32

Learning from Complex Rule-Governed Environments: On the Proper Functions of Nonconscious and Conscious Processes

Pierre Perruchet

ABSTRACT Improvement in performance of people faced with a complex rule-governed situation can occur without the emergence of any conscious knowledge of the rules. This chapter illustrates with a new example, and subsequently generalizes, the proposal that this phenomenon does not testify to the unconscious abstraction of the rules underlying the situation, as held by the prevalent, abstractionist interpretation. Indeed, performance improvement can be accounted for by a memory-based framework positing that subjects only learn specific fragments of the material, which constitute the basic functional unit of knowledge in most learning conditions. More important, findings lend weight to this framework when the experimental design contrasts its predictions with those derived from the abstractionist standpoint.

This theoretical shift has major consequences for the alleged implicitness of the processes engaged in this mode of learning. In particular, specific knowledge accounting for performance improvement appears to be usually available to conscious reflection upon explicit request in normal subjects. However, there is evidence that this knowledge can be used automatically to satisfy the demands of a subsequent task involving the same or similar material. Possible implications of these data on the role of nonconscious and conscious processes in adaptive behavior are discussed.

32.1 INTRODUCTION

Research on human learning has traditionally focused on simple experimental settings, such as classical conditioning and paired associate learning. Although some studies on concept or category learning have involved more complex situations, the level of complexity is generally adjusted so that subjects can break at least some of the rules structuring the material within one or a few experimental study sessions. These paradigms parallel the learning conditions in some real-world situations, such as scholastic ones, but there is no doubt that they provide a poor analogue of most natural settings. Indeed, the natural environment is structured by rules, whether physical, linguistic, social, or other, a majority of them so complex that they remain beyond the scope of common understanding. In addition, subjects in experimental situations of concept or category learning are typically asked for intentional rule searching, whereas laypeople rarely engage in active analysis of their natural environment in order to penetrate its deep structure.

Experimental paradigms involved in artificial grammar learning (Reber 1989), pattern sequence learning (Shanks, chap. 33, this volume), and learning in interactive situations (Berry, chap. 30, this volume) are devised to better fit natural conditions of adaptation. In these so-called implicit learning paradigms, subjects are faced with situations governed by complex, arbitrary rules, without being prompted for an explicit analysis of the task. In these conditions, performance improves, although subjects remain unable to articulate the rules governing the situations. This empirical outcome, first evidenced in the pioneering work of Arthur Reber (1967), has been unambiguously confirmed by all subsequent studies.

Demonstrating the ability of human subjects to improve their performance in a complex environment even though the structure of this environment is unavailable to the subjects' conscious awareness constitutes in itself a genuine contribution, and the novelty and implications of it are largely unnoticed. For example, the ability of people to learn from situations structured by complex arbitrary rules provides powerful support for the empiricist standpoint in the longstanding metatheoretical debate on the role of innate processes on behavior.

However, the major interest of disposing of laboratory analogues to complex learning in natural environments is that they offer an opportunity to explore the nature of the psychological processes involved in this form of adaptive behavior. The earlier interpretation (Reber 1967) accounts for improved performance by positing that human subjects unconsciously abstract rules embodied in the complex situations at hand and implicitly use this abstract knowledge to cope with subsequent situations. This interpretation is still advocated (Reber 1989; Lewicki, Hill, and Czyzewska 1992), and may even be prevalent, at least among investigators engaged primarily in related areas of research (e.g., Roediger 1990; Schacter 1987).

This interpretation has been the object of controversy, initiated by the papers of Brooks (1978) and Dulany, Carlson, and Dewey (1984) and extensively developed since then. At the heart of the debate are the two main components of the conventional position, centered on the abstract and the unconscious nature of the knowledge base acquired during learning.

It is important to make clear the relationships between the criticisms addressing the abstractionist stance and those addressing the claim about unconsciousness. From a logical standpoint, the availability in consciousness of the knowledge base underlying improved performance is independent of the nature of this knowledge. However, the two aspects cannot be handled separately in any empirical inquiry for an obvious reason: assessing whether the knowledge base responsible for improved performance is conscious or nonconscious requires this base to be clearly identified. Everyone presumably agrees that demonstrating subjects' unconsciousness of the abstract rules underlying the situation should be irrelevant if the change in performance is in fact attributable to quite a different form of knowledge. The argument developed in this chapter is that the main drawback of the conventional position is related primarily to the role this position confers to abstract knowledge in adaptive behavior, and only as a way of consequence, to its claims about unconsciousness or implicitness.

The shift away from an abstractionist stance has historical antecedents in research areas involving simpler learning paradigms. One example can be found in research on classical and instrumental conditioning in animals and humans. Rats' sensitivity to contingency relationships between stimuli in a conditioned suppression paradigm was originally demonstrated in the oft-cited paper by Rescorla (1968). The fact that a variety of animal species, including humans, are sensitive to contingency relationships in classical and instrumental conditioning paradigms has been abundantly confirmed since then (Wasserman 1990). The original interpretations naturally posited that this sensitivity to contingency testified to a genuine ability to abstract the covariations embedded in the conditioning situation. However, most current explanatory models of conditioning no longer assume that animals actually engage in contingency abstraction processes. As first demonstrated by Rescorla and Wagner (1972) and expanded on amply later, empirical results may be conveniently accounted for by simple associative processes (for recent reviews, see Papini and Bitterman 1990; Shanks 1991).

Categorization learning research was the theater for a similar conceptual shift. By and large, this field of research has been characterized by an upsurge in models postulating that the specific properties of items belonging to the same category are condensed into a small amount of abstract information (e.g., prototypes, occurrence frequency of independent features, or feature covariation). People were assumed to rely on this condensed form of representation when they had to make category assignments of new items in a subsequent test phase.

Since the seminal papers by Brooks (1978) and Medin and Shaffer (1978), such models have been challenged by models in which no abstraction takes place. The study exemplars are stored in memory with their specific properties, and subsequent categorization judgments of new items are made on the basis of their degree of similarity with the stored exemplars. These exemplar-similarity models have been tentatively applied to situations in which categories are ill defined (Perruchet, Pacteau, and Gallego 1993) or defined by logical rules (Nosofsky 1991). In all cases, stored exemplars were found to form a major, or at least substantial, component of the category representation, even when the classification rule was simple and made explicit for the subjects (Allen and Brooks 1991). Similar examples can be found in the problem-solving literature reviewed by Medin and Ross (1990).

This brief incursion into areas of research involving relatively simple experimental settings reveals a strikingly similar evolution in thought. The very patterns of performance that seemed at one time to be clear support for the engagement of abstractive processes have been subsequently reinterpreted within far more parsimonious frameworks, which no longer assume rule abstraction.

Recent work on implicit learning suggests that the evidence put forward for rule abstraction in more complex experimental settings is ready for a similar reappraisal. For the purposes of illustration, this chapter first presents a reinterpretation of a recent study on predictive behavior carried out in Reber's laboratory (Kushner, Cleeremans, and Reber 1991) in which subjects had to learn a rule that, according to the authors, was "far more complex than anything that has been studied to date." Then I will tentatively demonstrate that all the alleged evidence for abstraction amassed in the prior literature on implicit learning is inconclusive, because an alternative model accounts for the empirical data as well as, or better than, the mechanisms involved in the abstractionist framework. The model used here is similar to Brooks's (Brooks 1978; Vokey and Brooks 1992), insofar as it relies essentially on memory for specific events. However, it departs from Brooks's model regarding the notion of specific events, in a way that considerably extends the model's field of application.

This fundamental theoretical shift has major consequences with regard to the issue of implicitness, which will be discussed next. To anticipate, redirecting the postexperimental tests of awareness from an abstract form of knowledge to the memory of specific events leads us to find in the subjects' knowledge available to conscious awareness a sufficient basis to account for performance improvement in implicit learning situations. However, this does not imply that everything occurring in these situations relies on a conscious form of processing. I will argue that the involvement of nonconscious processes in the use of the memory for specific past events ought to be acknowledged.

32.2 AN ILLUSTRATIVE EXAMPLE

The Kushner, Cleeremans, and Reber Study

In the Kushner, Cleeremans, and Reber (1991) experiment, subjects viewed a computer screen on which a square stimulus could appear in three possible numbered locations, arranged to form the vertices of an invisible triangle. An event was defined as the appearance of the stimulus in a given location. A trial consisted of six successive events. The first five stimuli were displayed in rapid succession. After the fifth event had occurred, subjects were asked to predict where the sixth event would appear and had to enter their prediction by typing 1, 2, or 3 on a numeric keypad. Then the sixth stimulus appeared in its correct location. The location of the sixth event was determined on the basis of the relation between the locations at which the second and fourth stimuli had appeared. If these stimuli appeared at the same location, the sixth stimulus appeared in location 1. If they had been in a clockwise relation, the sixth stimulus appeared in location 2, and if they had been in a counterclockwise relation, the sixth stimulus appeared in location 3.

Discovering this set of rules from the displayed sequences involves very sophisticated operations: (1) picking up relevant events from the irrelevant context (the first, third, and fifth events in each trial were always irrelevant) and (2) abstracting a rule based on the relationships between the location of the relevant stimulus (considered in isolation, the location of stimulus on the second or on the fourth event had no predictive value). Despite this considerable difficulty, subjects did learn. Prediction accuracy, which started near chance level (33 percent correct responses), improved over the ten sessions of training and reached about 45 percent correct responses at the end.

From session 11 to session 14, the previous rules were modified by shifting the location of the sixth stimulus by one step: the sequences that previously ended in location 1 now ended in location 2, and so on. Subjects' prediction accuracy dropped back to chance level on session 11 and then improved again over the next sessions. In the four subsequent and final sessions, the location of the sixth stimulus was determined at random; as expected, subjects' performance remained around chance level.

Extensive postexperimental interviews revealed that in some specific cases, subjects had become aware of some of the regularities embedded in the material (for instance, the fact that sequences in which the five stimuli occurred in the same location always ended in location 1). However, these very fragmentary pieces of knowledge were clearly insufficient to account for performance. By and large, subjects were unable to report the correct rules and, in particular, the nature of the crucial events. When they were asked to rate each of the five stimuli in terms of their relevance to the prediction task on a five-point scale, the mean ratings were respectively 3.33, 3.5, 3.67, 2.67, and 1.83. At first glance, these results provide a compelling argument for the authors' conclusion: "Clearly, subjects have become sensitive to contingencies about which they are unable to report."

A Memory-based Reinterpretation

Suppose now that subjects memorize parts of the specific sequences of events, without abstracting any rules, and make predictions on the basis of their memory for earlier sequences. Each sequence was displayed once per session, so that the gradual improvement in performance across the first ten sessions would naturally be expected to occur under this assumption. The drop in prediction accuracy when rules are shifted in session 11 may also be readily explained, insofar as changing rules also change specific sequences. The fast subsequent improvement from sessions 11 to 14 may be accounted for by the familiarity subjects acquired with the material at this stage of the experiments, which is known to facilitate coding and hence memorization of new sequences. Finally, the drop in performance when the sixth event occurred randomly during the final phase of the experiment is also compatible with a pure memory account for trivial reasons.

In Kushner, Cleeremans, and Reber's experiment, the exhaustive set of possible sequences (243 [3^5]) was displayed during the training phase. This procedure makes abstractionist and memory interpretations impossible to disentangle, because abstracting rules and memorizing their product lead to the same outcome. In the experiment reported here, only two of the three possible instantiations of each rule were displayed during training, and the remaining possibilities were shown in a subsequent transfer phase. The particular sequences presented in each phase were chosen in order to pit the predictions of the two models against each other.

Table 32.1 shows that the rule "Same position, then issue A" was instantiated by 1-1 and 2-2 in the study phase. When 3-3 was displayed in the transfer phase, subjects who abstracted the rule were expected to predict A. However, subjects who memorized part of the sequences were expected to predict B or C, because the transfer sequence had one additional element in common with the study sequences ending with B or C than with the study sequences ending with A. Indeed, the former always included one stimulus in location 3, whereas the latter never did. The other predictions specified in table 32.1 were generated by applying similar reasoning.

Method

Subjects The subjects were six females, fourth-year university students majoring in psychology. They were paid 80 FF (around $14) and received a bonus of 0.10 FF (around 1.7 cents) per correct prediction.

Material The stimuli were white disks, 10 mm in diameter. They appeared within one of the three empty circles permanently displayed on the screen and numbered 1, 2, and 3. The circles were located at the vertices of an invisible equilateral triangle (10 cm per side).

Table 32.1 Summary Design of the Study and Transfer Phases of the Experiment

a. Study Phase (Sessions 1–9)
Observed Sequences

Second Event	Fourth Event	Sixth Event	Rule
1	1	A	Same position
2	2	A	
1	3	B	Clockwise
3	2	B	
2	3	C	Counterclockwise
3	1	C	

b. Transfer Phase (Session 10)

Observed Sequences		Model Predictions for the Sixth Event	
Second Event	Fourth Event	Rule Abstraction	Memory Based
3	3	A	B or C
2	1	B	A or C
1	2	C	A or B

Note: This table shows the observed (a) and predicted (b) locations of the sixth event as a function of the preceding events. This location is symbolized by letters rather than digits because allocation of the physical positions to each of the three rules was counterbalanced across subjects.

There were 162 (6 × 27) different sequences in the study phase. Indeed, only 6 out of the 9 (3^2) different possible combinations of events 2 and 4 were shown (see table 32.1), and each of these combinations was presented in 27 (3^3) different contexts, made up by the exhaustive combination of events 1, 3, and 5. In the transfer phase, there were 81 (3 × 27) different sequences, corresponding to combinations not displayed previously.

Procedure Subjects were recruited to participate in ten experimental sessions with the following constraints: no more than two immediately successive sessions, no more than four sessions a day, and no more than a 48-hour interval between any two sessions. All subjects completed the experiment within 5 days. As in Kushner, Cleeremans, and Reber, the study was presented as being about predictive behavior.

The 162 different study sequences were shown on each of the first nine sessions. For the first five stimuli of each trial, duration of both stimulus and interstimulus was set to 250 ms each. After subjects had pressed 1, 2, or 3 on the keyboard to enter their predictions, the stimulus was displayed in its correct location for 2000 ms. This location was counterbalanced across subjects, so that each subject had one of the 6 (3!) possible assignments of issues to the sequences. The next trial began 500 ms later. After trials 54 and 108, subjects were prompted on the screen to take a break; they had to press any key on the keyboard to continue with the experiment. The number of correct predictions was displayed on the screen at the end of each session.

Session 10 also consisted of 162 trials, with the 81 different transfer sequences presented twice. Subjects were kept unaware of the change in sequences; however, they were informed that they would no longer see the stimulus appearing in its correct location after entering their prediction. (Correctness of prediction in this session depended on the choice of a theoretical model, so it seemed preferable not to provide information at this stage.)

Order of trials was randomized without any constraints, both between sessions and between subjects, for study as well as for the transfer sessions.

Subjects were interviewed at the end of the experiment. They were questioned about their general strategy, their hypotheses about the structure of the task, and their ability to recall specific sequences. This interview was intended to replicate the Kushner, Cleeremans, and Reber procedure and not to provide a sensitive test of memory for the study sequences (which would have required the running of a recognition test before the transfer session). Subjects were debriefed about the objective of the study only after all of them had completed the entire experiment.

Results and Discussion

Figure 32.1 shows the proportion of correct predictions on the nine training sessions. Mean performance gradually improved from random level to a proportion of .415 correct (which significantly differs from .33, $t = 2.71, p =$

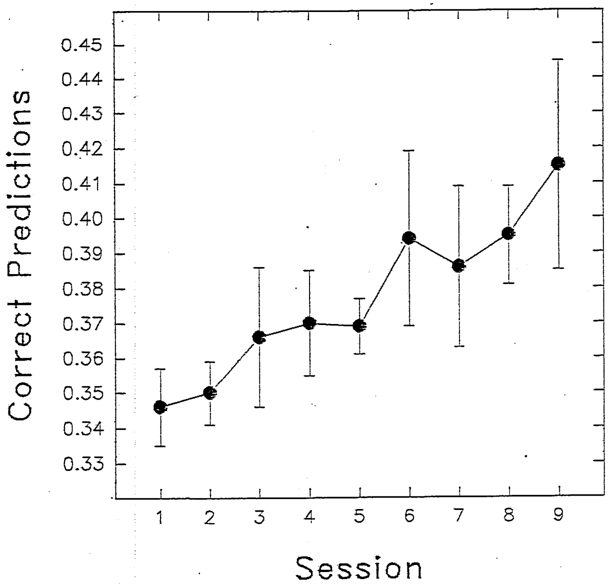

Figure 32.1 Proportion of correct predictions during the study phase. Error bars represent standard errors of the mean.

.042, $df = 5$) on the ninth session. Standard error also tended to increase over sessions, indicating that subjects did not learn at the same rate. The mean individual proportion of correct predictions on the whole study phase ranged from .340 to .407.

Results on the transfer session are shown in table 32.2. The data fitted nicely with the predictions of the memory model. Statistical analyses confirm that issues B and C, considered together, were predicted more often than chance after the sequences instantiating the rule "same position" ($t = 6.78$, $p = .001, df = 5$), and issues A and B were predicted more often than chance for the rule "counterclockwise" ($t = 3.99, p = .010, df = 5$). Although results also fit with the predictions of the memory model for the rule "clockwise," the difference from chance failed to reach significance ($t = 1.89, p = .116$, $df = 5$).[1]

Although performances exhibited the same general trend for all the subjects, individual patterns of responding were unequally differentiated. An individual degree of differentiation was computed as the overall proportion of predictions fitting with the memory model during the transfer phase. Interestingly, this score correlated strongly with the rate of correct predictions made during the first nine sessions (Pearson r: .804, $p = .027$). This result lends weight to the idea that predictions performed during the training and transfer phases tap the same knowledge base.

Table 32.2 Mean Proportion of Predictions Falling into A, B, and C, as a Function of the Preceding Events during the Transfer Phase

Observed Sequences		Predictions for the Sixth Event		
Second Event	Fourth Event	A	B	C
3	3	.204	.376	.420
2	1	.494	.213	.293
1	2	.426	.309	.265

Overall, these findings provide compelling evidence in favor of the memory model. One possible counterargument starts from considering that the present situation differs from Kushner, Cleeremans, and Reber's insofar as subjects were deprived of part of the information during training in the former but not in the latter situation. It may be that the whole set of possible exemplars is needed to abstract generative rules. This hypothesis has two damaging properties. The first is its unfalsifiability. The only way to test whether rules have been abstracted is to observe the ability of subjects to deal with new situations in transfer tasks. Presenting the whole set of exemplars generated by the rules during training makes the rule abstraction and pure memory models impossible to disentangle. Second, hypothesizing that subjects need all the possible instantiations of a rule to abstract this rule deprives the process of abstraction of its primary functional value: allowing adaptation to new situations.

One final experimental result pertains to the postexperimental interviews. They revealed extremely poor explicit knowledge about the task. Most subjects actively searched for regularities in the first few sequences, but gave up after their repeated failures to find an even partially successful rule. That is exactly what Kushner, Cleeremans, and Reber also reported about their study. However, it should be clear now that my conclusion is exactly the opposite from theirs. Whereas they inferred from the fact that performance improves without subjects being able to verbalize the rules that these rules were unconsciously abstracted, the conclusion of my experiment is that subjects did not abstract any rule at all, consciously or unconsciously (see Cleeremans, chap 31, this volume, for other arguments stemming from connectionist simulations of the same task).

32.3 GENERALIZING THE REFUTATION OF THE ABSTRACTIONIST MODEL

This section aims to show that virtually all the alleged evidence for abstraction may be accounted for by simple memory for the displayed items or, more generally, for a fragmentary part of them, and when predictions from abstractionist and memory-based models are pitted against each other, the results unambiguously stand in favor of the latter.

The Memory-based Model as an Alternative Framework

The Case of No or Inadequate Transfer Procedure The fact that performance of subjects repeatedly exposed to a structured situation improves across trials or sessions is obviously congruent with the hypothesis that subjects acquire a structural representation of this situation. Thus, the mere occurrence of learning is commonly cited in support, although rarely in isolation, for the abstractionist position. It is clear that a memory-based model makes exactly the same predictions, for trivial reasons. Making abstractionist and memory based models distinguishable implies in all cases the assessing of learning through a transfer procedure in which subjects are exposed to a new situation.

The Kushner, Cleeremans, and Reber (1991) study illustrates a case in which transfer consists in changing or suppressing the rules underlying the study material. This procedure has been used in a few studies on artificial grammar learning (e.g., Reber 1969) and, on a larger scale, on sequential pattern learning (e.g., Lewicki, Hill, and Bizot 1988). In all of these studies, this procedure results in a sharp drop in subjects' performance, which is taken as evidence for implicit rule knowledge. It is worth noting that a pure memory account again predicts the same outcome. This occurs because changing the rules also changes the product of the rules that subjects, supposedly, memorize. The decrease in performance when rules are changed must not be confused with a genuine transfer effect, which involves the reverse: (relative) stability of performance when the same rules apply to different material. All the studies reviewed below involved genuine transfer paradigm.[2]

Transfer to New Material with Similar Surface Features In most studies including a transfer procedure, the surface features of the stimuli are left unchanged between study and test. In artificial grammar learning, for example, transfer items are typically new strings made up of the same letters as the study strings. Subjects have been shown to perform above chance in assessing the grammaticality of these items.

At first glance, only an abstractive mechanism is able to account for this result, given that the test items were not previously displayed. However, the experimental data may still be accommodated within a memory-based model. One possibility has been put forward by Brooks (1978), who reasons that subjects perform the grammaticality task by assessing the degree of similarity of new test items with the stored representation of a specific study item. Grammaticality and similarity to specific items, Brooks argues, are confounded in typical procedures, so that judging from similarity leads to mimicking performances obtained by applying grammatical rules to the new material. Subsequent studies (McAndrews and Moscovitch 1985; Vokey and Brooks 1992) demonstrated that when the test material was devised to make grammaticality and specific similarity independent, a substantial part of the variance in grammaticality judgments was indeed accounted for by the similarity of test items with a specific study exemplar. However, another part of the variance was linked to the genuine grammaticality of test items. These results led Vokey and Brooks (1992) to conclude in favor of a dual model, reserving a place for abstractive processes in addition to memory-based mechanisms.

My colleagues and I explored a somewhat different way of accounting for transfer data (Perruchet, Gallego, and Pacteau 1992; Perruchet and Pacteau 1990). Our framework is close to Brook's in that grammaticality judgments also rely on memory for specific events, but the events of interest no longer match what are the logical units from the experimenter's standpoint. The level of analysis is shifted from whole exemplars to small chunks of letters such as bigrams or trigrams. These chunks typically vary in frequency across the set of items. In keeping with the ubiquitous laws of memory, the strength of the memory trace is thought to be sensitive to this frequency of occurrence. The sensitivity to frequency information provided by our model is crucial, because it can be shown that the frequency of occurrence of chunks conveys part of the structural information about items. Hence, the simple and automatic effect of occurrence frequency of chunks on performance mimics the effects of knowing and applying the abstract rules underlying the material.

This change in the size of the basic units of knowledge has dramatic consequences on the explanatory power of memory-based accounts of artificial grammar learning. I have shown elsewhere (Perruchet, n.d.) that taking trigrams as the basic unit of knowledge explains within a memory-based framework the variance in performance that Vokey and Brooks (1992) failed to account for with a memory-based model that posited the primacy of item knowledge (and that they consequently attributed to genuine abstraction processes). In addition to converging results from other laboratories (Dienes, Broadbent, and Berry 1991; Servan-Schreiber and Anderson 1990), these data lead to the conclusion that transfer on items built from the same letter set as the study items does not provide evidence for abstraction in the area of artificial grammar.

Studies in sequence pattern learning (Cleeremans and McClelland 1991) and control process situations (Marescaux, Dejean, and Karnas 1990) exhibit remarkable convergences with the conclusion drawn from artificial grammars.

Transfer to New Material with Dissimilar Surface Features The rationale for changing the surface features of the material between the study and transfer phases is straightforward. This procedure makes irrelevant an account for transfer in terms of components common to study and test material. Hence, the occurrence of positive transfer should ensure that performance is mediated by the knowledge of abstract rules. However, the experimental literature consistently shows that transfer on superficially different, although structurally similar, material is, at best, an uncommon phenomenon.

Studies in the grammar learning area (Brooks and Vokey 1991; Mathews et al. 1989; Reber 1969; Whittlesea and Dorken 1993) support this assertion, despite the optimistic claims of most of their authors. In these studies, the letters making up the items were changed in a consistent way between the study and the test phases, with the structure of the grammar left unchanged.

The proportion of correct responses in the changed letter set condition was typically found to be around .55, a proportion higher than the .50 correct proportion corresponding to random responding.

Unfortunately, none of these studies included a genuine control group receiving no prior training with structured strings, which is needed to assess the occurrence reliably, and the magnitude of transfer effects observed in the experimental group (see Perruchet, n.d., for a more detailed criticism). Even if further studies with adequate control showed some transfer effects, it remains to be proved that these effects stem from implicit processes. One possibility is that subjects engage in a controlled search for abstract rules during the study phase, even though there was no explicit request. In line with this hypothesis, Whittlesea and Dorken (1993) report that any trend for transfer to a new letter set disappears when the study strings are presented as distractors devised to prevent rehearsal in what is disguised as a number learning experiment. Another possibility stems from the fact that subjects at the beginning of the test phase are typically informed that the strings they saw during study were generated by a complex set of rules and that they have to assess the well-formedness of new items with regard to these rules. These instructions inevitably shift subjects to a rule discovery mental set. Thus, abstractive operations on the recalled representation of study items may be performed at this time.

The findings are still more clear-cut in other subareas of research. Total failure to obtain transfer to new material with dissimilar surface features is the rule in studies on serial reaction time tasks (Stadler 1989; Willingham, Nissen, and Bullemer 1989), as well as in studies on control process tasks (Berry and Broadbent 1988; Squire and Frambach, 1990).

Contrasting Predictions from Abstractionist and Memory-based Models

The literature reviewed up to now offers no straightforward evidence that a genuine abstraction process mediates performance in implicit learning paradigms, insofar as a memory-based interpretation conveniently accounts for the same data. However, both models have the status of alternative accounts. A memory-based model may be preferred for its economy or some other extraneous features, but the validity of an abstract standpoint remains intact. The search for a stronger test implies devising procedures aimed at contrasting the predictions of the two models. The experiment presented at the beginning of this chapter illustrates such a procedure. Recall that findings unambiguously supported the memory-based model.

Following the same strategy, my colleagues and I (Perruchet, Gallego, and Savy 1990) conducted further analyses of the sequence learning situation initially explored by Lewicki, Hill, and Bizot (1988). We pointed out that an interpretation positing that subjects were sensitive to the frequency of occurrences of small components of the study sequence accounted for the rough performance pattern as well as Lewicki, Hill, and Bizot's interpretation that subjects abstracted the generative rules of the sequence. However, we also derived from the two models a set of contrasted predictions pertaining to specific features in the fine-grained pattern of subjects' performance (Perruchet, Gallego, and Savy 1990, table 1). The empirical data exhibited a strikingly good fit with our predictions, while the predictions derived from the abstractionist framework were clearly discomfirmed. This result was due again to the acknowledgment of the primacy of small chunks of trials over the "logical" unit of the task.

To recapitulate, a memory-based model taking small fragments of stimulus material as basic units of knowledge not only provides an economical alternative account to the abstractionist model; the empirical data clearly lend support to the former when predictions based on the two models are experimentally contrasted.

Does a Model Based on Memory for Small Units Really Exclude Abstraction?

The preceding analysis shows to what extent the explanatory power of memory-based models is enhanced when psychological units are conceived of as fragments of physical or logical units. The formation of these subunits may be described in terms of abstraction from larger units. Similarly, the effect of the occurrence frequency of these small units may be described as the abstraction of the most relevant information from the whole initial data.

Abstraction as a descriptive concept, however, must not be confused with abstraction as a psychological process. For the sake of illustration, consider the status of the fragments of letters such as bigrams or trigrams in artificial grammars. These chunks may be described as abstracted from a string of letters. Similarly, strings of letters are abstracted from the list of strings, and letters are abstracted from chunks. This nested listing may be lengthened at both ends. Any event may be considered as abstracted from its context. This makes it clear that the descriptive definition of abstraction is psychologically irrelevant.

The term *abstraction* refers in this chapter to analytical operations performed to extract componential information from a primitive entity. Considering chunks as abstracted from strings in this psychologically relevant way is tantamount to confounding the experimenter's units with the subject's units. In fact, a number of arguments support the idea that in the standard conditions of training used in artificial grammar paradigms, subjects primarily encode fragments of items such as bigrams or trigrams (for arguments, see Perruchet and Pacteau 1991; Perruchet, n.d.). The resulting primacy of fragments over the whole exemplars in subjects' representions is beneficial to subsequent grammaticality judgments, but no special processing linked to the grammatical structure of the study items needs to be assumed. Chunking would occur also with random sets of letters, because it depends on fundamental and ubiquitous properties of the perceptual system.

The same line of reasoning applies to frequency effects. As argued by Mathews (1991), forgetting infrequent units can abstract knowledge. When forgetting infrequent units occurs on structured material, its consequences on grammaticality judgments parallel the effects of rule-based processing. However, the psychological process at hand is forgetting, not abstraction.

These remarks lead to the tentative characterization of the memory-based model put forward here as a model that involves only the perceptual and memory processes that are engaged and exert observable consequences upon performance in all situations, whether they are structured or random. Claiming that implicit learning in rule-governed environments is based on memory-based processes means that no processes are engaged other than those recruited in, for instance, list learning or other "pure memory" paradigms. The fact that engaging these simple and ubiquitous processes in rule-governed situations generates performances that generally parallel those expected from subjects fully informed of the rules leads to an emphasis on their powerful adaptive value but has no implications with regard to the intrinsic sophistication of these processes.

Summary

Thus, the entire body of data collected in the implicit learning area can be encompassed within a memory-based framework positing the primacy of small functional units of knowledge. More important, the results lend support to this framework when the experimental design contrasts its predictions with those derived from the abstractionist standpoint. These findings suggest that the simple and ubiquitous processes recruited in any memory paradigm are a far more general and efficient way of coping with structured environments than was previously acknowledged, thus extending the conclusion reached in other areas of research dealing with simpler pattern stimuli to very complex situations.

32.4 IMPLICATIONS FOR THE IMPLICITNESS OF PROCESSING

Such a radical change in our conception about how humans learn from rule-governed environments has major consequences for the issue of implicitness. Reber (1989, p. 219) defines implicit learning by three criteria: "(a) Implicit learning produces a *tacit knowledge base* that is abstract and representative of the structure of the environment; (b) such knowledge is *optimally acquired independently of conscious effort* to learn; and (c) it can be *used implicitly* to solve problems and make accurate decisions about novel stimulus circumstances" (emphasis added). Although there is no evidence for Reber's requirement for an abstract knowledge base, these criteria can still serve to define implicitness, although the knowledge base subtending performance improvement is conceived of as a pool of specific events rather than a set of abstract rules.

Is the Knowledge Base Unavailable to Conscious Reflection?

Studies in the abstractionist framework naturally focus on knowledge of the rules underlying the situation and, unsurprisingly, fail to reveal any conscious knowledge of these rules. Studies aimed at investigating memory for specific items or fragments of items provide a totally different picture.

Subjects typically exhibit a large amount of specific knowledge when tests are devised to tap this feature. Convincing demonstrations that the pieces of knowledge available to consciousness are sufficient to account for performance improvement have been made in the artificial grammar learning area (Dienes, broadbent, and Berry 1991; Druhan and Mathews 1989; Dulany, Carlson, and Dewey 1984; Perruchet and Pacteau 1990), as well as in process control situations (Marescaux and Karnas 1991; Sanderson 1989).

In the sequence learning area, the early evidence for unconsciousness was also dismissed when the level of inquiry was shifted from the knowledge of rules to the knowledge of specific sequences. For instance, Perruchet, Gallego, and Savy (1990) showed that the demonstration provided by Lewicki, Hill, and Bizot (1988) falls short when the measure of consciousness no longer bears on the abstract two-order dependency rules generating the sequences but rather on short segments of the sequences (Cleeremans, 1993 Cleeremans and McClelland 1991).

In some respects, the same analysis can be made on the Nissen and Bullemer (1987) experiments and on studies patterned after the same paradigm. In these studies, the sequences are not the product of abstract generative rules but nevertheless embed some kind of regularity. Indeed, a single, arbitrary sequence is repeated throughout the training session. Several authors (Cohen, Ivry, and Keele 1990; Nissen, Willingham, and Hartman 1989; Willingham, Nissen, and Bullemer 1989) claimed that in order to gain place reaction times improve with repetition of this sequence without subjects acquiring conscious knowledge of the sequence (or at least before they do so). Most of this evidence is again grounded on a confounding between the logical units of the task: namely, the whole repeated sequence and the psychologically relevant units. For instance, in the studies above cited postexperimental interviews were devised to assess if subjects realized that they saw the repetition of the same sequence, which is generally ten trials long. In a still more recent experiment, Willingham, Greeley, and Bardone (1993) assessed explicit knowledge of their sixteen-trial sequences through a recognition test in which they present the old sixteen-trial sequence mixed with new sixteen-trial sequences. The problem is that subjects can learn part of the sequence before noticing that the same sequence is continuously repeated or before being able to recognize the whole repeated sequence. Amorim and I (Perruchet and Amorim 1992) showed that when subjects were instructed to generate a sequence that looked like the one they had encountered in the training phase or were asked to recognize small fragments of this sequence, reliable explicit knowledge of salient fragments was revealed after an amount of practice that was hardly sufficient to improve mean motor performance. Importantly, the trials within the sequence on which reaction times were the lowest were the ending elements of the best-recalled or best-recognized chunks of trials.

Overall, shifting the object of inquiry from the abstract rules or from the logical units of the task to fragmentary specific knowledge leads to the conclusion that the knowledge base underlying performance improvement in implicit learning paradigms is normally available to consciousness.[3]

Is the Knowledge Base Optimally Acquired without Conscious Effort to Learn?

Obviously, Reber's criterion refers tacitly here to conscious effort to learn the rules. In a memory-based framework, the lack of a positive effect of this strategy would call for a straightforward interpretation: rules are useless for performance improvement. But empirical data reveal not only a lack of positive effect but a detrimental effect of rule-searching instructions. One explanation may be that omitting instructions for rule searching facilitates memory for specific study items. This would occur because the search for rules diverts subjects from paying full attention to the specific items. The instructions used by default, which typically call for rote memory of items,[4] are well suited to promote acquisition of specific knowledge and hence should be associated with better performance.

Is the (specific) knowledge base optimally acquired without conscious effort to learn exemplars? The evidence about the beneficial effect of rote memory instructions on performance is in line with a negative response. However, focusing on the role of conscious effort to learn may be misguided. The actual causal factor is more probably the nature of the processing that such a strategy elicits. Attempting to memorize the items would be efficient because this strategy leads to the allocation of attentional processing to the relevant features of the material. But any orienting task eliciting similar processing would be efficient as well, even if subjects are not conscious that they are learning. These kinds of ideas are reminiscent of the framework that evolved in the memory area during the past two decades around the notions of encoding specificity or transfer-appropriate processing. Whittlesea and Dorken (1993) have cogently argued for applying this framework to the artificial grammar learning area. In support of their so-called episodic processing account, they show through several experiments that performance in transfer tasks is altered in consonance with transfer-appropriate processing principles by manipulating the specific demands of the study task.

Is the Knowledge Base Used Implicitly?

In the abstractionist point of view, the implicit use of knowledge is a straightforward consequence of its alleged unconscious status. In the context of a memory-based framework, the availability in consciousness of fragmentary specific knowledge makes it possible for subjects to use this knowledge deliberately to cope with the test situation. This hypothesis has been advocated by Dulany, Carlson, and Dewey (1984) and endorsed by Perruchet and Pacteau (1990) in the artificial grammar area. In the standard procedure, subjects make their well-formedness assessment of test strings while they are fully informed of the rule-governed nature of the situation and without time pressure. These conditions make it likely that subjects engage in deliberate use of all the information they can explicitly retrieve about the study strings. However, what happens if such deliberate control processing is impeded, or at least not explicitly prompted by, the procedure?

In an early study on artificial grammar learning, Reber (1967) observed that subjects instructed to learn the letter strings by rote performed better when the strings were generated by the grammar than when strings were generated randomly, although they were never informed about the nature of the strings. Similarly, in sequential pattern learning paradigms, subjects are never instructed about the structure of the sequence. In addition, they have to respond as fast as possible to successive targets. These conditions make it unlikely that they engage in deliberate strategies of responding based on the explicit representation of the sequence, at least during the early stage of training. Introspective subjects' reports support this assertion. Nevertheless, reaction times are sensitive to the structure of the sequence after remarkably little practice. Moreover, the improvement in performance still occurs when subjects perform the task under attentional distraction (Cohen, Ivry, and Keele 1990; Nissen and Bullemer 1987; Perruchet and Amorim 1992).

These data suggest that performance improvement is at least partially independent of the deliberate use of retrievable information. However, this independence remains debatable, given that postexperimental tests consistently show that the knowledge base underlying performance improvement is available to normal subjects' awareness upon explicit request.

Studies conducted with patients who have no conscious access to the knowledge base as a result of neurological impairements provide a unique opportunity to assess the role of conscious knowledge on performance improvement. Knowlton, Ramus, and Squire (1992) submitted amnesic patients with various etiologies to a standard artificial grammar learning paradigm. They reported that patients performed as well as normal subjects on the grammaticality task, although they were impaired on a recognition test of letter strings. Likewise, several pattern sequence learning studies conducted with individuals with Korsakoff's syndrome reported that reaction times improved with training (Nissen and Bullemer 1987), even with a 1-week interval between study and test (Nissen, Willingham, and Hartman 1989), although patients lacked the corresponding explicit knowledge. Normal improvement in performance despite a severe impairment of explicit representations has also been found in patients with amnesia resulting from Alzheimer's disease (Knopman and Nissen 1987). Overall, these findings clearly run counter to the idea that changes in performance are necessarily mediated by the aware knowledge of the components of the situation in which these changes are observed. At least some of these changes could originate from the automatic use of memory (to borrow Jacoby's terms, e.g., Jacoby, chap. 26, this volume).

Of course, performance usually can also be influenced by conscious, deliberate use of memory. Some insight about the dual influence of automatic and intentional processes on performance is provided in the Knowlton, Ramus, and Squire (1992) study. This study reported no difference in grammaticality judgments between amnesic patients and control subjects, but this result was observed only when judgments were performed on the basis of subjective feelings. A difference emerged when subjects were encouraged to use their conscious memory of study strings. Performance of normal subjects, who presumably have conscious knowledge of some fragments of the study strings, benefited from this change in instructions, demonstrating that intentional use of memory can assist grammaticality judgments (see also Willingham, Nissen, and Bullemer 1989 for similar evidence in sequence learning). By contrast, performance of amnesic patients, who presumably have inadequate conscious knowledge of the study material, tended to deteriorate in the same conditions. Although it would be premature to propose a complete interpretation of these results, they suggest that automatic and conscious uses of memory interact in complex ways, according to laws that warrant further empirical and theoretical research.

This research should benefit from increased proximity with research in memory. The rejection of the abstractionist model of learning leads to drawing a close parallelism between the operations involved when subjects are faced with a rule-governed environment and with a set of unrelated items, as in typical memory paradigms. Studies on implicit memory as assessed by repetition priming tests (Moscovitch, chap. 25, this volume; Roediger 1990; Schacter 1987) teach us a fundamental principle: that a specific past event can exert an influence on subsequent performance without this event being explicitly remembered. In this domain, the joint influence of deliberate and automatic processes has generated an abundant literature. For instance, several theories posit that recognition partly relies on subjective familiarity with the test items, with familiarity being construed as the automatic consequence of the earlier processing of the same material (e.g., Mandler 1990). Conversely, the possible influence of explicit retrieval on performance in nominally implicit tests of memory also has be acknowledged (e.g., Schacter 1987). There is now a large consensus around the idea that there are no process-pure memory tasks (Jacoby 1991). Presumably, similar developments in the context of learning in complex situations are warranted.

Summary

Experimental data provide striking evidence that an amount of specific knowledge sufficient to account for performance improvement is available to conscious reflection upon explicit request, at least when normal subjects are submitted to standard experimental conditions. This knowledge is closely dependent on the operations engaged in by the subjects to satisfy the demands of the study task. However, these results do not entail that performance is causally determined by the intentional use of this knowledge. Introspective evidence in normal subjects suggests, and studies on neurologically impaired patients confirm, that changes in performance can occur without the involvement of concious, intentional processes, as in typical implicit memory tests.

32.5 THE FUNCTIONS OF NONCONSCIOUS AND CONSCIOUS PROCESSES

On the basis of the empirical evidence put forward so far, this section is devoted to more speculative developments on the function of nonconscious and conscious processing. I argue that the new perspective leads us to differentiate more clearly the role of nonconscious processes from the role of conscious processes than was the case in the conventional, abstractionist framework.

Conceiving Nonconscious Processes in the Model of Conscious Processes

Anyone attempting to learn from a new structured stimulus pattern invariably engages in a deliberate controlled strategy that searches for deep regularities in order to store in memory a small number of ready-to-use abstract laws and apply them when required by the situation. When researchers observed that subjects' performance improved in complex situations even though these operations were not intentionally engaged and available to retrospective awareness, they naturally hypothesized that the same operations were performed by an autonomous, unconscious processor. The only acknowledged differences between the two modes of processing pertain to their power (with nonconscious processing viewed as more powerful than conscious processing) and their rapidity (with nonconscious processing viewed as faster than conscious processing). As claimed by proponents of this view, "Our conscious thinking needs to rely on notes (with flowcharts or lists of if-then statements) or computers to do the *same job* that our nonconsciously operating processing algorithms can do instantly and without external help" (Lewicki, Hill, and Czyzewska 1992, 798; emphasis added).

Postulating unconscious activity in direct analogy with human conscious processing is by no means specific to this area of research. This view is standard in domains such as cognitive ethology, linguistics, psychoanalysis, formal models in artificial intelligence, and the information-processing approach in cognitive psychology.

A Crucial Difference

This standard view has been challenged in several ways. Searle (1990) has forcefully argued that this view anthropomorphizes the nonconscious processes in the brain. Dulany (1991) remarks that even Freud was concerned about this aspect, as shown by the following quotation: "The psychoanalytic assumption of unconscious mental activity appears to us, on the one hand, as a further expansion of the primitive animism which caused us to see copies of our own consciousness all around us" (Freud [1915] 1957, 171).

The framework outlined in this chapter leads to hypothesizing qualitative, and not only quantitative, differences between nonconscious and conscious processing. The evidence reviewed suggests that nonconscious processes are operative in the use of specific knowledge, while conscious processes can deal with both specific and abstract knowledge.

Reintroducing abstraction here may seem paradoxical, insofar as the general stance of this chapter runs counter to an abstractionist standpoint. The paradox, however, is only apparent. The memory-based model of learning put forward here is intended to be relevant only in usual paradigms of implicit learning, that is, when subjects are not prompted for intentional rule searching or, more generally, for abstract processing of the study material. Obviously other forms of knowledge can be acquired under other conditions. For instance, Turner and Fischler (1993) and Wittlesea and Dorken (1993) show that subjects can learn abstract rules of the grammar underlying strings of letters when they are instructed to search for rules, or when they are given incidental instructions orienting toward the structure of the material (see also Shanks and St. John in press, for additional evidence for abstraction in other areas of experimental research). At a more general level, the very existence of sciences such as logics, linguistics, and physics also testifies to the human ability to abstract the structure of complex environments. The point of this chapter is not that learning from complex situations never involves abstractive operations but that these operations require explicit thinking for them to be performed.

These claims about the specific properties of nonconscious and conscious processing are intrinsically speculative. Indeed, demonstrating that the unconscious is unable to perform some operations goes beyond any experimental inquiry; it can be shown only that the reverse proposal—the unconscious is endowed with such abilities—has no empirical support.

Is the Unconscious Smart or Dumb?

In the target paper of a set of contributions introduced under this title (Loftus and Klinger 1992) in the *American Psychologist*, Greenwald (1992) argued for an unconcious that would be far less sophisticated than commonly thought. This chapter unambiguously supports this standpoint.

The Greenwald claim, however, could be misguiding, insofar as the lack of sophistication is usually associated with the lack of power. Therefore, it is worth emphasizing that the analysis in this chapter leads to conceive the unconscious as both "dumb" and considerably powerful. Indeed, by and large, the reappraisal presented in this chapter consists in changing the nature of the adaptive operations performed without conscious awareness, without questioning their end product. These operations have limited sophistication insofar as they consist essentially in benefiting from the earlier processing of similar specific events. But they are performed with considerable efficiency, presumably due to a massively parallel mode of processing. In fact, they almost perfectly simulate the action of a sophisticated processor.

32.6 ABOUT EVOLUTION THEORY

Nonconscious processes are conceived in this chapter as oriented toward the exploitation in parallel of a large database of specific pieces of knowledge to satisfy immediate demands. Conversely, abstract operations are linked with conscious and controlled thought. This conception corresponds to a profound change in perspective with regard to the conventional position (e.g. Reber 1989), in which abstract knowledge is acquired and used by an unconscious, autonomous processor.

Recently Reber (1992) framed his position within a functionalist perspective, anchored within evolutionary biology. He forcefully argues that implicit systems ought to be conceived of as fundamental in adaptation, given that consciousness is a late arriver in ontogenetic and phylogenetic development. In keeping with this claim, he puts forward a set of observations leading support to the "primacy of the implicit" (Reber 1990). I entirely agree with this perspective.

Reber's line of reasoning, however, provides no support for his view concerning the nature of implicit processes. His contention—that viewing abstraction as a primitive and fundamental adaptive process is consonant with the basic tenets of evolutionary biology—in fact uses the claim that implicit learning generates abstract knowledge as one of its premises. Introducing evolution theory here provides no additional support for this premise. If it is replaced, as I think it must be, by "implicit learning engages memory-based processes only," then the conclusion that memory-based processes are primitive and fundamental adaptive processes follows as well.

It may be argued further that the position advocated in this chapter is closer to evolution theory principles than usual on at least two key points. The first pertains to the unrealistic view that species low in the phylogenetic scale are endowed with sophisticated processing systems. A memory-based model of implicit learning rests on far simpler (although equally powerful) processes, which are easier to account for by elementary neurophysiological mechanisms.

The second point may be more important. Reber's framework gives no place to consciousness in the learning process. Consciousness is construed as a simple epiphenomenon, when its intervention is not perceived as detrimental to efficient learning (Reber 1989). This conception is undermined by a serious flaw. As Dulany (1991) noted, "We may wonder why consciousness evolved at all if it is the poor thing it is often said to be" (p. 101). The recency of the emergence of consciousness in phylogenetic history must not overshadow the fact that this emergence has apparently been accompanied by a major improvement in adaptive abilities. The framework outlined here accounts for this feature by construing consciousness as a prerequisite for analytic reasoning and genuine abstraction.

NOTES

Preparation of this chapter was supported by the Centre National de la Recherche Scientifique and the Université René Descartes, Paris, France. I thank D. Berry, J. Gallego, J. Lautrey, and one anonymous reviewer for their helpful comments on an early version of the chapter.

1. For the three rules, *t*-tests compared the observed proportion of responses falling into one of the two issues predicted by the memory model with .66. For two of these rules (clockwise and counterclockwise relations), the probability of selecting one or another of these issues substantially differed. This trend was not anticipated by any of the competing models of the task, and its interpretation is not clear.

2. For alleged evidence of abstraction that relies on another, somewhat complex procedure, and its reinterpretation within a nonabstractionist framework, see Perruchet, Gallego, and Pacteau (1992).

3. For the sake of brevity, I pass over in silence the problems surrounding the nature of the tests devised to assess awareness, and especially the problems related to the sensitivity of these tests. Priority is given in this chapter to *what* to measure, rather than to *how* to measure. For an analysis including other aspects, see the overview by Shanks and St. John (N.d.).

4. In artificial grammar learning, instructions asking for rote learning are termed implicit, and instructions asking for rule searching are termed explicit. These labels are appropriate to the abstractionist standpoint. However, they have to be reversed to be consonant with a memory-based framework. If performance improvement is due to the memory for specific events, instructions stressing rote learning would be more correctly called explicit, and conversely, rule-searching instructions become implicit in nature.

REFERENCES

Allen, S. W., and Brooks, L. R. (1991). Specializing in the operation of an explicit rule. *Journal of Experimental Psychology: General, 120,* 3–19.

Berry, D. C., and Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology, 79,* 251–272.

Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch and B. B. Lloyd (Eds.), *Cognition and categorization,* 169–215. Hillsdale, NJ: Erlbaum.

Brooks, L. R., and Vokey, J. R. (1991). Abstract analogies and abstracted grammars: A comment on Reber, and Mathews et al. *Journal of Experimental Psychology: General, 120,* 316–323.

Cleeremans, A. (1993). *Mechanisms of implicit learning: A connectionist model of sequence processing.* Cambridge, MA: MIT Press.

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

Cohen, A., Ivry, R. I., and Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 17–30.

Dienes, Z., Broadbent, D., and Berry, D. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 875–887.

Druhan, B. B., and Mathews, R. C. (1989). THIYOS: A classifier system model of implicit knowledge of artificial grammars. Paper presented at the annual meeting of the Cognitive Science Society, Ann Arbor, MI.

Dulany, D. E. (1991). Conscious representation and thought systems. In R. S. Wyer and T. K. Srull (Eds.), *Advances in social cognition,* vol. 4, 97–120. Hillsdale, NJ: Erlbaum.

Dulany, D. E., Carlson, A., and Dewey, G. I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General, 113,* 541–555.

Freud, S. [1915] (1957). The unconscious. In J. Strachey (Ed. and Trans.), *The standard edition of the complete psychological works of Sigmund Freud,* vol. 14, 159–215. London: Hogarth Press.

Green, R. E. A., and Shanks, D. R. (1993). On the existence of independent learning systems: An examination of some evidence. *Memory and Cognition, 21,* 304–317.

Greenwald, A. G. (1992). New look 3: Unconscious cognition reclaimed. *American Psychologist, 47,* 766–779.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30,* 513–541.

Knopman, D. S., and Nissen, M. J. (1987). Implicit learning in patients with probable Alzheimer's disease. *Neurology, 37,* 784–788.

Knowlton, B. J., Ramus, S. J., and 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.

Kushner, M., Cleeremans, A., and 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, Chicago.*

Lewicki, P., Hill, T. and Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology, 20,* 24–37.

Lewicki, P., Hill, T., and Czyzewska, M. (1992). Nonconscious acquisition of information. *American Psychologist, 47,* 796–801.

Loftus, E. F., and Klinger, M. R. (1992). Is the unconscious smart or dumb? *American Psychologist, 47,* 761–765.

Mandler, G. (1990). Your face looks familiar but I can't remember your name: A review of dual process theory. In W. E. Hockley and S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock.* Hillsdale, NJ: Erlbaum.

Marescaux, P-J., Dejean, K., and Karnas, G. (1990). *Acquisition of specific or general knowledge at the control of a dynamic simulated system: An evaluation through a static situations questionnaire and a transfer control task.* Report 2PR2GK of the KAUDYTE project (ESPRIT BRA 3219).

Marescaux, P-J., and Karnas, G. (1991). *The implicit versus explicit knowledge distinction revisited: When finding associations between verbalizable knowledge and some performance criteria.* Report 4PR3GK of the KAUDYTE project (ESPRIT BRA 3219).

Mathews, R. C. (1991). The forgetting algorithm: How fragmentary knowledge of exemplars can abstract knowledge. *Journal of Experimental Psychology: General, 120,* 117–119.

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

McAndrews, M. P., and Moscovitch, M. (1985). Rule-based and exemplar-based classification in artificial grammar learning. *Memory and Cognition, 13,* 469–475.

Medin, D. L., and Ross, B. H. (1990). The specific character of abstract thought: Categorization, problem solving, and induction. In R. J. Sternberg (Ed.), *The psychology of human intelligence*, vol. 5, 189–223. Hillsdale, NJ: Erlbaum.

Medin, D. L., and Shaffer, M. M. (1978). A context theory of classification learning. *Psychological Review, 85*, 207–238.

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

Nissen, M. J., Willingham, D., and Hartman, M. (1989). Explicit and implicit remembering: When is learning preserved in amnesia? *Neuropsychologia, 27*, 341–352.

Nosofsky, R. M. (1991). Typicality in logically defined categories: Exemplar-similarity versus rule instantiation. *Memory and Cognition, 19*, 131–150.

Papini, M. R., and Bitterman, M. E. (1990). The role of contingency in classical conditioning. *Psychological Review, 97*, 396–403.

Perruchet, P. (n.d.). On defining the knowledge units of a synthetic language: Comment on Vokey and Brooks (1992). *Journal of Experimental Psychology: Learning, Memory, and Cognition*. In press.

Perruchet, P., and Amorim, M. A. (1992). Conscious knowledge and changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 785–800.

Perruchet, P., Gallego, J., and Pacteau, C. (1992). A reinterpretation of some earlier evidence for abstractiveness of implicitly acquired knowledge. *Quarterly Journal of Experimental Psychology, 44A*, 193–210.

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

Perruchet, P., and Pacteau, C. (1991). The implicit acquisition of abstract knowledge about artificial grammar: Some methodological and conceptual issues. *Journal of Experimental Psychology: General, 120*, 112–116.

Perruchet, P., Pacteau, C., and Gallego, J. (1993). Abstraction of covariation in category learning: A critical note on K. Richardson's studies. *British Journal of Psychology, 84*, 289–295.

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

Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology, 81*, 115–119.

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

Reber, A. S. (1990). On the primacy of the implicit: A comment on Perruchet and Pacteau. *Journal of Experimental Psychology: General, 119*, 340–342.

Reber, A. S. (1992). The cognitive unconscious: An evolutionary perspective. *Consciousness and Cognition, 1*, 93–133.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66*, 1–5.

Rescorla, R. A., and Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black and W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research*, 64–99. New York: Appleton-Century-Crofts.

Roediger, H. L., III (1990). Implicit memory: A commentary. *Bulletin of Psychonomic Society, 28*, 373–380.

Sanderson, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 729–747.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 501–518.

Searle, J. R. (1990). Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences, 13*, 585–642.

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

Shanks, D. R. (1991). Instrumental actions and the acquisition of causal beliefs. Report 9102. University of California, Department of Cognitive Science.

Shanks, D. R., and St. John, M. F. (n.d.). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*. In press.

Squire, L. R., and 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, and Cognition, 15*, 1061–1069.

Turner, C. W., and Fischler, I. S. (1993). Speeded tests of implicit knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1165–1177.

Vokey, J. R., and Brooks, L. R. (1992). Salience of item knowledge in learning artificial grammar. *Journal of Experimental Psychology: Learning, Memory and Cognition, 18*, 328–344.

Wasserman, E. A. (1990). Detecting response-outcome relations: Toward an understanding of the causal texture of the environment. In G. H. Bower (Ed.), *The psychology of learning and motivation*, vol. 26, 27–82. New York: Academic Press.

Whittlesea, B., and Dorken, M. D. (1993). Incidentally, things in general are incidentally determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 22*, 227–248.

Willingham, D. B., Greeley, T., and Bardone, W. (1993). Dissociation in a serial response time task using a recognition measure: Comment on Perruchet and Amorim (1992). *Journal of Experimental Psychology: Learning, Memory and Cognition, 19*, 1424–1430.

Willingham, D. B., Nissen, M. J., and Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15*, 1047–1060.