cued-report and forced-choice tasks were used, the knowledge was elicited (Berry & Broadbent, 1987; Dienes & Fahey, 1995; Marescaux et al., 1989). There is suggestive evidence that the knowledge may lie below a subjective threshold: Subjects could not report how they controlled the system; they often informally stated that they were just guessing (Berry & Broadbent, 1984; Stanley et al., 1989). Furthermore, the knowledge could not be used in a flexible way (Berry & Broadbent, 1988). Future research needs to more directly test the claim that a subjective threshold is relevant by asking for confidence ratings on the correctness of the control decisions (cf. Chan, 1992). Convincing evidence that the knowledge is resistant to secondary task still needs to be provided.

SEQUENCE LEARNING

It appears that people can become sensitive to sequential structure despite having little ability to articulate their knowledge; one everyday example is language. People's sensitivity to sequential constraints in language can be shown in three ways: First, people are faster to read a word when it occurs in an appropriate context than when it appears in an inappropriate context (Tyler & Marslen-Wilson, 1977). Second, people can choose or predict what word can occur next in a sequence in a constrained way (Miller & Selfridge, 1950). Third, people can, of course, freely generate sequences of words that follow the constraints of English. Despite these abilities, people find it difficult to say what the constraints are that guide their performance. A growing number of studies have tried to assess how and whether implicit knowledge of sequential structure can develop in the laboratory. Knowledge of the structure is shown by one or more of the ways listed above for showing knowledge of structure in natural language.

Work in this area has in fact involved two types of paradigm: (1) those paradigms in which the main measure of performance is the facilitated processing of sequences obeying rather than disobeying certain constraints (cf. Tyler & Marslen-Wilson's, 1977, study of natural language) and in which the aim is to see if this measure dissociates from other measures; (2) those paradigms in which the main measure of performance is the successful prediction of the next element (cf. Miller & Selfridge, 1950). Prominent examples of the first paradigm are the SRT task initiated by Nissen and Bullemer (1987), the Hebb (1961) digit task, and the matrix-scanning paradigm initiated by Lewicki and colleagues (e.g., Lewicki et al., 1987). Examples of the second paradigm are the probability-learning task of Reber and Millward (1968) and the sequence-prediction task of Kushner, Cleeremans, and Reber (1991). This review will largely focus on stud-

ies using the SRT task of Nissen and Bullemer and the matrix scanning task of Lewicki and colleagues.

In the Nissen and Bullemer (1987) task, a light appeared at one of four locations (arranged horizontally) on a video monitor. Subjects were required to press one key, out of four keys, that was directly below the position of the light. The sequence of lights either was determined randomly or appeared in a repeating 10-trial sequence. The results showed a rapid decrease in reaction time with training in the repeating-sequence condition, but not in the random condition. Furthermore, when subjects in the repeating condition were switched to a random sequence, reaction times increased substantially.

In the matrix-scanning task of Lewicki et al. (1987), subjects had to indicate, by pressing a button, which of four quadrants on a monitor contained a target digit. The task was structured in blocks of seven trials. During the first six trials of each block, only the target appeared on the screen (hence, the search task was simple). On the seventh trial, the target was embedded in a field of 35 distractors, making search much more difficult. The experiment was set up so that four of the six "simple" trials (first, third, fourth, and sixth) predicted the location of the target on the seventh ("complex") trial. On the complex trials, 24 possible target locations were used. Each was associated with a unique sequence of simple trials. (Hence, the task was based on 24 rules, where a rule consisted of a series of six locations, two of which were irrelevant.) In order to learn the relations between the simple and complex trials, subjects had to learn not only specific and relatively long sequences of target locations but also which particular trials should be attended to and which should be ignored. After extensive training, Lewicki et al. changed the relations between the simple and complex trials so that the target location on the complex trial was in the quadrant diagonally opposite its original location. The results showed that, prior to this change, performance was improving slowly but steadily. However, when the change was made, large negative transfer effects were observed.

Accessibility of Knowledge of Sequential Structure

Free report. Willingham, Nissen, and Bullemer (1989) trained subjects on the Nissen and Bullemer SRT task and then asked them if they had noticed a sequence or pattern and, if so, what it was. They found two subgroups of subjects: (1) the *unaware* subjects, who either claimed they had not noticed that there was a pattern or failed to specify more than three positions of the sequence correctly; and (2) the *aware* subjects, who claimed to have noticed a pattern and could reproduce the whole sequence. Both groups showed substantial procedural learning of the sequence; the reaction times of the unaware group decreased nearly 100 msec with training on the repeating sequence, and those of the aware group decreased about 200 msec. Although this result seems to suggest that unaware subjects had more knowledge than they could freely report, even knowledge of three elements could be sufficient to produce a substantial drop in reaction times.

Willingham, Greeley, and Bardone (1993) asked subjects, after performing the SRT task, whether the stimuli occurred completely randomly or with some degree of predictability. Nine out of 45 subjects said that the stimuli occurred randomly, and 11 subjects said that the next stimulus was only rarely predictable. Even the 9 subjects who claimed that the stimuli were completely random still showed reliable sequence learning.

There is another objection to concluding from these data that subjects can show sequence learning in the absence of ability to report their knowledge: Subjects may not even need to know the sequential dependencies between items in order to do well on the SRT task. Shanks et al. (1994) argued that subjects may have learned the probabilities of individual positions rather than the dependencies between successive positions. In previous studies, a comparison (either between or within subjects) has been made between trials in which subjects were exposed to the repeating sequence and trials in which the stimuli were chosen at random (with the constraint that stimuli could never repeat on consecutive trials). However, Shanks et al. pointed out that different stimuli have different probabilities of occurrence in the repeating sequence. They therefore introduced a pseudorandom series that matched the probabilities of occurrence in the repeating sequence. Following the procedure used by Willingham et al. (1989), Shanks et al. classified subjects in the repeating-sequence group as being aware or unaware on the basis of structured interviews. Shanks et al. found that performance of the unaware subjects was no better than that of subjects exposed to the pseudorandom sequence (although, as expected, it was better than that of subjects exposed to the truly random sequence). They suggested that most, if not all, of the knowledge in the unaware group was of the frequencies of different stimuli. Although, under the conditions of the Willingham et al. study, unaware subjects appear to learn only frequencies of individual stimuli, other paradigms have indicated that unaware subjects can be sensitive to higher order dependencies (e.g., Cleeremans & McClelland, 1991; McKelvie, 1987; Reed & Johnson, 1994), as we now discuss.

Reed and Johnson (1994) compared subjects' performance on a training sequence that differed from a transfer sequence only in terms of second-order conditional probabilities. Prediction of an element in the training sequence could be made only on the basis of the two preceding elements. Reed and Johnson found that after greater exposure than that given by Shanks et al. (136 repetitions of the sequence as opposed to 40), and with a simultaneous secondary tone-counting task (to reduce the amount of explicit sequence learning), subjects were faster on the training sequence than on the matched transfer sequence. That is, subjects must have been learning more than simple frequencies of elements; rather, they must have been learning the second-order structure of the sequence. Nonetheless, when asked, only 10 out of 60 subjects suggested that elements had appeared according to a fixed pattern (none of these could say what the pattern was).

Cleeremans and McClelland (1991) also increased the complexity of the stimuli in order to decrease the probability of the sequence being learned in an explicit way. Instead of using a fixed repeating sequence, they used a finite-state grammar to determine a structured but unrepeating sequence. The results showed that, during the course of 60,000 trials spread over 20 sessions, subjects became progressively sensitive to the sequential structure of the material. They responded significantly faster on the predictable trials than on the unpredictable trials. Detailed analyses showed that subjects were sensitive to up to three elements of context in predicting the next element. Analyses of verbal reports showed that subjects did become aware of some aspects of the structure. For example, many reported noting alternations that occurred in the grammar. However, Cleeremans and McClelland pointed out that there were many instances in which performance facilitation resulting from sensitivity to sequential structure (including sensitivity to three elements of context) was not accompanied by corresponding free report of the rules. They concluded that, overall, subjects could freely report very little of their knowledge of the sequential structure of the material.

In terms of the matrix-scanning tasks, Lewicki et al. (1987) found that, despite extensive practice, subjects were unable to report anything related to the underlying rules and denied awareness of their existence. These basic findings were replicated by Stadler (1989) using the same paradigm. He asked subjects a set of questions that began with general items and became increasingly more specific. Subjects did not describe any rules and appeared to be surprised that any rules existed. Similarly, after practice at the Lewicki, Hill, and Bizot (1988) task, subjects failed to report having noticed any pattern in the sequence of exposures, and none of them reported even suspecting that the sequential structure of the material had been manipulated. However, Perruchet, Gallego, and Savy (1990) argued that the improvement in performance on Lewicki et al.'s (1988) task could be accounted for without suggesting that subjects acquire knowledge of the composition rules and that subjects partition the sequence into logical blocks of five trials. Rather, they suggested that the results could be explained by the relative frequency of a few simple sequences of target locations. Perruchet et al. carried out an extended replication of the Lewicki et al. experiment and showed that differences in reaction time on the first two and last three trials of each logical block of five trials could be attributed to the relative frequency of particular target transitions throughout the experiment. However, they did not report whether subjects could describe their knowledge in free report.

In summary, there is evidence from both the SRT task and the matrix-scanning task that much of the knowledge that is acquired in sequence-learning tasks cannot be easily articulated by subjects. A similar conclusion was reached by Kushner et al. (1991), who used a sequence-prediction task. They found that many subjects could predict the next element in a sequence without being able to state what the underlying rules were, nor did subjects

offer specific sequences that they had memorized (see Perruchet, 1994; note that there have been two failures to replicate learning with this paradigm, by Reber and by Cleeremans; Reber, personal communication, October, 1995). However, in many of the studies, not all knowledge of sequential structure was inaccessible in this way—subjects could verbalize something about the sequential structure of the stimuli.

Objective threshold. Willingham et al. (1989) attempted to see if a cued-report test could elicit the knowledge underlying SRT performance of the unaware subjects. Specifically, they used a *generate* task, which involved displaying the repeating sequence but required subjects to predict the next stimulus position, rather than responding to the present stimulus position. The performance of the unaware subjects was at the same level as that of control subjects who had not been exposed to the repeating sequence; the performance of the aware group was greater than that of both. Apparently, unaware subjects could not produce their knowledge of the sequence on a relevant cued-report task.

However, this turned out not to be compelling evidence. Shanks et al. (1994), in a replication study, found that unaware subjects performed significantly above chance on the generate task and pointed out that there was a trend in this direction in the Willingham et al. data. Moreover, Perruchet and Amorin (1992) argued that Willingham et al. (and others) had simply used an insensitive version of the generate test. In the generate task, the stimulus is displayed until the subject makes the correct prediction for the next trial, or, in some experiments (Cohen, Ivry, & Keele, 1990), the stimulus is displayed at its correct location after the subject's response, whether correct or not. There were two main parts to their argument: (1) the instructions given to subjects before the generate task did not mention that subjects should reproduce prior sequences, and (2) correcting the subject's response could make it difficult to recall the immediately preceding sequence. Perruchet and Amorin adapted the generate task in two ways: by changing the instructions to emphasize the relation between the study and test phase, and by eliminating feedback on response accuracy. Using this modified free-generation test, they found that subjects could produce sequences that partially followed the constraints of the training stimuli after only two 100-trial blocks of training. In a second experiment, in addition to the generate task, Perruchet and Amorin assessed knowledge using a recognition procedure of the four-trial chunks composing the repeating sequence. As in the artificial grammar learning studies reviewed above, they found that subjects could recognize chunks that did follow the constraints of the training stimuli. Unfortunately, Perruchet and Amorin did not divide their subjects into groups based on their self-reported awareness of a pattern, so their results still leave open the claim that there is a subset of subjects that cannot produce their sequential knowledge on a generate task.

Reed and Johnson (1994) provided initially compelling evidence that there could be knowledge not elicited

by the generation task. Subjects were trained on one sequence on the SRT task for 17 blocks and simultaneously performed a secondary tone-counting task (in order to impair explicit learning). After training, one group of subjects (SRT test group) carried on performing the SRT task, but with a different transfer sequence. After training, another group of subjects (the generation test group) were tested on a cued-generation task that answered the criticisms of Perruchet and Amorin (1992): Subjects were explicitly instructed to reproduce the training sequence, and their answers were not corrected. A final group (control generation group) performed the generate task after training on a different sequence. The reaction times of the SRT test group increased substantially (81 msec) on the transfer sequence, relative to the training sequence, indicating good learning of the training sequence on the SRT task. On the other hand, the generation test group and the control generation group performed at the same level on the generation task. While these results were promising, Shanks and Johnstone (in press) pointed out that Reed and Johnson had used different sequences for the control generation group and for the SRT test group in the transfer phase: That is, the SRT and generation tests of knowledge were potentially testing different pieces of knowledge. Furthermore, in a replication study in which this problem was corrected, Shanks and Johnstone found significant amounts of knowledge were elicited on the cued-generation task used by Reed and Johnson and also on a free-generation task.

Taken together, the evidence suggests that the knowledge underlying SRT performance can be elicited by the generate task. Consistently, Perruchet and Amorin reported high correlations between reaction times and generation-task performance. Before accepting the conclusion that both tasks tap the same data base, however, we need to deal with the numerous studies reporting functional dissociations between SRT and generation performance—that is, cases in which an independent variable influences the SRT and generation tasks differently. These dissociations have been used as evidence that the generation and SRT tasks do, in fact, tap different knowledge bases.

A number of variables have been found to affect generation-task performance but not SRT performance. Cohen et al. (1990, Experiment 1) asked subjects performing the SRT task to perform simultaneously a secondary task, counting the number of high-pitched tones in a sequence of tones. Subjects received either an easy version of the tone task, in which there were few high-pitched tones, or a difficult version, in which there were many high-pitched tones. Results showed that SRT performance was unaffected by variation in distractor-task difficulty. However, subjects in the difficult-distractor condition performed less well on the generate task than did subjects in the easy-distractor condition, indicating a dissociation between the SRT and generation tasks. Howard, Mutter, and Howard (1992) showed that, if subjects initially simply observe the sequence, then their later responses are just as fast as those of subjects who re-

sponded to the sequence from the beginning. However, observing subjects had superior generation performance, relative to responding subjects, providing a further dissociation between generation and reaction time. Howard and Howard (1992) and Cherry and Stadler (1995) found that, on the SRT task, elderly subjects and young subjects performed at the same level; however, on the generation task, the elderly subjects were significantly impaired. Finally, Nissen, Knopman, and Schacter (1987) found that scopolamine did not affect performance on the SRT task but significantly impaired performance on the generation task.

Conversely, a number of variables have been found that affect SRT performance but not generation-task performance. Cohen et al. (1990, Experiment 4) used two different types of sequences: “unique” and “ambiguous” (discussed in more detail below). In unique sequences, each element was perfectly predictable from the preceding element, whereas, in ambiguous sequences, each element was preceded by two others equally often. Cohen et al. found that dual-task conditions interfered with learning the ambiguous sequence but not the unique sequence. There was no difference in performance on the generation task between single- and dual-task conditions for either sequence. Frensch, Buchner, and Lin (1994) also used a secondary tone-counting task and systematically manipulated the timing of the tone with respect to the onset of the primary SRT stimulus. As will be discussed below, this significantly affected the amount of learning shown on the SRT task but had no significant effects on generation performance. Frensch and Miner (1994) manipulated the response–stimulus interval (RSI) on the SRT task (500 vs. 1,500 msec). There was no evidence, on the SRT task, that subjects learned the sequence when RSI = 1,500 msec; however, there was good learning when RSI = 500 msec. In the generate task, performance was not significantly better in the 500-msec condition than in the 1,500-msec condition. Finally, Jiménez and Cleeremans (1994), using a finite-state grammar rather than a fixed repeating sequence, found that reaction time significantly correlated with the actual conditional probabilities of elements even when subjects’ generation performance had been partialled out. Jiménez and Cleeremans concluded that the SRT task was tapping knowledge not elicited by the generation task. We will consider another explanation below.

All of these dissociations, at first, appear to suggest that some knowledge elicited by the SRT task may in fact not be elicited by the generate task: The two tasks may tap different knowledge bases. One objection to this conclusion is that the dissociations may simply reflect the vagaries of Type II errors (e.g., Perruchet & Gallego, 1993), since they are all (except for the Jiménez and Cleeremans’s 1994 result) of the form of a significant result on one measure and a nonsignificant result on another. They do not, for example, provide the more compelling pattern of a reversed association (Dunn & Kirsner, 1988). There is also another explanation. The dissociations could be explained by assuming that the knowledge un-

derlying SRT performance can be elicited by an appropriately sensitive measure of generation performance, but generation performance is additionally affected by another source of variance (e.g., hypotheses that the subject might have explicitly formulated). If explicit knowledge is more likely to be applied to the generation, rather than incidental, reaction time tasks, dissociations can be expected between the two measures. If the explicit knowledge does not reflect actual conditional probabilities as accurately as the implicit knowledge does, generation performance will not correlate as well with conditional probabilities as reaction time performance does (as found by Jiménez and Cleeremans, 1994).

Finally, we turn to a different type of cued-report test. Willingham et al. (1993) trained subjects on the SRT task and then tested subjects' ability to recognize the right sequence (of 16 elements) among distractor sequences. Subjects' recognition ratings were above chance. However, when Willingham et al. selected subjects who were at chance on the recognition test, they still showed significant sequence knowledge on the SRT task. Perruchet and Gallego (1993) argued that these data did not indicate that the knowledge underlying performance on the SRT task could not be elicited by a relevant recognition test. The recognition test was a test of subjects' knowledge of the whole sequence; the SRT task would be sensitive to subjects' knowledge of only fragments of the sequence.

Researchers have also used the generate task in conjunction with matrix-scanning tasks. For example, Stadler (1989) trained subjects on Lewicki et al.'s (1987) task. He then provided subjects with six simple trials and asked them to predict the location of the target on the seventh trial. Although Stadler reported substantial learning with practice on the scanning task, subjects' performance was at chance on the generate test. Chance for the generate task was 25%, and subjects scored 24%. The upper limit of the 95% confidence interval for subjects' score was 31%, indicating that the experiment had reasonable sensitivity. However, Perruchet et al. (1990), using the same paradigm as Lewicki et al. (1988), found that subjects performed significantly above chance on a generate task.

In summary, the bulk of the evidence suggests that the knowledge used to control reaction time can be elicited by appropriate cued-report tests, such as the generation task. There is only one result inconsistent with this conclusion (Stadler, 1989) that bears further investigation.

Subjective threshold. Many subjects report that they were unaware that there was any structure in the stimuli, even though they were demonstrably influenced by the structure. For example, Willingham et al. (1993) found that subjects who claimed that the stimuli were completely random still showed reliable sequence learning on the SRT task. Shanks and Johnstone (in press) found that subjects could, with some accuracy, freely generate training sequences, but almost all subjects claimed that they were guessing. In another experiment, Shanks and Johnstone found that, when subjects were asked to freely generate sequences, the average confidence rating of subjects trained on a repeating sequence was no higher

than the average confidence rating of subjects trained on a nonrepeating sequence. Shanks and Johnstone argued that this was because their trained subjects had little metaknowledge about their knowledge.

Features of Implicit Learning

Transfer to different domains. Willingham et al. (1989) investigated the specificity of subjects' knowledge when superficial aspects of the task were changed. They adapted the reaction time task so that subjects saw a sequence of colors at different positions and mapped a given color onto a response. In one condition, there was a repeating sequence of stimulus positions; in another condition, there was a repeating sequence of motor responses. After training with either the perceptual sequence or the response sequence, subjects were transferred to the standard reaction time task, mapping positions onto responses. The sequence of locations was the same as that used in the perceptual sequence conditions, and the sequence of responses was the same as that used in the response sequence condition. Willingham et al. found very little evidence for positive transfer of learning from either training task to the transfer task. (In fact, very little learning occurred during training with the perceptual sequence.) They suggested that the transfer specificity was due to the fact that learning on the task was neither solely perceptual nor solely motor but represented instead the mapping rules governing performance. That is, the production rules required for performing the transfer task differed from those developed during training. Howard et al. (1992) argued that the learning could have been perceptual; in the perceptual sequence condition, location was incidental and so may not have been well attended and thus not well encoded. Consistently, Cohen et al. (1990) found that, if subjects were transferred from using different fingers for different buttons to responding with just the index finger of one hand, there was no detectable drop in performance. Cohen et al. concluded that the sequence specification was not tied to a particular motor effector.

In terms of matrix-scanning tasks, Stadler (1989) also investigated whether learning relied primarily on perceptual or motor processes. In the response transfer condition, subjects used a different response apparatus than had been used during training; in the position transfer condition, the position of the target within a quadrant during the six simple trials was different. Stadler reported greatly diminished transfer in the position transfer condition but not in response transfer condition. He suggested that learning in the task was based on the perceptual processes involved in locating the target rather than in the processing involved in responding to the target.

In summary, performance in sequence-learning tasks does show transfer specificity. The inflexibility of the subjects' knowledge could be further tested by informing subjects of the mapping in perceptual features between old and new stimuli (cf. Berry & Broadbent, 1988); if the knowledge is as inflexible as the implicit knowledge of controlling dynamic systems, then even subjects

informed of the mapping will not fully transfer their knowledge.

Learning tends to be associated with a focus on particular items rather than on the underlying rules. Performance on these tasks is largely associated with incidental learning conditions: Subjects just have to respond to the presented items. Subjects are not instructed to search for the underlying rules, and many subjects express considerable surprise when told that there is a sequential structure to the stimuli. However, Frensch and Miner (1994) found that asking subjects to search for the rules did improve performance on the SRT task. Future research could investigate whether rule-search instructions interact with stimulus complexity (e.g., the use of a finite-state grammar, as used by Cleeremans & McClelland, 1991, or a fixed second-order sequence rather than a first-order sequence), as in the artificial grammar and dynamic system control paradigms.

Robustness: Psychological and organic impairment. Nissen and Bullemer (1987) and Nissen, Willingham, and Hartman (1989) showed that amnesic patients became sensitive to the sequential structure inherent in the task and that their reaction times improved in the same way as those of normals. Similarly, Knopman and Nissen (1987) found that the level of specific sequence learning for Alzheimer's patients was nonsignificantly different from that of controls. However, Ferraro, Balota, and Connor (1993) found that Alzheimer's and Parkinson's patients were impaired on the SRT task, relative to controls, and Knopman and Nissen (1991) found that patients with Huntington's disease (HD) were impaired. If the neural basis of implicit learning is more anatomically distributed than is that of explicit learning, then widespread general neural degradation will eventually impair even implicit learning performance. Thus, in the Knopman and Nissen (1987) study, the patients had impaired explicit memory but had normal learning on the SRT task. Ferraro et al. suggested that they used more severely demented patients than did Knopman and Nissen, and their patients were impaired on both explicit memory and the SRT task. Indeed, Ferraro et al. showed that the severity of the dementia in their patients made a substantial difference to the level of learning on the SRT task.

Knopman and Nissen (1991) suggested that, if patients have to concentrate inordinately on the act of making a motor response, they may effectively be in a dual-task situation. As we will see below, there appears to be a component of SRT learning sensitive to attentional demands. However, Jackson, Jackson, Harrison, Henderson, and Kennard (1995) found that their Parkinson's patients started off with reaction times nonsignificantly different from those of the controls, suggesting that they had no difficulty in making the motor response per se. Even so, the Parkinson's patients learned significantly less on an SRT task than did controls. Jackson et al. argued that the execution of implicitly acquired knowledge of sequential order depends on the basal ganglia, which are damaged in Parkinson's disease. If Parkinson's and HD patients are not impaired in any general implicit learning

ability, but only in its motor expression, they should perform normally on the artificial grammar learning task. This is an issue that future research needs to investigate.

Nissen et al. (1987) showed that injecting scopolamine had no influence on learning as assessed by reaction time but did reduce generation performance. (Scopolamine also reduced free recall of a list of words but had no effect on repetition priming.) That is, scopolamine affected the putative explicit but not implicit measure.

Robustness: Other individual difference variables. Howard and Howard (1989, 1992), Frensch and Miner (1994), and Cherry and Stadler (1995) found that the elderly demonstrated as much learning on the SRT task as did the young. Frensch and Miner found that young adults were superior to the elderly on the SRT task when it was performed with a secondary tone-counting task. However, this may reflect the demands of the secondary task, requiring more attention by the elderly than the young. As we will see below, there is a component of subjects' knowledge of the SRT task that is sensitive to attention. Cherry and Stadler (1995) found that the elderly from a low socioeconomic status (SES) were impaired, relative to the young from a high SES, indicating that the task is not immune to the effects of individual difference variables. This still leaves open the claim that implicit knowledge is more robust than explicit knowledge. Consistent with a greater relative robustness of implicit knowledge over explicit knowledge, Cherry and Stadler found that both high- and low-SES elderly were impaired on the generation task, relative to the young.

Robustness: Secondary tasks. A number of studies have examined the effects of introducing secondary tasks. Nissen and Bullemer (1987) found that requiring subjects to perform a concurrent tone-counting task eliminated learning of the repeating sequence in the SRT task. Subjects who received the repeating sequence under dual-task conditions did not improve any more than did subjects who received the random sequence. Cohen et al. (1990) suggested that the detrimental effect of the secondary task may be due to the particular sequence of repeating stimuli used by Nissen and Bullemer. They suggested that only structures with at least some unique associations (e.g., Event A always follows Event C) could be learned under dual-task conditions. Conversely, they argued that structures with all items repeated in different orders in different parts of the structure required attention for learning. They constructed two sequences: One sequence (the unique sequence) had unique associations between successive stimuli; the other sequence (the ambiguous sequence) had no unique associations (each stimulus was followed by two others equally often). Under dual-task conditions, they found significant residual learning only of the unique sequence and argued that attention was necessary for learning ambiguous sequences (cf. the similar results of Dienes et al., 1991, in the artificial grammar learning paradigm). However, Cleeremans and McClelland (1991) pointed out that the dual-task conditions interfered with learning both the unique and the ambiguous sequences to roughly the same extent, suggesting that

they may in fact rely on the same mechanism. Since then, a number of studies have found significant learning of ambiguous sequences under dual-task conditions (Keele & Jennings, 1992; Reed & Johnson, 1994).

Whether or not unique and ambiguous sequences require different learning mechanisms, Curran and Keele (1993) provided evidence that there were nonetheless two different types of sequence learning, one requiring more attention than the other. In their first experiment, there were two groups of subjects: (1) intentional subjects, who were explicitly told what the sequence was, and (2) incidental subjects, who were not even told that there was a sequence. After initial exposure to the sequence under single-task conditions, the two groups differed substantially in the amount of specific sequence knowledge demonstrated by their reaction time performance (i.e., difference in reaction times between trials with random and fixed sequences). When a secondary task was then introduced, the two groups demonstrated equivalent amounts of sequence knowledge. Curran and Keele argued that “nonattentional learning” and “attentional learning” occurred in parallel, and the knowledge produced by the attentional learning could not be elicited under dual-task conditions. Their second experiment provided an even stronger demonstration. One group of subjects was told the sequence and was then given further exposure under single-task conditions; the other group of subjects was simply incidentally exposed to the sequence under dual-task conditions. Despite substantial differences in performance between the two groups at this stage, when both groups were subsequently tested under dual-task conditions, they demonstrated the same amount of specific sequence knowledge. In a third experiment, a single group of incidental dual-task subjects showed the same amount of knowledge when they were later tested under single-task conditions as they originally did under dual-task conditions. Curran and Keele argued that attentional and nonattentional learning operate independently, in parallel, and do not share any information. They further argued that awareness only enhances the attentional form of learning.

Stadler (1995) argued that dual-task conditions might interfere with learning on the SRT task because they disrupt consistent organization of the sequence. Consistently, Stadler found that introducing random long pauses between trials disrupted learning as much as did the concurrent tone-counting task. Furthermore, a concurrent memory span load (which should not produce inconsistent groupings of trials) disrupted performance substantially less.

Frensch et al. (1994) suggested that a secondary task interferes with sequential learning by limiting the amount of time that consecutive stimuli are simultaneously available in short-term memory. They varied the interval between the primary stimulus of the SRT task and the tone of the secondary task. Each primary stimulus followed a fixed interval after the subject’s response to the preceding primary stimulus. The more the tone delayed the subject’s responding to the primary stimulus, the worse the subject’s performance was. Frensch et al. argued that a

plausible interpretation of these results was that the slower the subject was, the greater the time between successive primary stimuli, and, so, the fewer primary stimuli were available in short-term memory (or, perhaps, the more time available for explicit hypothesis testing, disrupting an implicit learning mechanism; Reber, 1976).

Frensch and Miner (1994) directly tested the relevance of rate of presentation of the SRT stimuli. Under both dual-task and single-task conditions, learning on the SRT task was substantially greater when the interval between subjects’ responses and the next stimulus was 500 msec than when the interval was 1,500 msec (consistent with the results of Stadler, 1995). They argued that rate of presentation might be important because, at faster rates, more items remain active in working memory. Such an explanation would be consistent with the claim that learning higher order *n*-grams (i.e., combination of *n* items) is impaired under dual-task conditions (i.e., slow presentation rates), when less context would be available in working memory. Similarly, the different types of sequence learning isolated by Curran and Keele (1993) may be due to the same mechanism learning and applying different contents under dual-task and single-task conditions—that is, low-order *n*-grams under dual-task conditions and high-order *n*-grams under single-task conditions. Future research needs to test these speculations more directly.

In summary, sequence learning is sensitive to dual-task conditions, which may be solely due to the effect of dual-task conditions on rate of presentation or may also be due to other factors, such as organization, or other ways in which concentration affects the processing of stimuli. This is a matter for future research to resolve. In any case, some sequence learning is relatively insensitive to the presence of secondary task (Curran & Keele, 1993).

Conclusion

In terms of the distinction between subjective and objective thresholds, there is little evidence that the knowledge relevant to performance lies below an objective threshold. In most, but not all, cases when relevant cue-report tasks were used, the knowledge was elicited. There was evidence that the knowledge lies below a subjective threshold: Subjects were often surprised to hear that there was structure to the stimuli. When subjects believed there was literally no structure or that they had not learned it, subjects still had demonstrably acquired knowledge. In terms of features that may distinguish knowledge above and below threshold, the knowledge could not be used in a flexible way (e.g., Willingham et al., 1989), and the knowledge is robust to psychological impairment, the effects of age, and, in some cases, to secondary tasks. Future research still needs to determine whether it is the subjective threshold that is doing the useful work in separating the different types of knowledge.

DISCUSSION

The reviews of the literatures for artificial grammar learning, the control of complex systems, and sequence

learning all provided evidence that subjects could learn to perform well in a task without being able to freely report what they had learned or why they had made the right decision. Cued-report tests, on the other hand, typically elicited considerably more knowledge than did free-report tests. That is, the accessibility of the knowledge is relative to the test used to elicit it. Shanks and St. John argued that because cued-report tests elicit subjects' knowledge, we should regard the knowledge as conscious. However, consciousness can be defined according to either an objective threshold or a subjective threshold (Cheesman & Merikle, 1984).

If consciousness is defined according to a subjective threshold, we are conscious of a piece of knowledge only when we know that we know it. Cued-report tests are then seen not as tests of what subjects are conscious of but as tests of what subjects know. The elicitation of knowledge by cued-report tests but not free-report tests suggests that the knowledge is above an objective threshold, but it could be below a subjective threshold. The claim that learning can produce knowledge below a subjective threshold was directly tested in the artificial grammar learning paradigm in two ways. First, subjects were often just as confident in incorrect decisions as in correct decisions (the lack of correlation criterion). Second, subjects could classify substantially above chance even when they believed that they were literally guessing (the guessing criterion). In terms of the guessing criterion, subjects in the sequential reaction time paradigm will generate accurate sequences while claiming that they are literally guessing. In all paradigms, subjects frequently reported that they did not know why they made the decisions they did or that they were surprised that there was any structure in the stimuli.

The zero-correlation criterion has been applied only to the artificial grammar learning paradigm. In further applications of the zero-correlation criterion, it should be kept in mind that it is not an *ipso facto* indicator of a lack of metaknowledge in any interesting sense. However, we argue that it does indicate a lack of metaknowledge in the case of artificial grammar learning (Dienes & Perner, 1996). To see what it is that subjects lack metaknowledge about in this situation, it will be useful to distinguish between deterministic and stochastic responses. A subject responds deterministically if the same test string always elicits the same response; a subject responds stochastically otherwise. If subjects respond deterministically to all test strings, this may be because subjects are using partially correct rules (e.g., "say 'grammatical' whenever an M starts a string, and 'nongrammatical' otherwise"). In this case, subjects may consider all applications of these rules subjectively as cases of knowledge, regardless of whether the rules lead to correct or incorrect performance. Hence, subjects' confidence judgments will be the same for correct and incorrect responses, despite insights into the rules that subjects are using (metaknowledge). In other words, a lack of correlation between confidence and classification judgments does not indicate a lack of metaknowledge in this situation. Subjects

may be perfectly aware of the rules that they are applying and when they are applying them (and even why they have formulated those rules and why they are applying them).

In a typical artificial grammar learning experiment, subjects actually respond stochastically to at least some strings. Reber (1989) interpreted these results as indicating that subjects know some exemplars perfectly and guess the rest of the time. If this were the case, the lack of correlation between confidence and accuracy would indicate that subjects do not know when they are applying their knowledge (and when they are applying guesses). A more general interpretation is that subjects are using stochastic rules that specify the probability for saying "grammatical" for each exemplar (Dienes, 1992). That is, we need not assume, as Reber does, that the probabilities must be 0, .5, or 1. In this more general scenario, one highly plausible assumption is that the probabilities reflect the degree to which the exemplars satisfy the learned constraints (see Dienes, 1992, for an accurate prediction of the subjects' probabilities given this assumption). For example, we assume that if subjects say "grammatical" to an exemplar more than 50% of the time, then the exemplar is subjectively more likely to be grammatical than nongrammatical. If subjects had general access to these probabilities, then they should be more confident of "grammatical" decisions associated with high, rather than low, probabilities of saying "grammatical." The probabilities must correlate with the actual grammaticality of the exemplars if subjects perform above chance. Thus, if subjects could have used the different probabilities to inform their confidence ratings, confidence would have correlated with accuracy. Therefore, under these assumptions, a lack of correlation would indicate that subjects lack metaknowledge about the strength of their rules. In summary, a lack of correlation between confidence and accuracy indicates a lack of metaknowledge whenever subjects use stochastic rules that lead to above-chance performance, but not when subjects use deterministic rules.

One contrast between the application of the guessing criterion and the zero-correlation criterion is in terms of what a significant result means in each case. Finding performance significantly above chance on trials on which the subjects believe that they are guessing (guessing criterion) indicates the presence of some implicit knowledge but does not rule out that the subject also has some explicit knowledge in that domain. Finding that confidence is significantly greater for correct trials than for incorrect trials indicates the presence of some metaknowledge but does not rule out the presence of some implicit knowledge. Probably, in most learning situations, subjects will acquire a mixture of implicit and explicit knowledge.

If we are to be sure that we have found a psychologically real "threshold" of accessibility (be it subjective or objective), knowledge above and below the threshold should be qualitatively different. If the knowledge is below a subjective threshold, subjects may not know what type of (external or internally generated) cue will elicit their knowledge. Thus, one might expect that, if given a sim-

ilar task in a different domain, subjects may not apply their knowledge if the relevant cues happen not to be present. We reviewed evidence from all three literatures that there was indeed limited transfer of learning. Of course, people fail to apply much of their explicit knowledge in situations where it would be useful. In many cases, however, this is because subjects fail to realize that there is a useful mapping from their previous knowledge to the new situation (e.g., Gick & Holyoak, 1980). What is perhaps peculiar to implicit knowledge is that transfer might fail even when the subject is aware of the exact mapping between corresponding elements of the original and transfer tasks, as was found in the control of complex systems paradigm. In implicit learning paradigms, cued-report tests are transfer tasks that use components of the learning task, and, so, the corresponding elements of the two tasks are identical (and, thus, the calling conditions of knowledge merely below subjective threshold would be met).

Previous researchers have also argued that implicit knowledge may be optimally acquired by a focus on items rather than on rules, and it may be relatively robust. We reviewed evidence that artificial grammar learning, learning the control tasks, and sequence learning all can take place when the subjects are focused on items rather than on rules. Finally, in all three paradigms, the acquisition of the knowledge appears unaffected by amnesia and, at least in the case of artificial grammar learning, is unaffected by psychiatric impairment and variation in IQ. However, the knowledge is partly affected by secondary tasks. On the other hand, there are other tasks, arguably explicit, for which learning occurs only when subjects are searching for rules, for which subjects can say how they have arrived at the answer, for which confidence correlates with accuracy, and which are very much affected by amnesia, psychiatric impairment, and variation in IQ. The tendency of these features to go together in certain tasks suggests that there are distinct learning modes (that we suggest produce knowledge either above or below a subjective threshold).

Interestingly, most of the neuropsychological evidence on the implicit–explicit distinction can be interpreted as showing that knowledge can be below a subjective threshold: Amnesics are influenced by past events without knowing that they are; blindsight patients can discriminate visual features without knowing that they can; prosopagnosics also often do not know that they are influenced by familiar faces (see Berry & Dienes, 1993). We believe that the case for a psychologically real subjective threshold in the many domains that have been used to investigate unconscious learning is far from complete, but we believe that the existing data strongly indicate that this will be a fruitful direction for future research.

REFERENCES

ABRAMS, M., & REBER, A. S. (1988). Implicit learning: Robustness in the face of psychiatric disorders. *Journal of Psycholinguistic Research*, **17**, 425-439.
 ALLEN, R., & REBER, A. S. (1980). Very long-term memory for tacit knowledge. *Cognition*, **8**, 175-185.

ALTMANN, G. T. M., DIENES, Z., & GOODE, A. (1995). On the modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 899-912.
 BERRY, D. C. (1984). *Implicit and explicit knowledge in the control of complex systems*. Unpublished doctoral dissertation, University of Oxford.
 BERRY, D. C. (1991). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology*, **43A**, 881-906.
 BERRY, D. C., & BROADBENT, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, **36A**, 209-231.
 BERRY, D. C., & BROADBENT, D. E. (1987). The combination of explicit and implicit learning processes. *Quarterly Journal of Experimental Psychology*, **39A**, 585-609.
 BERRY, D. C., & BROADBENT, D. E. (1988). Interactive tasks and the implicit–explicit distinction. *British Journal of Psychology*, **79**, 251-272.
 BERRY, D. C., & BROADBENT, D. E. (1995). Implicit learning in the control of complex systems. In P. Frensch & J. Funke (Eds.), *Complex problem solving* (pp. 131-150). Hillsdale, NJ: Erlbaum.
 BERRY, D. C., & DIENES, Z. (1993). *Implicit learning: Theoretical and empirical issues*. Hove, U.K.: Erlbaum.
 BORNSTEIN, R. F., & PITTMAN, T. S. (1992). *Perception without awareness: Cognitive, clinical, and social perspectives*. New York: Guilford.
 BRIGHT, J., & BURTON, M. (1994). Past midnight: Semantic processing in an incidental learning task. *Quarterly Journal of Experimental Psychology*, **47A**, 71-90.
 BROADBENT, D. E. (1977). Levels, hierarchies, and the locus of control. *Quarterly Journal of Experimental Psychology*, **29A**, 181-201.
 BROADBENT, D. E. (1991). Recall, recognition, and implicit knowledge. In W. Kessen, A. Ortony, & F. I. M. Craik (Eds.), *Essays in honour of George Mandler* (pp. 125-134). Hillsdale, NJ: Erlbaum.
 BROADBENT, D. E., & ASTON, B. (1978). Human control of a simulated economic system. *Ergonomics*, **21**, 1035-1043.
 BROADBENT, D. E., FITZGERALD, P., & BROADBENT, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology*, **77**, 33-50.
 BROOKS, L. R., & VOKEY, J. R. (1991). Abstract analogies and abstracted grammars: Comments on Reber (1989) and Mathews et al. (1989). *Journal of Experimental Psychology: General*, **120**, 316-323.
 BUCHNER, A., FUNKE, J., & BERRY, D. (1995). Negative correlations between control performance and verbalizable knowledge: Indicators for implicit learning in process control tasks? *Quarterly Journal of Experimental Psychology*, **48A**, 166-187.
 CHAN, C. (1992). *Implicit cognitive processes: Theoretical issues and applications in computer systems design*. Unpublished doctoral dissertation, University of Oxford.
 CHEESMAN, J., & MERIKLE, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, **36**, 387-395.
 CHERRY, K. E., & STADLER, M. A. (1995). Implicit learning of a non-verbal sequence in younger and older adults. *Psychology & Aging*, **10**, 379-394.
 CLEEREMANS, A., & MCCLELLAND, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, **120**, 235-253.
 COCK, J., BERRY, D. C., & GAFFAN, E. A. (1994). The role of similarity processing in an incidental learning task. *Quarterly Journal of Experimental Psychology*, **47A**, 1015-1034.
 COHEN, A., IVRY, R., & KEELE, S. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 17-30.
 CURRAN, T., & KEELE, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 189-202.
 DIENES, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science*, **16**, 41-79.
 DIENES, Z., & ALTMANN, G. (1993, July). *The transfer of implicit knowledge across domains*. Paper presented at the meeting of the EPS and BBS, Toronto.
 DIENES, Z., & ALTMANN, G. (in press). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. Berry (Ed.), *How implicit is implicit learning?* Oxford: Oxford University Press.

- DIENES, Z., ALTMANN, G., KWAN, L., & GOODE, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 1322-1338.
- DIENES, Z., BROADBENT, D. E., & BERRY, D. C. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 875-887.
- DIENES, Z., & FAHEY, R. (1995). The role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 848-862.
- DIENES, Z., & PERNER, J. (1996). Implicit knowledge in people and connectionist networks. In G. Underwood (Ed.), *Implicit cognition* (pp. 227-256). Oxford: Oxford University Press.
- DULANY, D. E. (1962). The place of hypotheses and intentions: An analysis of verbal control in verbal conditioning. In C. W. Eriksen (Ed.), *Behavior and awareness* (pp. 102-129). Durham, NC: Duke University Press.
- DULANY, D. E., CARLSON, R. A., & DEWEY, G. I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General*, **113**, 541-555.
- DUNN, J. C., & KIRSNER, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, **95**, 91-101.
- ELLIS, N. (1993). Rules and instances in foreign language learning: Interactions of explicit and implicit knowledge. *European Journal of Cognitive Psychology*, **5**, 289-318.
- ERDELYI, M. H. (1986). Experimental indeterminacies in the dissociation paradigm of subliminal perception. *Behavioral & Brain Sciences*, **9**, 30-31.
- ERDELYI, M. H., & BECKER, J. (1974). Hypermnnesia for pictures: Incremental memory for pictures but not words in multiple recall trials. *Cognitive Psychology*, **6**, 159-171.
- FERRARO, F. R., BALOTA, D. A., & CONNOR, L. T. (1993). Implicit memory and the formation of new associations in nondemented Parkinson's disease individuals and individuals with senile dementia of the Alzheimer type: A serial reaction time (SRT) investigation. *Brain & Cognition*, **21**, 163-180.
- FRENSCH, P. A., BUCHNER, A., & LIN, J. (1994). Implicit learning of unique and ambiguous serial transitions in the presence and absence of a distractor task. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 567-584.
- FRENSCH, P. A., & MINER, C. S. (1994). Effects of presentation rate and individual differences in short-term memory capacity on an indirect measure of serial learning. *Memory & Cognition*, **22**, 95-110.
- GICK, M., & HOLYOAK, K. (1980). Analogical problem solving. *Cognitive Psychology*, **12**, 306-355.
- GOMEZ, R. L., & SCHVANEVELDT, R. W. (1994). What is learned from artificial grammars? Transfer tests of simple association. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 396-410.
- GREEN, R. E. A., & SHANKS, D. R. (1993). On the existence of independent explicit and implicit learning systems: An examination of some evidence. *Memory & Cognition*, **21**, 304-317.
- HAYES, N. A., & BROADBENT, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, **28**, 249-276.
- HEBB, D. O. (1961). Distinctive features of learning in the higher animal. In J. F. Delafresnaye (Ed.), *Brain mechanisms and learning* (pp. 37-51). London: Oxford University Press.
- HOWARD, D. V., & HOWARD, J. H. (1989). Age differences in learning serial patterns: Direct vs. indirect measures. *Psychology & Aging*, **4**, 357-364.
- HOWARD, D. V., & HOWARD, J. H. (1992). Adult age differences in the rate of learning serial patterns: Evidence from direct and indirect tests. *Psychology & Aging*, **7**, 232-241.
- HOWARD, J. H., MUTTER, S. A., & HOWARD, D. V. (1992). Serial pattern learning by event observation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 1029-1039.
- JACKSON, G. M., JACKSON, S. R., HARRISON, J., HENDERSON, L., & KENNARD, C. (1995). Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit. *Neuropsychologia*, **33**, 577-593.
- JACOBY, L. L., & BROOKS, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 18, pp. 1-47). New York: Academic Press.
- JIMÉNEZ, L., & CLEEREMANS, A. (1994). Direct and indirect measures of implicit learning. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society* (pp. 445-450). Hillsdale, NJ: Erlbaum.
- KEELE, S. W., & JENNINGS, P. (1992). Attention in the representation of sequence: Experiment and theory. *Human Movement Science*, **11**, 125-138.
- KNOPMAN, D., & NISSEN, M. J. (1987). Implicit learning in patients with probable Alzheimer's disease. *Neurology*, **37**, 784-788.
- KNOPMAN, D., & NISSEN, M. J. (1991). Procedural learning is impaired in Huntington's disease: Evidence from the serial reaction time task. *Neuropsychologia*, **29**, 245-254.
- KNOWLTON, B. J., RAMUS, S., & SQUIRE, L. R. (1992). Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science*, **3**, 172-179.
- KNOWLTON, B. J., & SQUIRE, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 79-91.
- KOLERS, P., & ROEDIGER, H. L., III (1984). Procedures of mind. *Journal of Verbal Learning & Verbal Behavior*, **23**, 425-447.
- KUSHNER, M., CLEEREMANS, A., & REBER, A. (1991). Implicit detection of event interdependencies and a PDP model of the process. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 215-220). Hillsdale, NJ: Erlbaum.
- LEWICKI, P., CZYZEWSKA, M., & HOFFMAN, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 523-530.
- LEWICKI, P., HILL, T., & BIZOT, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, **20**, 24-37.
- MANZA, L., & REBER, A. S. (in press). Representation of tacit knowledge: Transfer across stimulus forms and modalities. In D. Berry (Ed.), *How implicit is implicit learning?* Oxford: Oxford University Press.
- MARESCAUX, P.-J., LUC, F., & KARNAS, G. (1989). Modes d'apprentissage sélectif et non-sélectif et connaissances acquises au contrôle d'un processus: Evaluation d'un modèle simulé [Selective and nonselective learning modes and acquiring knowledge of process control: Evaluation of a simulation model]. *Cahiers de Psychologie Cognitive*, **9**, 239-264.
- MATHEWS, 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, & Cognition*, **15**, 1083-1100.
- MCCANDREWS, M. P., & MOSCOVITCH, M. (1985). Rule-based and exemplar-based classification in artificial grammar learning. *Memory & Cognition*, **13**, 469-475.
- MCGEORGE, P., & BURTON, A. M. (1989). The effects of concurrent verbalization on performance in a dynamic systems task. *British Journal of Psychology*, **80**, 455-465.
- MCGEORGE, P., & BURTON, A. M. (1990). Semantic processing in an incidental learning task. *Quarterly Journal of Experimental Psychology*, **42A**, 597-610.
- McKELVIE, S. J. (1987). Learning and awareness in the Hebb digits task. *Journal of General Psychology*, **114**, 75-88.
- MERIKLE, P. M. (1992). Perception without awareness: Critical issues. *American Psychologist*, **47**, 792-795.
- MICHAS, I., & BERRY, D. C. (1994). Implicit and explicit processes in a second language learning task. *European Journal of Cognitive Psychology*, **6**, 357-381.
- MILLER, G. A., & SELFRIDGE, J. A. (1950). Verbal context and the recall of meaningful material. *American Journal of Psychology*, **63**, 176-185.
- MYERS, C., & CONNOR, M. (1992). Age differences in skill acquisition and transfer in an implicit learning paradigm. *Applied Cognitive Psychology*, **6**, 429-442.
- NEAL, A., & HESKETH, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin & Review*, **4**, 24-37.
- NISSEN, M. J., & BULLEMER, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, **19**, 1-32.

- NISSEN, M. J., KNOPMAN, D. S., & SCHACTER, D. L. (1987). Neurochemical dissociation of memory systems. *Neurology*, **37**, 789-794.
- NISSEN, M. J., WILLINGHAM, D., & HARTMAN, M. (1989). Explicit and implicit remembering: When is learning preserved in amnesia? *Neuropsychologia*, **27**, 341-352.
- PERRUCHET, P. (1994). Learning from complex rule-governed environments: On the proper function of conscious and unconscious processes. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 811-835). Cambridge, MA: MIT Press.
- PERRUCHET, P., & AMORIN, M. A. (1992). Conscious knowledge and changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Human Learning & Memory*, **18**, 785-800.
- PERRUCHET, P., & GALLEGO, J. (1993). Association between conscious knowledge and performance in normal subjects: Reply to Cohen and Curran (1993) and Willingham, Greeley, and Bardone (1993). *Journal of Experimental Psychology: Human Learning & Memory*, **19**, 1438-1444.
- PERRUCHET, P., GALLEGO, J., & SAVY, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. *Cognitive Psychology*, **22**, 493-516.
- PERRUCHET, P., & PACTEAU, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge. *Journal of Experimental Psychology: General*, **119**, 264-275.
- REBER, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior*, **6**, 855-863.
- REBER, A. S. (1969). Transfer of syntactic structures in synthetic languages. *Journal of Experimental Psychology*, **81**, 115-119.
- REBER, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning & Memory*, **2**, 88-94.
- REBER, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, **118**, 219-235.
- REBER, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.
- REBER, A. S., & ALLEN, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, **6**, 189-221.
- REBER, A. S., KASSIN, S. M., LEWIS, S., & CANTOR, G. (1980). On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 492-502.
- REBER, A. S., & LEWIS, S. (1977). Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition*, **5**, 333-361.
- REBER, A. S., & MILLWARD, R. B. (1968). Event observation in probability learning. *Journal of Experimental Psychology*, **77**, 317-327.
- REBER, A. S., WALKENFELD, F., & HERNSTADT, R. (1991). Implicit and explicit learning: Individual differences and IQ. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 888-896.
- REED, J., & JOHNSON, P. (1994). Assessing implicit learning with indirect tests: Determining what is learned about sequence structure. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 585-594.
- ROEDIGER, H. L., III, & BLAXTON, T. A. (1987). Effects of varying modality, surface features, and retention interval on priming in word-fragment completion. *Memory & Cognition*, **15**, 379-388.
- SANDERSON, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 729-747.
- SANDERSON, P. M. (1990). *Implicit and explicit control of a dynamic task: Empirical and conceptual issues* (Tech. Rep. No. EPRL-90-02). Urbana, IL: University of Illinois, Engineering Psychology Research Laboratory.
- SCHACTER, D., MCANDREWS, M., & MOSCOVITCH, M. (1988). Access to consciousness: Dissociations between implicit and explicit knowledge in neuropsychological syndromes. In L. Weiskrantz (Ed.), *Thought without language*. Oxford: Oxford University Press, Clarendon Press.
- SEGER, C. A. (1994). Implicit learning. *Psychological Bulletin*, **115**, 163-196.
- SHANKS, D. R., GREEN, E. A., & KOLODNY, J. A. (1994). A critical examination of the evidence for unconscious (implicit) learning. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 837-860). Cambridge, MA: MIT Press.
- SHANKS, D. R., & JOHNSTONE, T. (in press). Implicit knowledge in sequential learning tasks. In M. A. Stadler & P. A. Frensch (Eds.), *Handbook of implicit learning*. Thousand Oaks, CA: Sage.
- SHANKS, D. R., & ST. JOHN, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral & Brain Sciences*, **17**, 367-448.
- SQUIRE, L. R., & FRAMBACH, M. (1990). Cognitive skill learning in amnesia. *Psychobiology*, **18**, 109-117.
- STADLER, M. A. (1989). On learning complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1061-1069.
- STADLER, M. A. (1995). Role of attention in implicit learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 674-685.
- STANLEY, W. B., MATHEWS, R. C., BUSS, R. R., & KOTLER-COPE, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology*, **41A**, 553-577.
- TYLER, L. K., & MARSLER-WILSON, W. D. (1977). The online effects of semantic context on syntactic processing. *Journal of Verbal Learning & Verbal Behavior*, **16**, 683-692.
- WHITTLESEA, B. W. A., & DORKEN, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General*, **122**, 227-248.
- WILLINGHAM, D. B., GREELEY, T., & BARDONE, A. M. (1993). Dissociation in a serial response time task using a recognition measure: Comment on Perruchet and Amorin (1992). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 1047-1060.
- WILLINGHAM, D. B., NISSEN, M. J., & BULLEMER, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1047-1060.
- WRIGHT, R. (1993, July). *The role of "explicit" variables on artificial grammar learning*. Paper presented to the Third Annual Meeting of the Canadian Society for Brain, Behaviour, and Cognitive Science and the Experimental Psychology Society, Toronto.

NOTE

1. If control knowledge had transferred to different targets, then there would be evidence that the subjects had implicitly acquired knowledge asked for by the questionnaires. However, the evidence for transfer to different targets is ambiguous. For example, Berry and Broadbent (1987) found transfer between different targets, but Marescaux et al. (1989) did not.

(Manuscript received August 9, 1995;
revision accepted for publication May 17, 1996.)