Implicit Learning and Tacit Knowledge

Arthur S. Reber
Brooklyn College and the Graduate Center
City University of New York

I examine the phenomenon of implicit learning, the process by which knowledge about the rule-governed complexities of the stimulus environment is acquired independently of conscious attempts to do so. Our research with the two, seemingly disparate experimental paradigms of synthetic grammar learning and probability learning is reviewed and integrated with other approaches to the general problem of unconscious cognition. The conclusions reached are as follows: (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 efforts to learn; and (c) it can be used implicitly to solve problems and make accurate decisions about novel stimulus circumstances. Various epistemological issues and related problems such as intuition, neuroclinical disorders of learning and memory, and the relationship of evolutionary processes to cognitive science are also discussed.

Some two decades ago the term *implicit learning* was first used to characterize how one develops intuitive knowledge about the underlying structure of a complex stimulus environment (Reber, 1965, 1967). In those early writings, I argued that implicit learning is characterized by two critical features: (a) It is an unconscious process and (b) it yields abstract knowledge. Implicit knowledge results from the induction of an abstract representation of the structure that the stimulus environment displays, and this knowledge is acquired in the absence of conscious, reflective strategies to learn. Since then, the evidence in support of this theory has been abundant, and many of the details of the process have been sharpened. This article is an overview of this evidence and an attempt to extend the general concepts to provide some insight into a variety of related processes such as arriving at intuitive judgments, complex decision making, and, in a broad sense, learning about the complex covariations among events that characterize the environment.

Put simply, this is an article about learning. It seems curious, given the pattern of psychological investigation of the middle decades of this century, that the topic of learning should be so poorly represented in the contemporary literature in cognitive psychology. The energies of cognitive scientists have been invested largely in the analysis and modeling of existing knowledge rather than in investigations of how it was acquired. For example, in an important recent article on the general topic of unconscious memorial systems, Schacter (1987) never came to grips with the distinction between implicit learning and implicit memory. The latter, the focus of his review, was dealt with historically, characterized, out-

lined, and analyzed, but virtually no attention was paid to the processes by which these implicit memories "got there." This general lack of attention to the acquisition problem may be one of the reasons why much recent theorizing has been oriented toward a nativist position (e.g., Chomsky, 1980; Fodor, 1975, 1983; Gleitman & Wanner, 1982). Failure to explicate how complex knowledge is acquired invites the supposition that "it was there all the time."

What follows is an exploration of implicit learning from the point of view that the processes discussed are general and universal. Implicit acquisition of complex knowledge is taken as a foundation process for the development of abstract, tacit knowledge of all kinds. The stepping-off place is the presumption that there is, at this juncture, no reason to place any priority on particular biological determinants of a specific kind. All forms of implicit knowledge are taken as essentially similar at their deepest levels. This position needs to be pushed as far as it can go; it has considerably more explanatory power than has been generally recognized.

Experimental Procedures

Research on implicit learning is properly carried out with arbitrary stimulus domains with complex, idiosyncratic structures. In order to obtain insight into a process such as implicit learning, it is essential to work with novel, synthetic systems and to focus on the capacity of one's subjects to induce knowledge of a deep sort from such stimulus fields. Over the years, a number of different techniques have been used. My colleagues and I have chosen, in our laboratory, to work with two procedures that we have found to be extremely useful: artificial grammar learning and probability learning. The former is well known in the literature and has been used by many; the latter is somewhat obscure but, as will become clear, is an extremely sensitive technique that has provided some intriguing data. It is useful to provide a short overview of each here and to outline the general procedures for its use. Various other techniques that have found their way into the laboratory, such as the various procedures developed by such

workers as Lewicki (see 1986a) and Broadbent (see Broadbent, FitzGerald, & Broadbent, 1986), are introduced later.

Grammar Learning

Figure 1 shows one of the first synthetic grammars used along with the basic types of "sentences" that it can generate. This grammar was first used by Reber (1965, 1967). It is a Markovian system derived from a simpler system that formed the basis of George Miller's *Project Grammarama* (see Miller, 1967) and has subsequently been used to generate the stimuli for a number of other studies (Howard & Ballas, 1980; Millward, 1981; Reber & Lewis, 1977; Roter, 1985). It can be taken as representative of the grammars used in a variety of other experiments.

Although there have been many variations on a theme here, the basic procedure used in these grammar learning studies is to have an acquisition phase, during which subjects acquire knowledge of the rules of the grammar, and a testing phase, during which some assessment is made of what they have learned. Additional details are supplied as follows when needed.

Several points, however, need to be kept clear about these synthetic languages and how they have been used to examine implicit, unconscious cognitive processes. First, they are complex systems, too complex to be learned in an afternoon in the laboratory, as Miller (1967) noted. Miller saw this as a liability, which it is if one wishes to examine explicit concept learning. This complexity, however, should be regarded as a virtue in the current context, for a rich and complex stimulus domain is a prerequisite for the occurrence of implicit learning. If the system in use is too simple, or if the code can be broken by conscious effort, then one will not see implicit processes. Second, the grammars given here are finite-state systems that generate strings of symbols in a left-to-right, nonhierarchical fashion. This fact should not be taken as reflecting any prejudices about the structural underpinnings of natural languages or their acquisition. We elected to use finite-state grammars for several reasons independent of the-
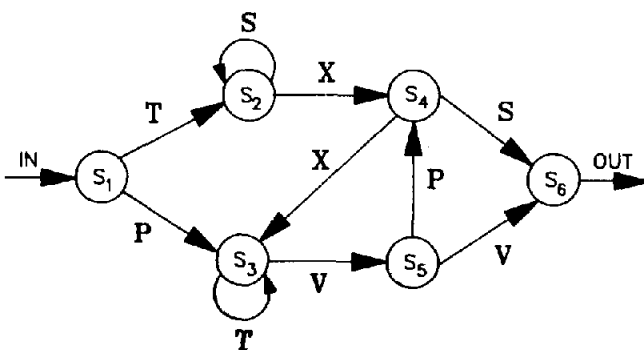
oretical issues in linguistics or natural language learning: They are mathematically tractable; they have an intrinsic probabilistic structure that is well known; they can generate a relatively large number of strings to use as stimuli; and, as mentioned, they are sufficiently complex so that the underlying formal structure is not within the conscious memorial domain of the typical subject upon the subject's entering the laboratory. Finally, as will become clear, there is nothing special about these stimulus generators in any interesting psychological sense. The basic components of implicit learning emerge in a wide variety of different empirical settings with a range of different stimulus environments.

Probability Learning (PL)

The format that we adopted departs noticeably from the traditional two-choice procedure in which each trial consists of a "ready" signal, a prediction response, and an outcome event. The procedure used to explore implicit processes derives from the proposition that the essential nature of a PL experiment has little to do with the explicit learning of probabilities of events. Rather, what passes in the literature for probability learning is actually a much more subtle process in which subjects learn implicitly about the stochastic structure of an event sequence to which they have been exposed. In the course of making predictions, they mimic its structure and thereby generate a sequence of responses, one by-product of which is an approximate matching of the probabilities of the events—that is, *probability learning*.

Accordingly, the PL procedure was modified as follows (Reber, 1966; Reber & Millward, 1968). The subject begins an experimental session simply by observing the occurrence of a sequence of rapidly presented events. There is no ready signal, and the subject makes no prediction responses. In this situation, a passively observed event is functionally equivalent to a traditional trial, and a learning session consisting of 2 or 3 min of observing events at a rate of two per second is sufficient to put a subject at an asymptotic rate of responding; that is, subsequent prediction responses made by subjects who have had this learning experience show all of the characteristics of ordinary subjects who have had an equivalent number of traditional trials. We dubbed this procedure the *instant asymptote* technique.

The typical experiment with this modified PL procedure consists of an acquisition phase, during which subjects observe event sequences that may have any of a variety of stochastic structures, and a testing phase, during which subjects make prediction responses. As with the grammar learning studies, there are many variations on this basic theme; they are introduced later as needed.

Despite the many superficial differences between the probability learning paradigm and the grammar learning experiment, there are two essential commonalities. First, in both cases the subject is confronted with a stimulus environment about which knowledge must be acquired in order to respond effectively during the testing session. Second, neither of the structural systems being used is part of or even remotely similar to the epistemic contents of the typical subject's long-term memory. These points are fundamental; the whole pur-

Figure 1. Schematic diagram of a finite-state grammar. (Stimuli are generated by following any path of arrows leading from the initial State 1 to the terminal State 6. The following are the five basic strings of the grammar with the loops or recursions in brackets: 1. T[S]XS; 2. T[S]XX[[T]VPX]VV; 3. T[S]XX[[T]VPX]VPS; 4. P[[T]VPX]VV; 5. P[[T]VPX]VPS.)

pose of examining implicit learning in the laboratory is to develop understanding of how rich and complex knowledge is initially obtained independently of overt, conscious strategies for its acquisition. This process is ubiquitous in human experience and yet, as a focus of psychological inquiry, it is largely ignored.

What follows is a brief review of our work with these two procedures presented in the form of basic issues that the literature addresses. These issues are integrated with the growing body of literature on unconscious, nonreflective, implicit processes and, finally, summed up in a systematic attempt to see how such cognitive systems could evolve and how they fit into contemporary struggles with a number of classic problems in pure and applied psychology.

The work in our laboratory that is the focus of this overview (beginning with the very first studies by Reber, 1965, 1966) was carried out with a particular research strategy: to use a limited number of techniques to examine a wide variety of effects. The virtue of this approach is that by developing a few techniques and building a robust data base, one can explore a large number of issues and not be terribly concerned about the vagaries that get introduced with alternative procedures. Given that the problems of implicit learning and tacit knowledge can be explored through these two procedures, this heuristic says that they should be used in as many circumstances as makes scientific sense. Biological fans of *E. coli* will recognize this strategy.

There is, of course, an alternative strategy: to examine these nonconscious cognitive and perceptual processes in as wide a variety of experimental environments as possible. The virtue of this strategy is that one is not likely to be seduced by idiosyncratic properties of particular procedures; generalizations come easier. If implicit learning is real, it should emerge in contexts conceptually remote from synthetic grammars and structured event sequences. Ideally, both strategies should be carried out. As will become clear, those whose research programs have taken the latter tack, such as Lewicki (1985, 1986a, 1986b; Lewicki, Czyzewska, & Hoffman, 1987), have typically produced data that are congruent and complementary.

Empirical Studies of Implicit Learning

On the Exploitation of Structure

When a stimulus environment is structured, people learn to exploit that structure in the sense that they come to use it in order to behave in a relevant fashion in its presence. This proposition seems noncontroversial as a generalization about human cognition; in fact it lies at the core of several approaches to perception (Gibson, 1966; 1979; Mace, 1974), decision making, and information processing (Garner, 1974; Hasher & Zacks, 1984) and can be seen as underlying, in a broad sense, any of a number of general theoretical analyses such as Anderson's (1983) production systems, Nelson's (1986) and Schlesinger's (1982) models of natural language acquisition, Fried and Holyoak's (1984) model of category induction, Lewicki's (1986a) analysis of socialization, and,

interestingly, Holland, Holyoak, Nisbett, and Thagard's (1986) model of induction and Rumelhart and McClelland's (1986) connectionist systems. The latter two are of special importance because each can be seen as advances in the cognitive sciences' hopefully awakened interest in knowledge acquisition.

In an early study on this general problem (Reber, 1967), subjects were shown to become interestingly sensitive to the constraints of a synthetic grammar simply from exposure to exemplary strings. In that experiment, subjects were not informed that they were working with rule-governed stimuli. They were merely requested to memorize strings of letters in what was touted as a rote memory experiment. With practice they became increasingly adept at processing and memorizing strings, whereas a control group working with nonordered letter strings showed no such improvement. Furthermore, after this neutral learning task, subjects were able to use what they had apprehended of the rules of the grammar to discriminate between new letter strings that conformed to the grammatical constraints and letter strings that violated one or more of the rules of the grammar. In simplest terms, these subjects can be said to have been exploiting the structure inherent in the stimulus display. This basic finding is a robust one and has been reported by numerous authors (e.g., Brooks, 1978; Dulany, Carlson, & Dewey, 1984; Howard & Ballas, 1980; Mathews et al., in press; Millward, 1981; Morgan & Newport, 1981).

Similar observations concerning the exploitation of structure have been made in somewhat different contexts by other researchers. Broadbent and his co-workers showed that knowledge of complex rule systems governing simulated economic/production systems is also acquired and used in an implicit fashion (Berry & Broadbent, 1984; Broadbent & Aston, 1978; Broadbent et al., 1986). In those studies, subjects were presented with an imaginary manufacturing situation such as a sugar production plant and instructed simply to maneuver variables such as wages, labor peace, worker output, and the like to yield a particular overall production standard. The systems, in fact, operate according to sophisticated, complex rule systems that relate these factors to each other. Achieving the required production standards requires that the rules be "known," in some sense of that word. Broadbent and his colleagues consistently reported that subjects induced the rule systems implicitly and made appropriate adjustments in the relevant variables and did so in the absence of conscious knowledge of the rules themselves. The pattern of these findings strongly parallels those in the synthetic grammar learning studies.

Several of the PL studies have also yielded analogues of this process. Reber and Millward (1971) found evidence that subjects can accurately anticipate the changing probabilities of events even when the anticipatory response requires an integration of information across 50 preceding events. In this particular case, the probability of each individual event on any Trial *n* was systematically increased and decreased as *n* moved through a period of 50 trials. Subjects were first given 1,000 instant asymptote trials with this sawtooth event sequence and then requested to predict successive events. Under these conditions, rather than shadowing the changing event

probabilities, subjects ultimately learned to anticipate the shifts in the likelihood of events so that their predictions of events rose and fell coincidentally with the actual event sequences. They had learned the underlying structure of the stimulus environment and were capable of exploiting it to direct their choices.

Millward and Reber (1972), using event sequences with short- and long-range contingencies between events, reported an even more impressive ability of subjects to exploit stochastic structures. Subjects were exposed to sequences that contained event-to-event dependencies such that the actual event that appeared on any Trial n was stochastically dependent on the event that had occurred on some previous Trial $n - j$, where $j = 1, 3, 5,$ or 7. Training consisted of several hundred trials of the instant asymptote procedure with the particular stochastic dependency for that session. During the first session, $j = 1$; during the second, $j = 3$; and so forth.

During testing, subjects displayed a clear sensitivity to these dependencies, a sensitivity that reflected an ability to exploit structure that required knowledge of event dependencies as remote as seven trials. What makes this finding interesting is that this capacity appears to be beyond what were found in earlier work (Millward & Reber, 1968; Reber & Millward, 1965) to be limits on explicit recall. In those experiments, subjects were asked to recall which event had occurred on a specified previous trial. Beyond five trials, they were virtually reduced to guessing.

Parallel findings were recently reported by Lewicki and his co-workers (Lewicki et al., 1987; Lewicki, Hill, & Bizot, 1988), who, using reaction time measures, showed that subjects implicitly knew the future location of a stimulus event even though, when they were asked to report explicitly where the event would occur, their performances were no better than chance. Lewicki and his colleagues used a complex rule that specified which of four quadrants of a field would contain a target number. The actual location was based on the pattern of quadrants in which the number had appeared on particular earlier trials. The situation, like the PL studies, was based on an arbitrary relationship between events. With extended practice, subjects showed dramatic decreases in reaction time to respond to the location of the target number. The improvement was clearly due to tacit knowledge of the stochastic relationship and not simply to increased facility with the task. Changing the rules produced an abrupt increase in reaction time, but when, in a postexperimental debriefing session, subjects were given an opportunity to consciously predict the critical quadrants, their performances were no better than chance. An intriguing element in these studies was that all of Lewicki et al.'s (1988) subjects were faculty members in a department of psychology, and all knew that the research they were involved in was oriented toward the study of nonconscious cognitive processes.

Clearly, subjects learn to use the structural relationships inherent in these various complex stimulus domains. No real surprises here. In many ways these various studies function basically as complicated existence demonstrations; they show that it is possible to obtain this kind of unconscious, nonreflective, implicit learning in a controlled laboratory setting, that it can occur in a relatively short time span, and that it can be seen to emerge when the stimulus is a structured domain whose content is arbitrary and distinctly remote from typical day-to-day experiences with the real world.

On Implicit Versus Explicit Processes

The experiments reviewed were all run under instructional sets in which subjects were unaware that the stimuli were structured or rule defined. In these cases, the point was to maximize the emergence of implicit learning.

It is important to be clear about this issue and the kinds of manipulations that have been used. It is universally accepted that the college undergraduate whose cognitive processes form (for better or worse) the foundations of our science is an active and consciously probing organism, especially in regard to things with structure and patterns. The aforementioned researchers carried out their experiments by carefully circumventing this pattern-searching tendency. In Reber and Millward's PL studies and in Lewicki's target location experiments, they accomplished this by "blitzing" the subjects with information at rates beyond those at which conscious code-breaking strategies could operate. The grammar learning studies and Broadbent's production system experiments were successful because the structure of the stimuli was highly complex and the instructions to the subjects were calculated to be vague. An obvious question is, What effect would explicit instructions have? What happens when subjects are informed, at the outset, that the materials that they will be working with reflect regularities and patterns?

The first manipulation of the factor of explicitness used the PL technique (Reber, 1966; Reber & Millward, 1968). The procedure consisted simply of telling some of the subjects exactly what was going on in the experiment. Specifically, by informing one group of subjects of the relative probabilities of the two events, the researchers gave them concrete instructions about the frequency characteristics of the event sequence that they would be asked to predict. These subjects were then run in a standard PL procedure and compared with a control group run with the same event sequences but without the explicit information.

The information about the event frequencies had virtually no effect on behavior. The two groups were statistically indistinguishable from each other, even on the first block of 25 prediction trials in which the impact of the instructions would have been most likely to be felt. Clearly, probability learning is more than the learning of probabilities. Rather, what really goes on is the apprehension of deep information about the structure of the sequence of events.

Postexperimental debriefings were revealing. Subjects were quite clear about knowing which light would be the dominant one, and all said that they believed the instructions. But, in virtually every case, they claimed that somehow the specific information lacked meaning that they felt they could use. It took real experience with the event sequence to acquire a knowledge base that was usable for directing choices on individual trials. This experiment used Bernoulli sequences; there was no "structure" in the usual sense of the word. Nevertheless, subjects reported achieving a sense of the nature of the event sequence from experience with events that they did not derive from the explicit instructions. Of importance

is that this occurred despite the fact that, in principle, there is nothing to be extracted from the event sequence other than the relative frequencies of the two events.

Several studies using the grammar learning procedure have also explored the boundary between implicit and explicit knowledge. In the first experiment, Reber (1976) used the simple device of encouraging one group of subjects to search for the structure in the stimuli while a comparable group was run under a neutral instructional set. Both groups were given the same learning phase, during which they had to memorize exemplars from a synthetic grammar, and an identical testing phase, during which they were asked to assess the well-formedness of novel letter strings. *Well-formedness* is defined by whether the grammar could generate the string. Informed subjects were told only about the existence of structure; nothing was said about the nature of that structure.

The explicitly instructed subjects in this study performed more poorly in all aspects of the experiment than did those given the neutral instructions. They took longer to memorize the exemplars, they were poorer at determining well-formedness of test strings, and they showed evidence of having induced rules that were not representative of the grammar in use. The suggestion is that at least under these circumstances, implicit processing of complex materials has an advantage over explicit processing.

However, as gradually became clear, what this study actually showed is that explicit processing of complex materials has a decided disadvantage in relation to implicit processing. This is no mere play on words. The implicit/explicit distinction is rather more complex than it first appeared. Analysis of the fine grain of the data from Reber's (1976) article revealed that the explicit instructions seemed to be having a particular kind of interference effect. Specifically, subjects were being encouraged to search for rules that, given the nature of finite-state grammars with their path-independent, Markovian properties and given the kinds of attack strategies that the typical undergraduate possesses, they were not likely to find. Moreover, they tended to make improper inductions that led them to hold rules about the stimuli that were, in fact, wrong. The simplest conclusion seems to be the right one: Looking for rules will not work if you cannot find them.

In a number of other studies, instructions of various kinds have been shown to have any of a number of effects. Brooks (1978) used finite-state grammars similar to the one used by Reber (1976) and a paired-associates learning procedure in which strings of letters from grammars were paired with responses of particular kinds (e.g., animal names, cities). He found that informing subjects about the existence of regularities in the letter strings lowered overall performance. Reber, Kassin, Lewis, and Cantor (1980) found poorer performance with explicit instructions when the stimuli were presented in a large, simultaneous array in which letter strings were posted on a board in haphazard fashion. Howard and Ballas (1980) reported detrimental effects of explicit instructions with structured sequences of auditory stimuli when there was no system- atically interpretable pattern expressed by the stimulus sequences. In all these cases, the original finding was basically replicated.

However, Millward (1981) found no difference between explicitly and implicitly instructed subjects in an experiment that, in principle, looked like a replication of Reber's (1976) with the seemingly modest variation that the strings used during learning were up to 11 letters long, whereas in Reber's study stimuli were no longer than 8 letters. Abrams (1987), using a strict replication of Reber's procedure with the exception that the study was run on a computer, failed to find the instructional effect. Dulany et al. (1984) reported no significant differences between the two instructional sets, although in this case a rather different testing procedure may have masked differences. Mathews et al. (in press) found a complex pattern of differences between instructional groups, depending on whether the letter set used to instantiate the grammar was modified over the several days of the experiment. They also used a grammar somewhat more complex than is typical. Danks and Gans (1975) reported no differences when they used a synthetic system that was considerably simpler than the Markovian systems used by others.

Last, several studies showed an advantage for the explicitly instructed subjects. Howard and Ballas (1980) reported that the explicit instructions that debilitated performance when introduced under conditions of semantic uninterpretability could also function to facilitate performance when the stimuli expressed semantically interpretable patterns. Reber et al. (1980) showed that it was possible to shift performance about rather dramatically by intermixing instructional set with the manner of presentation of the stimulus materials and with the time during learning when the explicit instructions were introduced.

There are two factors here that help make these data somewhat less haphazard than they appear to be. The first is psychological salience; the second is the circumstances under which the instructions are given to the subjects. The first of these is the more interesting and the one from which insight into process can be gained.

In two of the instances in which explicit instructions facilitated performance, the manner of presentation of the stimuli was such that the underlying factors that represent the grammar were rendered salient. In Howard and Ballas's (1980) study, the semantic component focused the subjects on the relevant aspects of the patterned stimuli. The effectiveness of such a semantic component has often been noted in artificial grammar learning studies (Moeser & Bregman, 1972; Morgan & Newport, 1981). In Reber et al.'s (1980) study, the simple expedient of arranging the exemplars of the grammar according to their underlying form produced the instructional facilitation. Moreover, several other researchers apparently arranged matters, inadvertently, so that structural properties became more salient. In Millward's (1981) study, for example, the use of longer strings provided many opportunities for subjects to be exposed to the loops or recursions in the grammar (see Figure 1) and thereby increased the psychological salience of the underlying structure. In Danks and Gans's (1975) study, the relatively simple nature of the stimuli likely acted to equate the mode of processing of the stimuli in both groups; that is, both groups were likely using a reasonably explicit mode independently of the instructions. Hence the converse of the earlier conclusion: Looking for rules will work if you can find them.

Some cases appear to be genuine failures to replicate the original finding: specifically, those of Dulany et al. (1984) and

Abrams (1987). In Dulany et al.'s study, the procedure used during learning should, in principle, have yielded a difference during testing. These are some interesting aspects of this experiment, but there are no obvious reasons why the effect failed to emerge. On the surface, Abrams's study, which was carried out in our laboratory, should also have produced an instructional effect. In that study, however, there were reasons for suspecting the computer used to run the study as the culprit. This remains to be explored, but simply presenting the explicit instructions on the computer screen may not have the compelling quality that a "real" experimenter reading them has. As other work suggests (Reber & Allen, 1978), implicit learning is the default mode; therefore, if the subjects do not understand or do not believe the instructions, then no differences would be expected. This suggests a possible important methodological factor that has gone largely unnoticed: the sophisticated equipment that we use commonly in our laboratories may be having untoward effects on our studies.

In any event, the literature on the implicit/explicit problem is clearly complex, and it takes but a moment's reflection to appreciate the fact that there are still other important issues lurking behind these findings. First, it seems clear that any number of confounding factors may influence, either positively or negatively, the impact of explicit instructions (cf. Lewicki, 1986a). Such instructions may introduce an element of stress or anxiety, may evoke a sense of motivation, may encourage one or another conscious strategy, and the like. To date, few of the researchers mentioned have taken such factors into account, and not much is known about how conscious, explicit processing systems interact with the implicit and unconscious. For example, a recent study suggests that using instructions that engage the explicit system may also elicit anxiety and that anxiety may be related to poor performance on a standard grammar learning task (Rathus, Reber, & Kushner, 1988).

Second, it seems clear that we are still dealing with a rather limited kind of analysis of complex learning, particularly if one wishes to view this research in its constrained laboratory setting as representing a general metaphor for real-world acquisition processes. In the real world nearly all complex skills are acquired with a blend of the explicit and the implicit, a balance between the conscious/overt and the unconscious/covert. There is surely a difference between simply informing a learner that the stimulus materials have structure, as researchers in the aforementioned experiments did, and telling the learner something definitive about that structure. The next section deals with this issue.

Effects of Providing Specific Information

This issue concerns the impact of giving subjects precise information about the nature of the stimulus display that they will be exposed to. This is an important question for a number of reasons. For one, this issue broaches on some of the classic questions in pedagogic theory about how best to convey highly complex and richly structured information to students. It also emerges in various studies on the acquisition of expertise in

such areas as medical diagnosis, in which the relationship between specific knowledge presented to medical students and their emergent tacit knowledge base is turning out to be most complex (see, e.g., Carmody, Kundel, & Toto, 1984).

Reber et al. (1980) attempted to address this issue by using the standard grammar learning procedure. In that study, subjects were presented with the actual schematic structure of the grammar; that is, they were presented with Figure 1. Each subject was handed the diagram and given a 7-min "course" in how such a structure can be used to generate strings of symbols. This procedure was supported by an observation period during which a set of exemplars was shown to the subjects. In this training format a maximally explicit learning procedure was thus mixed with a maximally implicit one.

Reber et al. (1980) explored the manner of interaction between these two modes of apprehension by introducing the explicit training at different points in the observation period. One group of subjects received the explicit instruction at the outset before any exemplars were seen; one group received it part way through the observation period; and for a third group, the explication of structure was delayed until after they observed the full set of exemplars. As in the typical grammar learning study, knowledge acquired during learning was assessed by means of a well-formedness task.

The key finding was that the earlier during the observation training the explicit instructions were given, the more effective they were. From the previous discussions it is clear that increasing the salience of the relationships between symbols increases the effectiveness of subjects' attentional focus. It is also clear that instructions that encourage the subject to deal with the stimuli in ways that are discoordinate with underlying structure have detrimental effects on acquisition. Thus the explication of the precise nature of the structure underlying the stimuli must have differential impact on learning, depending on which of these two processes is encouraged.

The most plausible interpretation, and the one that has interesting implications for theories of instruction, is that the function of providing explicit instructions at the outset is to direct and focus the subjects' attention. It alerts them to the kinds of structural relations that characterize the stimuli that follow and permits appropriate coding schemes to be implemented. These instructions did not teach the grammar in any full or explicit fashion; rather, they oriented the subjects toward the relevant invariances in the display that followed so that the subjects, in effect, taught themselves.

Accordingly, when such explicit instruction is introduced later in the observation period, its effects are different because two sources of difficulty are introduced. First, it imposes a formalization of structure that is, in all likelihood, discoordinate with the tacit system that was in the process of being induced. Second, it reduces the number of exemplars that can be used as a base for extracting invariance patterns. In the extreme case in which the instructions were delayed until the completion of the observation period, this informational base is virtually eliminated.

These points can use some further exploration. There is every reason to suspect that subjects' tacit representation of rules is idiosyncratic in various characteristics. The induction

routines hypothesized by Holland et al. (1986) predict this personalized aspect at the outset. From earlier work, it is clear that subjects are known to use a wide variety of coding schemes in focusing their attention on the stimuli (see Allen & Reber, 1980; Reber & Allen, 1978; and Reber & Lewis, 1977, for details). So long as these schemes do not entail inappropriate rule formation, their impact is superficial. Independently of individualistic mnemonics, attentional focusing priorities, or preferred rehearsal strategies, the implicit learner will emerge from the training session with a tacit, valid knowledge base coordinate with the structure of the stimulus environment.

The degree to which the explicit instructions introduce difficulties will thus be dependent on the extent to which the subject's tacit representation of structure matches the formalization provided by the schematic of the grammar and the accompanying characterization of its generational properties. In Dulany et al.'s (1984) terms, subjects are learning "correlated grammars" whose properties are, in all likelihood, not commensurate in any simple way with the Markovian system in use. Recent results from Mathews et al. (in press) strongly corroborate this interpretation.

The deep difficulty here is that there is a potentially infinite number of formalizations that could account for the structure displayed in any given subset of strings from one of these grammars; present the "wrong" one to a subject, and the instructions will not have a salutary effect. The problem is apparent: How much can we expect a subject to benefit from the specific information that the set of exemplars just observed and tacitly coded as, say, bigram covariation patterns is "in reality" to be formalized as a Markovian process? To take an obvious analogy, most of us with our extensive observation and generation of utterances in English have failed to derive any facilitative effect of explicit instruction with transformational grammar that, at least in principle, can be posed as a legitimate formalization of our tacit knowledge. Moreover, such explicit awareness of structure can actually be a nuisance when one tries to fulfill the kinds of demands placed on subjects in these experiments, as in the discrimination of well-formed, novel instances from instances that contain some violation of the formal system.

In a study that addressed this point directly, Bialystok (1981) found that subjects learning French as a second language could rapidly and accurately detect ungrammatical sentences and could do so largely independently of the complexity of the grammatical rules violated. However, when asked to characterize the nature of the violated rules, the complexity factor played a significant role. Complexity, of course, is defined in such cases by the kinds of grammars that are taught in "French as a second language" courses.

In summary, although there are not a lot of hard empirical data here, those that are available point toward an interesting conclusion. Specific instruction concerning the materials to be learned in complex situations will be maximally beneficial when it is representationally coordinate with the tacit knowledge derived from experience. Because this issue is ultimately critical for theories of instruction, it is one much in need of close examination.

On Deep and Surface Structure

The issue here is the degree to which implicit learning can be seen as acquisition of knowledge that is based on the superficial physical form of the stimuli or as knowledge of the deeper, more abstract relations that can, in principle, be said to underlie them.

In an early article, Reber (1969), reported evidence for the proposition that implicit knowledge is abstract and not dependent in any important way on the particular physical manifestations of the stimuli. This study consisted of two sessions during which subjects memorized letter strings from a grammar. When the second session began, the stimulus materials were, without warning, modified. For some subjects the same letters continued to be used, but the rules for letter order were now those of a different grammar (the "syntax" was altered). For other subjects the underlying structure was not tampered with, but the letters used to represent the grammar were replaced with a new set (the "vocabulary" was changed). The two obvious control groups, one for which both aspects were altered and one for which neither was changed, were also run. The various manipulations had systematic effects on subjects' ability to memorize stimuli in the second session. Modification of the rules for letter order produced decrements in performance; modifications of the physical form had virtually no adverse effects. So long as the deep rules that characterized the stimuli were left intact, their instantiations in the form of one or another set of letters was a factor of relatively little importance.

The recent study by Mathews et al. (in press) supported this general finding. Their experiment was run over a 4-week period. Subjects who received a new letter set each week (which was based on the same underlying syntax) performed as well as subjects who worked with the same letter set throughout the course of the experiment. The effect was quite striking; the transfer from letter set to letter set occurred smoothly and apparently without the need for any conscious translation.

Reber and Lewis (1977) reported an equally striking example of the abstract nature of tacit knowledge. They assessed knowledge of the grammar by having subjects solve anagram problems. After a standard training session during which subjects memorized exemplars from the language, they solved anagrams from the synthetic language over a 4-day period. For reasons to be discussed, it is convenient to code letter strings in the form of bigrams and to note the rank order of frequency of occurrence of each possible bigram. For example, the string PTTTVV contains bigrams PT, TV, and VV once each and TT twice. Given how this experiment was run, three rank orders of frequency of occurrence of bigrams exist: (a) one based on the actual solutions offered by the subjects (corrected, of course, for guessing), (b) one based on the frequency of occurrence of each acceptable bigram within the artificial language itself (within the string lengths used), and (c) one based on the actual bigrams that appeared in the learning stimuli.

Rank-order correlations among these three were revealing. The correlation between (a) and (b) was .72, whereas that

between (a) and (c) was only .04. The interesting point about these results is that the comparison between (a) and (b) is a comparison between subjects' usable knowledge and a deep representation of the frequency patterns of the grammar. Rank-order (b) was formed on the basis of the full set of acceptable strings that the grammar could, in principle, generate within the specific string lengths. Subjects, however, never saw this full set of strings; they were exposed only to the exemplars chosen for the training sessions. These particular strings were selected to ensure that each subject saw at least one string of each possible length and, for each length when it was possible, at least one example of each of the grammar's three loops. This procedure yielded a set of strings that displayed all of the deep structural characteristics of the finite-state language but, in terms of specific frequency of bigrams, was distinctly idiosyncratic.

The comparison between (a) and (c) is a reflection of the degree to which subjects are simply keying on the raw frequency data as displayed in the exemplars. The failure for this correlation to be different from zero suggests that subjects were not solving the anagrams on the basis of superficial knowledge of frequency of bigrams or on the basis of a fixed set of memorized instances. They clearly acquired knowledge that can be characterized as deep, abstract, and representative of the structure inherent in the underlying invariance patterns of the stimulus environment.

This finding is analogous to Posner and Keele's (1968, 1970) oft-cited abstraction of prototype effect. The underlying prototypes that their subjects extracted from the exemplary dot patterns are specifiable only in terms of an averaging of the spatial relations among the various components of the patterns. But, psychologically, such an averaging is not just a simple piling up of the features of the exemplars. If memory behaved like that, the resulting representation would be, not distinct prototypes that Posner and Keele found, but a blob.

The induction routine that appears to be operating in situations such as these is necessarily one that results in an abstract representation. Moreover, it is one that is applicable to the classification of novel instances and not specifically characterizable by a raw compilation of experienced instances.

This issue is one of considerable complexity. The point of the preceding argument is not that all memorial systems must be viewed as founded on induced abstractions. The evidence of Brooks (1978) and others (cf. Smith & Medin, 1981) shows that memories are frequently based on instantiations, fairly uninterpreted representations of the stimulus inputs. The point is that when implicit acquisition processes are operating, the resulting memorial system is abstract. As was shown elsewhere (Allen & Reber, 1980; Reber & Allen, 1978), the same subjects working with the same grammars can emerge from a learning session with either an instantiated memory system or an abstract one, depending on the learning procedures used. In those articles, that old war horse *functionalism* was shown to provide the best characterization of this issue; that is, the specific functions that need to be carried out invite the learner to assume a cognitive stance that is functional, that will accomplish the task at hand. Under some circumstances (such as the paired-associates learning procedure used by Brooks, 1978, and by Reber & Allen, 1978), a rather concrete, instantiated memorial system will be established;

under others (the instant asymptote technique in the PL studies by Reber & Millward, 1968, and the observation procedure in Reber & Allen's 1978 study), a distinctly abstract representation will emerge. In the pure implicit, unconscious acquisition mode, the default position is abstraction.

On Mental Representation

As Rosch and Lloyd (1978) pointed out, sooner or later every discourse on mental process and structure must come to grips with the problem of the form of the representation of knowledge held. Such discussions must begin with some presumptions. The ones introduced here are, of course, open to emendation as understanding processes. They are taken as the starting point simply because they have considerable explanatory power, more than most contemporary cognitivists have granted.

First, the general argument put forward by such diverse theorists as Gibson (1966, 1979), Garner (1974, 1978), and Neisser (1976), that the stimulus is more than the physical setting for the occurrence of a response, is taken as a given. This point is more than a simplistic swipe at behaviorism; it is an argument that stresses that the stimulus domain within which we function is extraordinarily rich and complex and, in all likelihood, much more so than most cognitivists have been willing to recognize. The underlying causative nature of the stimulus environment is rarely explored; most theorists are satisfied with characterizations that are theory driven.

Second, there is general agreement with the arguments put forward (in rather different forms, to be sure) by Palmer (1978) and Anderson (1978, 1979, 1983) to the effect that most theoretical attempts to deal with the representation issue are misguided. Palmer maintained that the confusion derives from a failure to deal directly with metatheoretical factors concerning existing models. Anderson argued that in principle, there are no ways in which behavioral data can be used to identify uniquely any one particular mental representation. There are some reasons for perhaps disputing these claims (see, e.g., Hayes-Roth, 1979; Pylyshyn, 1979, 1980), but they are not a concern here.

From the point of view that I take as presumptive here, it matters not at all whether the following interpretations of mental representation are supported by a well-structured consideration of metarepresentational factors or whether they can be shown to be uniquely specifiable. The point of view that I take reflects that of classical functionalism as introduced in the preceding section. Functional theories are typically regarded these days as formulations (abstract, to be sure) of what is possible for a person to process and why. This seems right, and as has been argued elsewhere (Allen & Reber, 1980; Reber & Allen, 1978), the main consideration should be with characterizing representations, in terms of how the individual can be seen as behaving in an adaptive fashion, rather than in terms of pure representational theory. For example, as discussed earlier, there are good empirical reasons for regarding the functional representation of the mental content of a finite-state grammar as an ordered set of bigrams (and trigrams; see Mathews et al., in press) and not as a formal Markovian system.

Third, the oft-dismissed position of representational realism is accepted as a first approximation. What is in the stimulus world is what ends up in the mind of the perceiver/cognizer. The point is that a good way to start dealing with the representation problem is with the physical stimulus itself. Under various constraints of processing and various task demands, enrichment and/or elaborative operations are certainly used, and the resulting coded representation may very well not be isomorphic with the stimulus field. Nevertheless, as Mace (1974) put it, a good initial strategy is to "ask not what's inside your head, ask what your head's inside of."

Several findings from studies with artificial grammars are relevant to the issue of representation. Table 1 gives the summary data from 14 separate experiments that reveal some interesting patterns. Some details on procedure are needed: In all of these studies, knowledge acquired during learning was assessed through the well-formedness task in which subjects are presented with a number of test strings (typically 100) that must be classified as either grammatical or nongrammatical. In the typical experiment, the 100 trials consist of 50 unique items, each of which is presented twice without feedback about the correctness of the response.

This procedure yields data that directly address the representational issue. The logic of the analysis is simple. There are four possible outcomes for each individual item for each subject: The subject may classify it correctly on both presentations (CC), classify it correctly on only one of the two (CE or EC), or misclassify it on both presentations (EE). Assume that the subject operates by using a simple decision-making strategy: When the status of the item is known, it is always classified correctly; when it is not known, a guess is made. This simple model is quite powerful and allows for a surprisingly deep analysis of the representation problem.

Specifically, under this model, (a) the values of CE, EC, and EE should be statistically indistinguishable from each other, and all should be significantly lower than the value of CC. This pattern is expected on the grounds that the items that contribute to CE, EC, and EE are those about which the subject's knowledge base is not relevant. (b) A value of EE significantly greater than the values of EC and CE is prima facie evidence for the elaboration of nonrepresentative rules on the part of subjects. Thus if subjects emerge from the learning phase with rules (either explicit or implicit) that are not accurate reflections of the grammar, this knowledge base will consistently lead them to misclassify particular items. (c) The robustness of representative knowledge can be assessed from the relationship between the values of EC and CE. If the value of CE is detectably larger than EC, we can reasonably suspect that forgetting was occurring during testing; correspondingly, if EC is larger than CE, we can infer that learning was occurring during testing. (d) One can estimate knowledge of the grammar by looking at the value of CC, which contains only those items whose status was known by the subjects plus those guessed correctly on both presentations. Last, (e) one can derive an overall measure of consistency of responding by taking the sum of CC and EE.

Of the values from 14 experimental conditions (see Table 1), the uninteresting ones can be dispensed with first. There are no cases in which the values of CE and EC are significantly different from each other. Thus there is no evidence of loss

Table 1

Correct (C) and Error (E) Response Patterns to Individual Items During Well-Formedness Tasks

	Pattern				
Condition/training procedure	CC	CE	EC	EE	Consistency
Reber (1967)					
1. Simple memorization	.69	.07	.12	.12	.81
Reber (1976)					
2. Simple memorization	.66	.10	.11	.13	.79
3. Memorization/rule search	.53	.12	.12	.23[a]	.76
Reber and Allen (1978)					
4. Simple observation	.73	.08	.09	.11	.84
5. Paired associates	.65	.12	.07	.16[a]	.81
Reber, Kassin, Lewis, and Cantor (1980, Experiment 1)					
6. Random display/implicit instructions	.51	.16	.14	.19	.70
7. Random display/explicit instructions	.48	.12	.14	.25[a]	.73
8. Structured display/implicit instructions	.52	.16	.16	.16	.68
9. Structured display/explicit instructions	.68	.10	.10	.11	.79
Reber, Kassin, Lewis, and Cantor (1980, Experiment 2)					
10. Rules at beginning of observation	.67	.11	.12	.11	.78
11. Rules in middle of observation	.58	.12	.14	.16	.74
12. Rules at end of observation	.57	.13	.15	.16	.73
13. Rules only	.54	.11	.16	.18[a]	.72
14. Observation only	.48	.15	.13	.24[a]	.72

[a] EE value significantly higher than the mean of the CE and EC values.

of knowledge during the well-formedness task and no evidence that any additional learning was taking place.

The interesting results are those concerning comparisons between the values of EE and those of EC and CE. When no difference is found between EE and the mean of CE and EC, it is reasonable to conclude that there was no evidence of nonrepresentative rules in use. Values of EE that are large in relation to those of EC and CE indicate that subjects were using rules that are not representative of the grammar. The footnoted values in Table 1 are the five conditions that yielded evidence that subjects emerged from the learning phase with notions about structure that were not commensurate with the stimulus display.

It is instructive to look closely at these five cases. In Condition 13, the subjects were given only the schematic diagram of the grammar but no opportunity to observe exemplars. It appears that, not surprisingly, such a procedure encourages subjects to invent specific rules for letter order and, in the absence of complete learning, to elaborate rules about permissible letter sequences that are not reflective of the grammar.

Conditions 3 and 7 illustrate what happens when subjects are under an instructional set that encourages the use of rule search strategies but in which the letter strings are given to them in a haphazard order. Such a set of demand characteristics encourages subjects to invent a sufficient number of inappropriate rules to inflate the EE values.

In Condition 5 a paired-associate task was used to impart knowledge. As is discussed elsewhere (Allen & Reber, 1980; Reber & Allen, 1978), the very nature of such a task leads subjects to set up an instantiated memorial system composed of parts of items and some whole items along with their associated responses. Hence the inflated EE value is not due to the application of inappropriate rules; rather, it is due to subjects' tendency to misclassify test strings because inappropriate analogies exist in instantiated memory. The fifth condition with an inordinately high EE value was 14. There is no obvious explanation for this outcome. This datum is an anomaly; 1 such outcome out of 14, however, is not bad at all.

The remaining 9 conditions all yielded response patterns that fit with the proposition that whatever subjects are acquiring from the training sessions can be viewed as basically representative of the underlying structure of the stimulus domains. These consist of "neutral set" conditions, in which the subjects are led to approach the learning task as an experiment in memory or perception and no mention is made of the rule-governed nature of the stimuli (Conditions 1, 2, 4, 6, and 8), and "structured set" conditions, in which subjects are provided with information concerning rules for letter order but in a manner than ensures that conscious rule searching will be coordinate with the kinds of rules in use (Conditions 9, 10, 11, and 12).

Taken together, these experiments lend general support to the proposition that implicit learning functions by the induction of an underlying representation that mirrors the structure intrinsic to the environment. Such an induction process takes place naturally when one is simply attending in an unbiased manner to the patterns of variation in the environment or when one is provided with an orientation that is coordinate with these variations.

This characterization of the appropriateness of mental representation entails nothing about the sheer amount of knowledge that one takes out of a learning session. In fact, it is relatively easy to show that there is little to enable one to distinguish explicit from implicit processes here. The consistency values in Table 1 reveal surprisingly little variation from condition to condition, particularly when compared with the range of CC and EE values. These consistency values can be seen as a raw estimate of the total number of rules that subjects can be said to be using during decision making, for they are simply the sums of the CC and EE values. Taking a simple (and only quasi-legitimate) average across conditions reveals that the overall mean consistency values for the footnoted conditions and the nonfootnoted ones are .75 and .76, respectively. Thus there is no evidence that either set of conditions produces more rule learning; the difference is that explicit learning results in the emergence of a number of inappropriate rules, whereas implicit learning tends to yield representative, veridical rules.

This same model of representation is supported by data from other tasks. In Reber and Lewis's (1977) anagram solution task, subjects worked with the same problem sets over 4 days. In that study there was improvement over time, so a stochastic model was fit to the data and used to predict the pattern of error and correct responses to individual test items that would be expected under the assumption that subjects were not using inappropriate rules. (For details on the model, see the original article.) The results were in keeping with the general theme here. The EEEE value (the proportion of items solved incorrectly on all 4 days of the study) was no higher than would be expected under the assumption that subjects either knew the solution to a particular anagram or made nonsystematic guesses for problems not within the domain of their knowledge base.

Several experiments in which the PL procedure was used are also of interest. The relevant data are the recency curves. In the standard analysis of a PL experiment, a recency curve represents the probability of a given response plotted against the length of the immediately preceding run of that event. Recency curves may take on any of a number of shapes, depending on the conditions of the experiment. Negative recency is most common, particularly early in an experiment. Under various circumstances, however, even positive recency may be observed (see Friedman et al., 1964, for details). The concern here is with the recency curves from experiments with 500 or more trials with a Bernoulli event sequence with probability of the more likely event set at .80.

Figure 2 presents the pooled recency data from five such experiments (see Reber, 1967; Reber & Millward, 1968). All subjects were run though a learning period with either the traditional PL or the instant asymptote technique. The subject-generated curve has been adjusted downward by exactly .04 at all points to correct for a ubiquitous overshooting effect that is observed in all of these many-trial experiments (see Reber & Millward, 1968, for a discussion of this issue). This adjustment in no way modifies the startling aspect of these two curves.

With few exceptions, the curves sit on top of each other. There is no evidence whatsoever for either the positive recency predicted by the early conditioning models or the negative recency reported by many. There is, however, overwhelming evidence for a mental representation that reflects the structure of the stimulus environment. The simplest characterization of this curve, which is based on a total of 44,000 responses, is that it reveals that subjects mimic the structure of the event sequence. Subjects' prediction responses show flat recency curves because the event sequences themselves display flat recency curves—as they must, being Bernoulli in nature.

This is not a new point; it was made earlier by Derks (1963) and by Jones and Meyers (1966), who showed that experiments can encourage either positive or negative recency by presenting event sequences with either many long or many short runs of events. But the precision with which subjects' response patterns can reflect the event patterns has never really been appreciated. To take this point to a further extreme, data like those in Figure 2 are so robust that they can actually be used as a check on one's experimental procedure. In one PL study (Millward & Reber, 1972), the subjects' overall response proportions were .523 and .476 for the two events, a result that was perplexing because each event had been programmed to occur in exactly half of the trials. The anomaly turned out to be in the computer program used to generate the sequences. A check revealed that the two events had actually been presented to subjects with proportions of .520 and .480!

Although the preceding analyses seem to provide support for the representational realist position, it is still unclear just how far one can legitimately push such a proposition. In many of the experiments reported here and in other related areas of study (see Schacter, 1987), subjects respond in ways that indicate that their mental content may not be quite so neatly isomorphic with that of the stimulus field. However, it also seems reasonably clear that when such transforms or constructions of representations are observed, "secondary" processes are responsible; that is, the "primary" process of veridical representation of environmental structure becomes colored either by elaborative operations, as in experiments in which instructional sets encouraged invention of inappropriate rules (Howard & Ballas, 1980; Reber, 1976), or by restrictive operations, as in studies in which task demands led to the narrowing of attentional focus (Brooks, 1978; Cantor, 1980; Reber et al., 1980). Also, careful scrutiny of the EE values in Table 1 reveals that even in the nonfootnoted conditions there was a tendency for some nonrepresentational elaboration to take place. In all nine of these cases, the EE value is equal to or higher than the EC or CE values ($p < .05$ by a sign test).

The problem of mental representation is clearly no easy nut to crack. The position taken here seems to be a reasonable one, although it will probably be shown to be wrong in the final analysis. Tacit knowledge is a reasonably veridical, partial isomorphism of the structural patterns of relational invariances that the environment displays. It is reasonably veridical in that it reflects, with considerable accuracy, the stimulus invariances displayed in the environment. It is partial in that not all patterns become part of tacit knowledge. It is structural
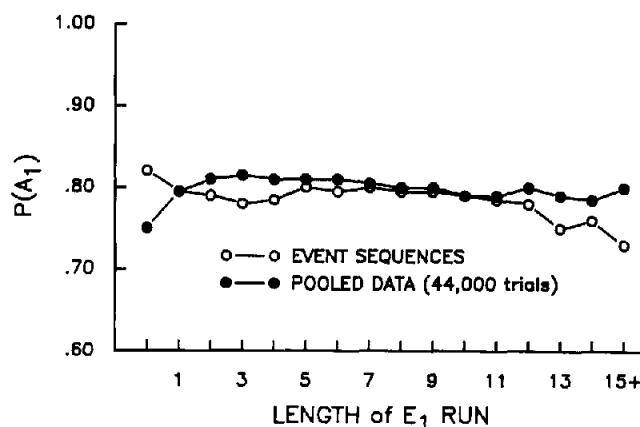

Figure 2. Pooled recency data from five separate probability learning experiments. (The open circles give the average probability that the more frequent event will occur over all event sequences used; filled circles give the average probability that the subjects will make the more frequent response, adjusted for overshooting. The data are based on 44,000 trials from 88 asymptotic subjects tested with the probability of the more frequent event set at .80.)

in that the patterns are manifestations of abstract generative rules for symbol ordering.

On the Availability of Tacit Knowledge

The conclusion reached in the first studies on implicit learning (Reber, 1965) was that the knowledge acquired was completely unavailable to consciousness. The many experiments carried out since have shown that position to have been an oversimplification. The picture that is emerging, though perhaps somewhat less striking, is certainly more interesting. Specifically, knowledge acquired from implicit learning procedures is knowledge that, in some raw fashion, is always ahead of the capability of its possessor to explicate it. Hence although it is misleading to argue that implicitly acquired knowledge is completely unconscious, it is not misleading to argue that the implicitly acquired epistemic contents of mind are always richer and more sophisticated than what can be explicated.

In Reber and Lewis's (1977) study, data were first presented to support this position. Over the 4 days of that study, during which subjects solved anagram puzzles on the basis of the syntax of an artificial grammar, there was a general increase in the ability of subjects to communicate their knowledge of the rule system in use. There was also an increase in the ability to solve the anagrams, but the former never caught up with the latter; that is, as subjects improved in their ability to verbalize the rules that they were using, they also developed richer and more complex rules. Implicit knowledge remained ahead of explicit knowledge.

In a recent study, Mathews et al. (in press) used a novel yoked-control technique to explore this issue. Subjects were interrupted at intervals during a well-formedness judgment task and asked to explicate the rules that they were using. The information was then given to yoked-control subjects, who

were then tested in the same well-formedness task. So equipped, these control subjects managed to perform at roughly half the level of accuracy of the experimental subjects. Moreover, as the experiment progressed and each experimental subject improved, so did each yoked control, but the controls never caught up with the experimental subjects.

The most direct attempt to deal with the issue of the degree to which implicitly acquired knowledge is available to consciousness was carried out by Dulany et al. (1984). After a standard learning procedure, subjects were asked to mark each well-formedness test item as acceptable or not and to specify what features of that item led them to classify it as they did. Dulany et al. argued that the features so marked accounted for the full set of decisions that each subject made, a result that, if correct, supports the notion that used knowledge of the grammar was held consciously. Reber, Allen, and Regan (1985), however, argued that the nature of the task that Dulany et al. used carried its own guarantee of success; that is, the task forced the data to appear as though they carried the implication of consciousness, whereas actually the subjects were reporting only vague guesses about the appropriateness or inappropriateness of letter groups. The issue continues to be disputed. Dulany, Carlson, and Dewey (1985) presented reasons for doubting Reber et al.'s analysis, whereas Hayes (in press) recently produced evidence in support of the interpretation of Reber and his co-workers.

One of the problems with this line of research is that it fails to distinguish between knowledge that is available to consciousness after attempts at retrieval and knowledge that is present in consciousness at the time that the decisions themselves are being made. Carmody et al. (1984) noted this problem in assessing the knowledge base that physicians are taught to use versus what they actually use in diagnosis, and Schacter (1987) argued that conclusions reached about the availability of implicit information must take account of a variety of task constraints that have their own impacts. Nevertheless, if it is not yet clear, the discussions that follow will emphasize even further the central thesis of this line of research. To wit: A considerable portion of memorial content is unconscious, and, even more important, a goodly amount of knowledge acquisition takes place in the absence of intent to learn.

Entailments and Implications

The preceding discussion is a reasonably thorough review of the current state of affairs as regards the general issues of the acquisition, usage, representation, and availability of tacit knowledge. As Schacter (1987) pointed out recently, one of the intriguing aspects of the history of work on this issue is that there is such an amazing variety of implicit processes that have been observed and yet there is nothing approaching a satisfactory theoretical account of them. What follows may or may not improve on this state of affairs. The following is a small flurry of speculation concerning the possible entailments and implications of the research. Each of the topics is touched upon only briefly; the point here is to provoke new avenues of study, not to draw any hard conclusions.

On the Origins of Unconscious Cognition

Usually the header here is the "Origins of *Conscious* Cognition," not "*Unconscious.*" Traditionally, the focus has been on consciousness with the implication that defining and characterizing consciousness will solve the problem; unconscious processes will be handled by the invoking of exclusionary clauses. The history of the variety of ways in which the unconscious has been represented (Ellenberger, 1970) shows this clearly. Consciousness assumes epistemic priority because it is so introspectively obvious, whereas the unconscious must be struggled with in derivative fashion.

The point to be defended here is that this ordering of priorities has been an error. The theoretically important exercise should be on the origins of unconscious cognitive processes. Consciousness, evolutionarily speaking, is a late arrival on the mental scene. Perhaps it is not of such recent origin as some have argued (Jaynes, 1976), but surely it postdates a number of fairly rich and elaborative cognitive processes that functioned and still function in our phylogenetic predecessors (Griffin, 1981, 1984). There is, moreover, absolutely no reason to suppose that these presumably adaptive mental capacities ought to have been lost. In fact, there are a number of reasons for supposing that they continue to flourish interpenetrated by an emerging executive system, conscious mentation.

Taking such a perspective gives unconscious cognition the empirical and theoretical priority that it deserves but has not enjoyed since the era of the philosophical emergentists. Unconscious cognitive functions should not have to be defended against arguments that deny their role in action (see the debate between Dulany et al., 1984, 1985, and Reber et al., 1985). The proper stance is to assume that unconscious mental processes are the foundations upon which emerging conscious operations are laid. The really difficult problem, then, is to discern how these components of mind interact.

This perspective has some interesting entailments. One is that it suggests a novel way to see how the work on implicit learning fits in with a good deal of other research on the cognitive unconscious. Another is that it allows for a new framework for:

Parsing the Cognitive Unconscious

A conspicuously large number of processes and functions have been assigned to the unconscious over the past century or so. They have come in a variety of forms, some concerned with perceptual processes, some with dynamic, some with motivational and emotional, and some with cognitive. A number of schemata have also been proposed for defining and classifying the subcategories of unconscious functions and operations (see Ellenberger, 1970; Erdelyi, 1985). Herein is one more.

As a first approximation, assume a relatively high-level parse that separates unconscious mentation into two classes, one that most aptly can be called the *primitive* and one that, for reasons to be spelled out, can be thought of as the *sophisticated*. The primitive unconscious encompasses a va-

riety of basic functions, all of which are carried out more or less automatically and are more or less devoid of meaning, *affect, or interpretation.* Included here is a range of processes such as Gibson's (1979) direct pickup of information in perception, Hasher and Zacks's (1984) automatic encoding of information about the frequency of events as they occur, Lewicki's (1985, 1986a, 1986b) studies showing unconscious apprehension of feature covariation, Broadbent and his co-workers' (Berry & Broadbent, 1984; Broadbent & Aston, 1978; Broadbent et al., 1986) studies on simulated economic and manufacturing systems, and, of course, the studies reviewed on implicit learning of complex covariations displayed by synthetic grammars and structured event sequences.

The operations of this primitive unconscious seem to be about as fundamental for a species' survival as any nonvegetative function could be. Virtually every organism must be able to perform a basic Hasher and Zacks's (1984) type of counting of the occurrences of ecologically important events. Ground squirrels presumably count small holes and keep a kind of log of their locations, and lions count gnus and their various properties. Rats, pigeons, dogs, and other laboratory subjects count covariations between events even in the most basic circumstances. The essence of Pavlovian conditioning is the apprehension of a genuine covariation between the conditioned stimulus and the unconditioned stimulus (Rescorla, 1967, 1988). The reason that the work of such researchers as Hasher and Zacks, Lewicki, and others is typically regarded as cognitive in nature and somehow different from Rescorla's is that it is typically carried out with mature adult subjects and with materials that have a rich cognitive underpinning, such as words, sentences, and pictures. Yet, there is no a priori reason to regard these high-level cognitive "counts" as different in any fundamental way from the very simple countings of Rescorla's subjects. What is different is the process by which each organism comes to categorize the items whose frequency and covariation patterns are being logged, not the mechanism for representing the raw data.

From this perspective, the grammar learning experiments, the PL studies, and the rest of the literature on implicit learning can be viewed as epistemic kin of the most basic of the primitive unconscious functions. For example, Lewicki (1985) showed that in the limiting case, only one exposure to a target person with a salient personality characteristic (e.g., kindness, capability) and particular physical characteristic (e.g., long or short hair) is sufficient to set up a tacit knowledge base that reflects these covariations and affects decisions made about novel people. In cases with a richer data base, such as the structured event sequences of the synthetic languages and the probability learning experiments, the kinds of structural covariations that are apprehended are deeper and more abstract. Yet, they can be viewed as categorical extensions in that the basic process is, in principle, still one of counting, only what is being counted are complex interdependent covariations among events—or, as they are commonly known in the literature, *rules.*

These various manifestations of the functions of the primitive unconscious have a number of additional factors in common. First, and most basic, the pickup of information takes place independently of consciousness or awareness of what is picked up. Put another way, adding the factor of consciousness changes the very nature of the process (Reber, 1976; Reber & Allen, 1978; Reber et al., 1980). As stated at the outset, this may be taken as the defining feature of implicit learning.

Second, although much of what is acquired may eventually be made available to conscious expression, what is held or stored exceeds what can be expressed. This is displayed in one of two fashions: Either predictions of performance made on the basis of available knowledge fall short of actual performance (Mathews et al., in press; Nisbett & Wilson, 1977; Reber & Lewis, 1977) or differential effects of variables are seen when the implicit–explicit instructional set is varied (Graf & Mandler, 1984; Reber, 1976; Reber et al., 1980; Schacter & Graf, 1986). Schacter (1987) noted that this inequality between implicit knowledge that is inaccessible and implicit knowledge that can be articulated explicitly may have to do with the degree of elaborative encoding that is allowed. Indeed, most of the studies in which a substantial proportion of once-tacit knowledge is made available to consciousness are those in which considerable overt encoding is carried out (e.g., Dulany et al., 1984).

Third, the memorial content of the primitive unconscious has a causal role to play in behavior. This proposition goes almost without saying, given the preceding discussion, but it needs to be specified; if there were no causal component to unconscious cognition, we might as well simply return to a radical behaviorism. Put simply, the primitive unconscious processes are for learning about the world in very basic ways. They are automatic and ineluctable; they function to pick up critical knowledge about categories and about covariations of aspects of categories. They do not, however, have any functions that involve meaning or affect; these are the province of the *sophisticated* unconscious.

In this latter class are included such phenomena as unconscious perception of graphic and semantic information (Marcel, 1983), perceptual vigilance and perceptual defense (Erdelyi, 1974), the implicit pickup of affective information that is based on phonological factors (Corteen & Wood, 1972) or geometric features (Kunst-Wilson & Zajonc, 1980; Seamon, Brody, & Kauff, 1983; Seamon, Marsh, & Brody, 1984), and repetition priming effects with various linguistic and nonlinguistic materials (Jacoby & Dallas, 1981; Scarborough, Cortese, & Scarborough, 1977; Scarborough, Gerard, & Cortese, 1979, among many others). In some ways the evidence for the unconscious element is stronger here than it is with the primitive unconscious. The use of forced-choice recognition tests as a measure of sensitivity in the work of Marcel (1983), Kunst-Wilson and Zajonc (1980), Seamon et al. (1983), and Seamon et al. (1984) supports the strong claim that these processes are occurring virtually independently of awareness. Although there may be some problems with methodology (see Holender, 1986), this procedure is, in principle, superior to that used by most researchers, who mainly browbeat their subjects into telling what they know (Allen & Reber, 1980; Berry & Broadbent, 1984; Brooks, 1978; Reber & Allen, 1978).

What makes these various processes intriguing and what differentiates these sophisticated processes from the primitive

is that all share a basic operating property: They all depend on a rich, abstract knowledge base that asserts itself in a causal manner to control perception, affective choice, and decision making independently of consciousness. This component of the cognitive unconscious depends on previously acquired knowledge, as opposed to the primitive component, which operated to acquire such knowledge. The very epistemic base that makes these sophisticated processes functional can be seen as that derived from the primitive processes.

These sophisticated systems also differ from the primitive in other ways. First, they are components of mind that are generally available to consciousness. In other words, there is awareness of the knowledge base itself; a subject in one of Marcel's (1983) experiments surely knows the target word and, moreover, surely knows that he or she knows it. What is crucial is that this overt knowledge base has a higher threshold for engagement than the covert one does. Second, they are based on knowledge systems that have become highly automatized. They share this automatic quality with the primitive functions in the limit, but there are good reasons for thinking that much of this interpretive and semantic knowledge derived from explicit processes that became automatic only after pained, conscious action. Interestingly, this line of argument parallels that taken by Dulany et al. (1984, 1985) in their criticism of the synthetic grammar learning studies; however, Dulany et al. targeted the wrong level for invoking it. Last, these systems all function on a symbolic level. All of the critical components of the sophisticated unconscious involve semantic and affective properties of stimuli. This aspect seems to be largely missing in the primitive domain. It seems that these sophisticated processes are more uniquely the stuff of humanity than are the primitive processes, which are operating systems that we share with virtually all corticated species and are found rather far down the phylogenetic scale.

The Robustness of Implicit Processes

There has been a good deal of work to suggest that implicit systems are robust in the face of disorders that are known to produce serious deficits in conscious, overt processes. Support for this functional separation of conscious and unconscious cognitive processing has come from the study of various patient populations. Classic cases are amnesia (see Milner, Corkin, & Teuber, 1968, for the early work and Schacter, 1987, for a recent overview), blindsight (Weiskrantz, 1986), prosopagnosia (Bauer, 1984), and alexia (Shallice & Saffran, 1986). In all these cases there is compelling evidence of effective performance in the absence of awareness.

The model of mental parsing suggested earlier provides a novel interpretation of this work. There is a standard heuristic in evolutionary biology that older primitive systems are more robust and resistant to insult than are newer, more complex systems. The hypothesis that the implicit cognitive processes are the functional components of the evolutionarily older, primitive system predicts that they should show greater resistance than should explicit processes. By extension, all of the various phenomena that have been cited as manifestations of primitive unconscious processes would be expected to display similar robustness under conditions in which parallel explicit processes have been diminished or even lost entirely.

The strongest evidence in support of such an interpretation comes from cases in which direct comparison between implicit and explicit processes has been made in clinical settings. In an extended series of studies, Warrington and Weiskrantz (see 1982 for an overview) found no deficits in amnesics when the task involved memory for words based on word-stem and word-fragment cues, but performance was seriously impaired when overt word recognition and recall procedures were used. A similar pattern emerged with Hasher and Zacks's frequency encoding task, in which performance was found to be robust in the face of clinical depression (Hasher & Zacks, 1979; Roy, 1982) and even Korsakoff's syndrome (Strauss, Weingartner, & Thompson, 1985). Last, a recent study of Abrams and Reber (1989) suggested that even the acquisition of knowledge is undiminished so long as the task is a nonreflective, unconscious one. They used an implicit grammar learning task and an explicit short-term memory task with a mixed population of institutionalized depressives, schizophrenics, and alcoholics with organic brain damage. The patients performed more poorly than a normal control group on the memory task, but the performances of the two groups were statistically indistinguishable on the implicit learning task. This last study is particularly important, for it is one of the few that shows that implicit *learning* is robust in the face of serious psychological and/or neurological disorders (see Graf & Schacter, 1985, for another example involving a word-completion task).

On Intuition

One of the gains of this line of research on implicit processes is that it provides the opportunity to reclaim intuition for cognitive psychology. There is probably no cognitive process that suffers from such a gap between phenomenological reality and scientific understanding. Introspectively, intuition is one of the most compelling and obvious cognitive processes; empirically and theoretically, it is one of the processes least understood by contemporary cognitive scientists.

The basic argument is simple: The kinds of operations identified under the rubric of implicit learning represent the epistemic core of intuition; that is, the introspective qualities that most people—from Bergson (1913) and Croce (1922) to Jung (1926), Polanyi (1958), and Westcott (1968)—identify when discussing intuition are those processes that have emerged in the studies of implicit acquisition of complex knowledge. Perhaps the most compelling aspect of intuition, and the one most often cited in the various definitions that have been given (see Westcott, 1968), is that the individual has a sense of what is right or wrong, a sense of what is the appropriate or inappropriate response to make in a given set of circumstances, but is largely ignorant of the reasons for that mental state. This, of course, is how the typical subject has been characterized after a standard acquisition session in an implicit learning experiment.

The point is that intuition is a perfectly normal and common mental state/process that is the end product of an implicit learning experience. In other words, intuition ought not to be embedded in personality theory as it was with Jung

(1926), and although it is a topic of some philosophical interest, it is probably best not dealt with as an a priori topic as it was by Croce (1922). It is a cognitive state that emerges under specifiable conditions, and it operates to assist an individual to make choices and to engage in particular classes of action. To have an intuitive sense of what is right and proper, to have a vague feeling of the goal of an extended process of thought, to "get the point" without really being able to verbalize what it is that one has gotten, is to have gone through an implicit learning experience and have built up the requisite representative knowledge base to allow for such judgment.

Summary

This article is an attempt to come to grips with an essential, although oft-ignored, problem in contemporary cognitive psychology: the acquisition of complex knowledge. At the heart of the presented thesis is the concept of implicit learning wherein abstract, representative knowledge of the stimulus environment is acquired, held, and used to control behavior. The operations of implicit learning are shown to take place independently of consciousness; their mental products have been demonstrated to be held tacitly; their functional controlling properties have been shown to operate largely outside of awareness. The strong argument is that implicit learning represents a general, modality-free Ur-process, a fundamental operation whereby critical covariations in the stimulus environment are picked up.

The key problem in all of this is to specify, as clearly as possible, the boundary conditions on the process of implicit learning—that is, to outline the circumstances under which it emerges and those under which it is suppressed or overwhelmed. A substantial part of the empirical work reviewed here should be seen in that light. Last, there has been an attempt to show how such a process can be seen as functioning in the context of other, complex cognitive operations and to speculate on how it might be viewed in a variety of other frameworks, from that of evolutionary theory to those of various clinical syndromes affecting cognitive function to those of some novel considerations of intuition.

References

Abrams, M. (1987). Implicit learning in the psychiatrically impaired. Unpublished doctoral dissertation, City University of New York.

Abrams, M., & Reber, A. S. (1989). Implicit learning in special populations. Journal of Psycholinguistic Research, 17, 425–439.

Allen, R., & Reber, A. S. (1980). Very long term memory for tacit knowledge. Cognition, 8, 175–185.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. Psychological Review, 85, 249–277.

Anderson, J. R. (1979). Further arguments concerning representations for mental imagery. Psychology Review, 86, 395–406.

Anderson, J. R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.

Bauer, R. M. (1984). Autonomic recognition of names and faces in prosopagnosia: A neuropsychological application of the guilty knowledge test. Neuropsychologia, 22, 457–469.

Bergson, H. (1913). Introduction to metaphysics. New York: Liberal Arts Press.

Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. Quarterly Journal of Experimental Psychology, 36, 209–231.

Bialystok, E. (1981). Some evidence for the integrity and interaction of two knowledge sources. In R. Anderson (Ed.), New directions in research on the acquisition and use of a second language. Rowley, MA: Newbury.

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. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), Cognition and categorization (pp. 169–211). New York: Wiley.

Cantor, G. W. (1980, April). On symbol string-type and the generalization of the implicit learning paradigm. Paper presented at the annual meeting of the Eastern Psychological Association, New York.

Carmody, D. P., Kundel, H. L., & Toto, L. C. (1984). Comparison scans while reading chest images: Taught, but not practiced. Investigative Radiology, 19, 462–466.

Chomsky, N. (1980). Rules and representations. New York: Columbia University Press.

Corteen, R. S., & Wood, B. (1972). Autonomic responses to shock-associated words in an unattended channel. Journal of Experimental Psychology, 94, 308–313.

Croce, B. (1922). Aesthetic (2nd ed.). New York: Norwood.

Danks, J. H., & Gans, D. L. (1975). Acquisition and utilization of a rule structure. Journal of Experimental Psychology: Human Learning and Memory, 1, 201–208.

Derks, P. L. (1963). Effect of run length in the gambler's fallacy. Journal of Experimental Psychology, 65, 213–214.

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.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1985). On consciousness in syntactical learning and judgment: A reply to Reber, Allen, and Regan. Journal of Experimental Psychology: General, 114, 25–32.

Ellenberger, H. F. (1970). The discovery of the unconscious. New York: Basic Books.

Erdelyi, M. H. (1974). A new look at the New Look: Perceptual defense and vigilance. Psychological Review, 81, 1–25.

Erdelyi, M. H. (1985). Psychoanalysis: Freud's cognitive psychology. New York: Freeman.

Fodor, J. A. (1975). The language of thought. Cambridge, MA: Harvard University Press.

Fodor, J. A. (1983). Modularity of mind: An essay on faculty psychology. Cambridge, MA: MIT Press.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 234–257.

Friedman, M. P., Burke, C. J., Cole, M., Keller, L., Millward, R. B., & Estes, W. K. (1964). Two-choice behavior under extended training with shifting probabilities of reinforcement. In R. C. Atkinson (Ed.), Studies in mathematical psychology (pp. 250–291). Stanford, CA: Stanford University Press.

Garner, W. R. (1974). The processing of information and structure. Hillsdale, NJ: Erlbaum.

Garner, W. R. (1978). Aspects of a stimulus: Features, dimensions, and configurations. In E. Rosch & B. B. Lloyd (Eds.), Cognition and categorization (pp. 99–133). Hillsdale, NJ: Erlbaum.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Gleitman, L. R., & Wanner, E. (1982). Language acquisition: The state of the art. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 3–48). New York: Cambridge University Press.

Graf, P., & Mandler, G. (1984). Activation makes words more accessible but not necessarily more retrievable. *Journal of Verbal Learning and Verbal Behavior, 23,* 553–568.

Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 501–518.

Griffin, D. R. (1981). *The question of animal awareness: Evolutionary continuity of mental experience.* New York: Rockefeller University Press.

Griffin, D. R. (1984). *Animal thinking.* Cambridge, MA: Harvard University Press.

Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General, 108,* 356–388.

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information. *American Psychologist, 39,* 1372–1388.

Hayes, N. A. (in press). Consciousness cannot explain syntactical judgment: A rebuttal of Dulany, Carlson, and Dewey. *Memory & Cognition.*

Hayes-Roth, F. (1979). Distinguishing theories of representation: A critique of Anderson's "Arguments Concerning Mental Imagery." *Psychological Review, 86,* 376–392.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *The Behavioral and Brain Sciences, 9,* 1–66.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning and discovery.* Cambridge, MA: MIT Press.

Howard, J. H., & Ballas, J. A. (1980). Syntactic and semantic factors in the classification of nonspeech transient patterns. *Perception and Psychophysics, 28,* 431–439.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110,* 306–340.

Jaynes, J. (1976). *The origins of consciousness in the breakdown of the bicameral mind.* Boston: Houghton Mifflin.

Jones, M. R., & Myers, J. L. (1966). A comparison of two methods of event randomization in probability learning. *Journal of Experimental Psychology, 72,* 909–911.

Jung, C. (1926). *Psychological types* (H. G. Baynes, Trans.). London: Routledge & Kegan Paul.

Kunst-Wilson, W. R., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science, 207,* 557–558.

Lewicki, P. (1985). Nonconscious biasing effects of single instances on subsequent judgments. *Journal of Personality and Social Psychology, 48,* 563–574.

Lewicki, P. (1986a). *Nonconscious social information processing.* New York: Academic Press.

Lewicki, P. (1986b). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 135–146.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and 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.

Mace, W. M. (1974). Ecologically stimulating cognitive psychology: Gibsonian perspectives. In W. Weimer & D. Palermo (Eds.), *Cognition and the symbolic processes.* Hillsdale, NJ: Erlbaum.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197–237.

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

Miller, G. A. (1967). *The psychology of communication.* New York: Basic Books.

Millward, R. B. (1981). Models of concept formation. In R. E. Snow, P.A. Frederico, & W. E. Montague (Eds.), *Aptitude, learning, and instruction: Cognitive process analysis.* Hillsdale, NJ: Erlbaum.

Millward, R. B., & Reber, A. S. (1968). Event-recall in probability learning. *Journal of Verbal Learning and Verbal Behavior, 7,* 980–989.

Millward, R. B., & Reber, A. S. (1972). Probability learning: Contingent-event sequences with lags. *American Journal of Psychology, 85,* 81–98.

Milner, B., Corkin, S., & Teuber, H. L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. *Neuropsychologia, 6,* 215–234.

Moeser, S. D., & Bregman, A. S. (1972). The role of reference in the acquisition of a miniature artificial language. *Journal of Verbal Learning and Verbal Behavior, 11,* 759–769.

Morgan, J. L., & Newport, E. L. (1981). The role of constituent structure in the induction of an artificial language. *Journal of Verbal Learning and Verbal Behavior, 20,* 67–85.

Neisser, U. (1976). *Cognition and reality.* San Francisco: Freeman.

Nelson, K. (1986). *Making sense: The acquisition of shared meaning.* New York: Academic Press.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review, 84,* 231–259.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259–303). Hillsdale, NJ: Erlbaum.

Polanyi, M. (1958). *Personal knowledge: Toward a post-critical philosophy.* Chicago: University of Chicago Press.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77,* 353–363.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology, 83,* 304–308.

Pylyshyn, Z. (1979). Validating computational models: A critique of Anderson's indeterminacy of representation claim. *Psychological Review, 86,* 383–394.

Pylyshyn, Z. (1980). Computation and cognition: Issues in the foundation of cognitive science. *The Behavioral and Brain Sciences, 3,* 111–169.

Rathus, J., Reber, A. S., & Kushner, M. (1988). *Implicit learning and anxiety levels.* Manuscript in preparation.

Reber, A. S. (1965). *Implicit learning of artificial grammars.* Unpublished master's thesis, Brown University.

Reber, A. S. (1966). *A perceptual learning analysis of probability learning.* Unpublished doctoral dissertation, Brown University.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior, 77,* 317–327.

Reber, A. S. (1969). Transfer of syntactic structure 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 and Memory, 2,* 88–94.

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

Reber, A. S., Allen, R., & Regan, S. (1985). Syntactical learning and judgment, still unconscious and still abstract: Comment on Dulany, Carlson, and Dewey. *Journal of Experimental Psychology: General, 114,* 17–24.

Reber, A. S., Kassin, S. M., Lewis, S., & Cantor, G. W. (1980). On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 492–502.

Reber, A. S., & Lewis, S. (1977). Toward a theory of implicit learning: The analysis of the form and structure of a body of tacit knowledge. *Cognition, 5,* 333–361.

Reber, A. S., & Millward, R. B. (1965). Probability learning and memory for event sequences. *Psychonomic Science, 3,* 431–432.

Reber, A. S., & Millward, R. B. (1968). Event observation in probability learning. *Journal of Experimental Psychology, 77,* 317–327.

Reber, A. S., & Millward, R. B. (1971). Event tracking in probability learning. *American Journal of Psychology, 84,* 85–99.

Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review, 74,* 71–80.

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43,* 151–160.

Rosch, E., & Lloyd, B. B. (1978). Representations. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 213–215). Hillsdale, NJ: Erlbaum.

Roter, A. (1985). *Implicit processing: A developmental study.* Unpublished doctoral dissertation, City University of New York.

Roy, F. A. (1982). Action and performance. In A. W. Ellis (Ed.), *Normality and pathology in cognitive functions.* New York: Academic Press.

Rumelhart, D. E., & McClelland, J. L. (Eds.) (1986). *Parallel distributed processing.* Cambridge, MA: MIT Press.

Scarborough, D. L., Cortese, C., & Scarborough, H. S. (1977). Frequency and repetition effects in lexical memory. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 1–17.

Scarborough, D. L., Gerard, L., & Cortese, C. (1979). Accessing lexical memory: The transfer of word repetition effects across task and modality. *Memory & Cognition, 7,* 3–12.

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

Schacter, D. L., & Graf, P. (1986). Effects of elaborative processing on implicit and explicit memory for new associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 432–444.

Schlesinger, I. M. (1982). *Steps to language: Toward a theory of native language acquisition.* Hillsdale, NJ: Erlbaum.

Seamon, J. G., Brody, N., & Kauff, D. M. (1983). Affective discrimination of stimuli that are not recognized: Effects of shadowing, masking, and cerebral laterality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 544–555.

Seamon, J. G., Marsh, R. L., & Brody, N. (1984). Critical importance of exposure duration for affective discrimination of stimuli that are not recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 465–469.

Shallice, T., & Saffran, E. (1986). Lexical processing in the absence of explicit word identification: Evidence from a letter-by-letter reader. *Cognitive Neuropsychology, 3,* 429–458.

Smith, E.E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Strauss, M. E., Weingartner, H., & Thompson, K. (1985). Remembering words and how often they occurred in memory-impaired patients. *Memory & Cognition, 13,* 1507–1510.

Warrington, E. K., & Weiskrantz, L. (1982). Amnesia: A disconnection syndrome? *Neuropsychologia, 20,* 233–248.

Weiskrantz, L. (1986). *Blindsight.* New York: Oxford University Press.

Westcott, M. R. (1968). *Toward a contemporary psychology of intuition.* New York: Holt, Rinehart & Winston.

1990 APA Convention "Call for Programs"

The "Call for Programs" for the 1990 APA annual convention will be included in the October issue of the *APA Monitor.* The 1990 convention will be held in Boston, Massachusetts, from August 10 through August 14. Deadline for submission of program and presentation proposals is December 15, 1989. This earlier deadline is required because many university and college campuses will close for the holidays in mid-December and because the convention is in mid-August. Additional copies of the "Call" will be available from the APA Convention Office in October.