THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 1992, *44A* (2) 193–210

A Reinterpretation of Some Earlier Evidence for Abstractiveness of Implicitly Acquired Knowledge

Pierre Perruchet
Université René-Descartes, Paris, France

Jorge Gallego
Université Pierre et Marie Curie, Paris, France

Chantal Pacteau
Université René-Descartes, Paris, France

Reber and Lewis (1977) exposed subjects to a subset of letter strings generated from a synthetic grammar, then asked them to reorder scrambled letters to generate new grammatical strings. The distribution of the frequency of the bigrams composing their solutions correlated better with the frequency of the bigrams composing the full set of strings generated by the grammar than with the frequency of the bigrams composing the subset of strings displayed in the study phase. In his recent overview on implicit learning, Reber (1989) develops this experimental result into one of the main supports for his contention that studying grammatical letter strings gives access to the abstract structure of the grammar.

However, this result can be accounted for by a set of biases inherent to the Reber and Lewis procedure. In the present experiment, a group of subjects learned from a list of the bigrams making up the study strings, a condition which precludes the abstraction of any high-level rules. The pattern of correlations outlined above also emerged in this condition, thus lending support to our re-interpretation.

The notion of implicit learning originates from the field of artificial grammar learning, as explored over the last twenty years by Reber and his associates (e.g. Reber, 1967; Reber, Allen, & Regan, 1985; see Reber, 1989, for a review). In a typical experiment, subjects are first exposed to a set of letter strings generated from a synthetic grammar that defines authorized letters and the permissible transitions between them. Subjects are given no informa-

Requests for reprints should be sent to P. Perruchet, Laboratoire de Psychologie Différentielle, 28 rue Serpente, 75006 PARIS, France.

tion about the rule-governed nature of the stimuli and are asked to perform a task that diverts them from the search for regularities such as rote learning (e.g. Reber, Kassin, Lewis, & Cantor, 1980), short-term matching (Mathews et al., 1989), or liking judgement (McAndrews & Moscovitch, 1985) of letter strings. Subjects are subsequently asked to categorize new letter strings as grammatical or non-grammatical; non-grammatical items are formed from the same subset of letters, but violate transition rules. Results consistently show that subjects perform this categorization task better than chance. As they are apparently unable to verbalize the rules underlying their decision of well-formedness, their performance is thought to be the end-product of an unconscious abstraction process.

In recent years, these findings have received further support from studies based on other paradigms. For instance, Lewicki, Hill, and Bizot (1988) studied the ability of subjects to improve their performance when the location of a response signal was determined by the pattern of its location on specific earlier trials. After extended practice, subjects exhibit a substantial decrease in reaction time to target signals with predictable locations, although they are unable to articulate any of the complex rules that regulate the sequence of trials (cf. also Lewicki, Czyzewska, & Hoffman, 1987; McKelvie, 1987; Millward & Reber, 1972; Nissen & Bullemer, 1987). Implicit learning has been also shown for interactive process control tasks, in which the learner influences the nature of the stimulation he or she receives. For instance, in one of Berry and Broadbent's (1988) experiments, subjects were required to imagine that they were in charge of a city transportation department and were instructed to manoeuvre the time interval between buses to reach and maintain a specified load of passengers per bus. With training in a computer-simulated interactive situation, subjects were able to make appropriate adjustments, although they were unable to verbalize the actual function relating the two variables (cf. also Broadbent, Fitzgerald, & Broadbent, 1986; Hayes & Broadbent, 1988; Stanley, Mathews, Buss, & Kotler-Cope, 1989). On the whole, these studies apparently provide consistent evidence that people can unconsciously abstract environmental regularities and use this tacit knowledge to improve performance. This ability is considered to be highly relevant to accounts of how humans cope with the complex environment of everyday life and is thought to encompass perceptual and social behaviour, language, and so on.

Another set of studies suggests, however, that the knowledge base underlying improvement in performance in implicit learning situations may be available to consciousness. As noted by Broadbent (in press), some of the discrepancies between studies could be due to the method of operationalization of consciousness. Typically, authors positing that some knowledge always remains inaccessible to conscious reflection, such as Reber and

Lewicki, rely on recall-like verbal protocols, whereas authors who challenge this claim use recognition-like tests (e.g. Dulany, Carlson, & Dewey, 1984; Perruchet & Pacteau, 1990). Another issue may, however, be of greater importance. Authors stressing the availability to consciousness of the knowledge underlying performance also claim that this knowledge is far less complex and abstract than is commonly thought.

In artificial grammar learning, several studies indicate that specific and fragmentary knowledge of study items components, in particular first and last letters and permissible bigrams (i.e. pairs of letters) or trigrams, is sufficient to account for grammaticality judgements (Dulany et al., 1984; Mathews et al., 1989; Perruchet & Pacteau, 1990). Perruchet and Pacteau (1990) obtained explicit knowledge on valid bigrams through a recognition test and then performed quantitative simulations of grammatical judgements on the basis of these data. The resulting performance pattern matched nicely with observed results when simulated judgements of non-grammaticality were generated for all test strings containing at least one unrecognized pair of letters. In a more sophisticated way, Druhan and Mathews (1989) developed a computational model based on Holland et al.'s (1986) induction theory. They used verbal instructions generated by subjects studying letter strings in standard conditions as input. When the rules derived from verbal instructions were automatically tuned by an optimization algorithm using feedback on correctness of responses, simulated performance approached that of the original subjects.

Similar conclusions may be drawn from studies involving other implicit learning paradigms. Perruchet, Gallego, and Savy (1990) observed that the rules constituting the series of trials in the Lewicki et al. (1988) paradigm also altered the relative frequency of particular transitions on the whole sequence of trials. They provide evidence, through replication and re-analysis of the Lewicki et al. data, that subjects' performance may simply reflect this far more elementary informational content. Likewise, Sanderson (1989) showed that performance in interactive control tasks could be due to familiarization with peripheral aspects of the task rather than to the internalization of the underlying structure of the situation.

In fact, these recent studies do *not* demonstrate that people really use conscious knowledge of surface features of the situation. By showing that this knowledge could be sufficient to account for performance, they merely provide an alternative interpretation of the traditional abstractionist point of view. Comparing the explanatory power of these two interpretations clearly requires other experimental evidence. In his recent overview on implicit learning, Reber (1989, pp. 225 ff) puts forward two main arguments to show that implicit learning gives access to deep, abstract structure of the stimuli, in addition to fragmentary pieces of knowledge on their surface features. The present paper is aimed at examining the empirical and theoretical validity of these two arguments. Because both refer to the field of grammar learning, the following discussion will be limited to this area of research.

Reber's First Argument: The Transfer to Superficially Different, but Structurally Similar Stimuli

The first argument draws on results initially collected by Reber (1969), and more recently replicated and extended by Mathews et al. (1989). Briefly, these studies demonstrate that learning from letter strings generated by one grammar transfers to letter strings generated by a grammar sharing the same formal rules, but using another set of letters. These results are interpreted as showing that subjects are able to abstract automatically the "syntax" of the displayed material, and they can use this structural knowledge independently of the "vocabulary".

Although Mathews et al. (1989) adhere to Reber's abstractionist position, certain features of their results run counter to a strict version of this point of view. For instance, transfer to a different letter set is better under intentional than under implicit instructions, whereas the theory posits unconsciousness of abstraction. Furthermore, transfer to a different letter set operates equally well for yoked subjects who learn from verbal instructions given by experimental subjects as for the experimental subjects themselves. Most of these verbal instructions are to select or avoid specific bigrams or trigrams in specific locations in the strings. Hence, transfer can hardly be put forward as an argument for an abstraction process generating high-level rules from entire letter strings.

The results of Mathews et al. leave open one possibility that people abstract "local syntaxes". However, even this weak version of the abstractionist viewpoint may be challenged. On formal grounds, the artificial grammar setting is a categorization task. Research conducted over the last decade or so in the field of categorization has convincingly demonstrated that the capacity to transfer to new test items does not prove that people have elaborated abstract representations from the study examplars. This hinges on two core issues. The first pertains to the moment of processing and discriminates between what Estes (1986) has termed "late" versus "early computation" models. In early computation models, processes underlying categorization are accomplished during the study phase, and their end products are merely examined at the test time. In late computation models, studied examplars are thought to be stored in memory without transformation, and all the computations needed for making a decision are performed when the subjects encounter new items. The postulate of Reber and Matthews et al., that learning is implicit, is obviously at variance with late computation models, which have received increasing support in the current literature. The second issue concerns the mode of processing. In the traditional view, people generate abstract representations, such as schemas or prototypes. More recent interpretations, inspired by seminal papers by Medin and Schaffer (1978) and Brooks (1978), focus on non-analytic, analogical processing (for a wider-ranging overview, see Jacoby & Brooks, 1984; Medin & Ross, 1990).

To illustrate how these recent developments apply to the field of grammar learning, assume that subjects have been asked to learn trigram CVC, and then exhibit positive transfer with trigram DJD. In the traditional view, transfer is taken as evidence that subjects automatically use the repetitive presentation of CVC to abstract something like "one letter surrounded by two identical letters", then store this formal rule in memory and apply it when coping with a new letter string. An alternative explanation is that subjects store the specific trigram CVC, and deal subsequently with DJD in the same way as they dealt with CVC on the basis of the global similarity between both items. The key point here is that this alternative interpretation accounts for transfer without assuming that people form an abstract representation while studying letter strings (for further discussion on these points, see Mathews, in press; Perruchet & Pacteau, 1991).

Reber's Second Argument: The Performance Pattern Reflects the Underlying Grammar, Not the Studied Stimuli

Reber's (1989) second argument is based on an intriguing experimental result initially reported by Reber and Lewis (1977). In this study, the standard categorization test was replaced by a task in which subjects had to reorder scrambled letters to generate grammatical strings. Nature of knowledge resulting from exposure to grammatical letter strings is assessed by detailed analysis of performance on this task. The bigrams occurring in subjects' solutions to the anagrams were tabulated, and their frequencies were compared against (a) the frequency of occurrence of the bigrams composing the strings actually displayed in the study phase, and (b) the frequency of occurrence of the bigrams composing the full set of strings generated by the formal grammar (hereafter: the virtual strings). The correlation coefficients were 0.04 and 0.72 respectively. Reber (1989) concludes that the failure of the correlation involving the study phase bigrams "to be different from zero suggests that subjects were not solving the anagrams on the basis of superficial knowledge of frequency of bigrams". Rather, the high correlation with the bigrams making up the virtual strings indicates that subjects "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" (Reber, 1989, p. 226).

Although the rationale for this argument is sound, a number of biases potentially flaw the empirical findings on which it is grounded. Let us start by considering the selection of the study strings from the virtual strings. Selection has two direct consequences. The first is to reduce drastically the variability of bigram frequency. The variance of the distribution of the bigram frequency of selected items is about one ninth of the corresponding value calculated from the whole set (3.85 vs. 35.89). Low variance obviously makes it difficult to obtain high correlations. The second effect of selection is that the distribution of frequencies of the bigrams making up the study and the virtual items differs. This constitutes a necessary condition for obtaining different correlations with both sets of data, as Reber and Lewis did. However, close scrutiny of the material shows that a bias may have been introduced. Consider the under-represented bigrams—that is, those bigrams having a lower proportion in the study items than in the whole sample. The observed pattern of correlations shows that these bigrams are especially easy to learn; indeed, the higher correlation of subject generated bigrams with virtual than with study items implies that the under-represented bigrams are generated in anagrams more often than could have been predicted from their effective occurrences in studied items (the reverse is true for the over-represented bigrams, which are not dealt with here for simplicity's sake). The four under-represented bigrams exhibiting the largest differences between study and virtual frequencies are TV, TT, VV, and XX. TV is an extremely common abbreviation, and other bigrams are doublets. These bigrams may have been especially easy to learn because of their relative salience rather than because of their frequency in the virtual set of grammatical strings.

Another more damaging bias than those deriving from study item selection arises from choice of the test material. Subjects start from strings of letters to solve anagrams, which places severe constraints on the nature of the pairs that may be produced. Suppose subjects learn in the study phase that VV is a valid bigram. They have little means of demonstrating this knowledge if there are few Vs in the test material, and overall, the frequency of VV in their response will be partly dependent on the number of available Vs. In the Reber and Lewis study, the letter strings used in the anagram task were obtained by scrambling the grammatical items that were not initially displayed. As a consequence, subjects are prompted to fit their bigram production with the bigrams making up these grammatical items. Indeed, if these items contained a large number of, say, VV bigrams, subjects had a large number of Vs at their disposal in the scrambled strings, which facilitated VV production. The crucial point is that the frequency of bigrams composing the grammatical test strings before scrambling closely matched the bigram frequency in the virtual strings (the two distributions correlated to 0.949) and were only moderately related to the bigram frequency in the study strings ($r = 0.234$). This correlational pattern is not simply an unfortunate chance effect, but, rather, reflects a logical constraint. The fact that the virtual strings were exhaustively partitioned into study and test strings makes the frequency of bigrams composing the virtual strings the (weighted) mean of bigrams composing study and test strings. It can be demonstrated that the correlation between two given series of values (at least when this correlation is positive) is lower than the correlation between one of the series and the mean of the two series. Thus, to sum up, the constraints created by the test material make it impossible for subjects to produce bigrams mimicking the distribution of bigrams in the study strings, whereas they can almost perfectly parallel the distribution of bigrams in the virtual strings.

Overview of the Present Study

The experiment described in this article was designed to strengthen our contention that the Reber and Lewis (1977) correlational result cannot serve as evidence for Reber's (1989) abstractionist interpretation, by showing that an identical pattern of results can be obtained when subjects are not given information needed for abstracting high-level rules.

The procedure was the following: The first group of subjects (string group) studied grammatical strings of letters in the study phase and were then given an anagram task in the test phase. The material used in both phases was identical to the one used in the original Reber and Lewis (1977) study. The string group was designed to replicate the original correlational results. However, our procedure differed from Reber and Lewis in several ways, in particular in the amount of training given. In the original study, subjects were exposed to four learning sessions. This has at least two negative consequences: First, the study items are displayed before the anagram solving trials on each of the sessions. Thus, subjects have a clear idea of the task requirements when dealing with the relevant information and may have engaged in explicit modes of learning, including hypothesis formulation, rule testing, mnemonic strategies, and so on. The second touches more specifically on the correlational results. As subjects progressively solve more anagrams correctly, correlations between bigrams frequencies approach fixed values, which depend exclusively on experimental material. If subjects solve all anagrams correctly, the correlations of the frequency of bigrams making up the anagram solutions (which are, in this case, the subset of grammatical letter strings not initially displayed) with, for instance, the bigrams composing virtual strings is 0.949, as stated above. To correct for these problems, duration of learning was restricted to a single session (the power of inferential tests was ensured by collecting data on a far larger sample of subjects than Reber and Lewis did).

The second group of subjects (bigram group) was given in the study phase the individual bigrams making up the strings displayed to the string group.

As in Perruchet and Pateau (1990, expt. 1), this was intended to convey information on valid bigrams and their frequency of occurrence, but to preclude the abstraction of any high-level rules. The test phase consisted of the same anagram task as for the string-group subjects. The mean scores for the bigram group were expected to be much lower. In the abstractionist framework, this would occur because subjects cannot abstract the deep structure of the grammar during the study phase. The studies showing that subjects exposed to letter strings acquired specific knowledge pertaining, for instance, to the first or last valid letters (e.g. Mathews et al., 1989) provide another interpretation. Thus bigram-group subjects, who only observed letter pairs in study phase, could not acquire this knowledge, which is obviously useful for solving anagrams. The correlational results are of more fundamental interest, in that they discriminate concurrent interpretations. In keeping with Reber's hypothesis, there is no reason to expect that the frequency of bigrams composing subjects' generated strings should correlate better with the frequency of the bigrams composing virtual strings than with the frequency of bigrams composing study strings. However, if the correlation pattern is due to one or a combination of several of the biases pointed out above, stronger correlations with bigrams composing virtual than study strings must emerge in both the bigram group and the string group.

Method

Subjects. Subjects were 53 first-year university students majoring in psychology. Twelve additional subjects participated to the experiment, but they were subsequently eliminated from data analysis because they failed to follow instructions on the anagram tasks. These subjects generated at least 10 anagrams where the number of constituent letters did not match the number of letters in the original strings.

Subjects were run in 3 groups of 21, 21, and 23. They were randomly allocated on an alternating basis to experimental conditions within test groups. To make this possible, the instructions specific to each group were provided on separate sheets of paper.

Materials. The stimuli consisted of letter strings generated by the Reber and Lewis (1977) grammar shown in Figure 1. The study items for the S group were the 15 letter strings used by Reber and Lewis. The strings were typed on a sheet of paper in a single column in random order (see Table 1).

The bigram group was shown the 78 bigrams that can be extracted from the 15 letter strings. These were typed in six columns, in random order, as displayed in Table 2.

The test phase used the 28 remaining letter strings that the grammar could generate. These letter strings were scrambled and presented in two columns in random order (see Table 3). A dotted line was provided opposite each letter string for the subject's response.

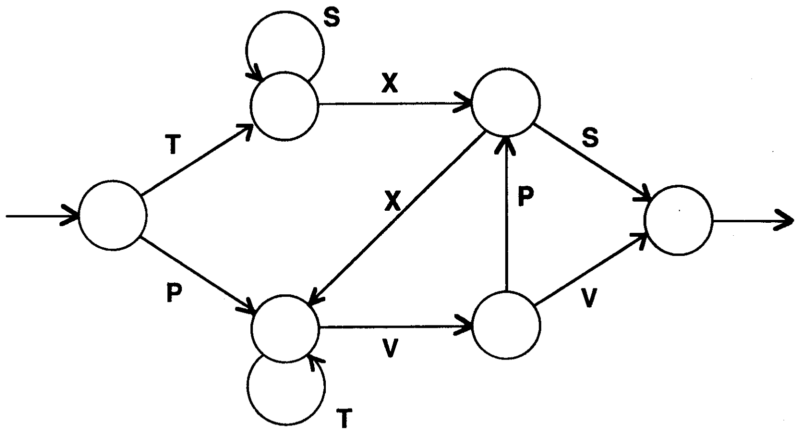

FIG. 1. Schematic diagram of the grammar used in the present experiment (taken from Reber & Lewis, 1977).

TABLE 1
List of Letter Strings Displayed in the Study Phase to the String Group

PTTTVPS
TSSSXS
PVV
PVPXVPS
TSSXXVV
TXXVPXVV
PTVPS
PVPXVV
TSSSSSXS
PTTTVV
TSXXVPS
TXS
TSXS
PVPXTVPS
PTVPXTVV

Procedure

For all subjects, the session comprised a learning phase and a test phase. Subjects first received two booklets in a folder. They were orally asked to remove the first booklet from the folder and to read the instructions printed on the front page. Instructions for the string group ($n = 26$) and the bigram group ($n = 27$) were similar, except for variations adapted to the nature of the displayed material:

TABLE 2
List of Bigrams Displayed in the Study Phase to the Bigram Group

XS	SX	SS	PS	XV	PX
XT	VV	VP	TX	XT	TT
VV	VP	TV	VV	PV	TS
VP	XV	PS	SX	TV	VP
SS	TT	PX	XS	VP	TX
XX	PX	PT	VP	PS	XV
TS	SS	XV	TV	SX	PT
PX	TS	VV	TT	VP	TV
PV	PT	VP	SS	SS	XX
TV	VP	XX	PV	PX	VV
SX	SS	TT	XV	TS	VP
XV	PS	XS	TV	VV	SX
PS	VP	TS	PT	XS	SS

TABLE 3
List of Scrambled Letter strings Displayed in the Test Phase

XPPVVTV	(PTVPXVV)	XSVVTSSX	(TSSSXXVV)
VPTPST	(PTTVPS)	VTPTVTTT	(PTTTTTVV)
PVVXVTPT	(PVPXTTVV)[a]	XVXVTS	(TSXXVV)
SSTSXSS	(TSSSSXS)	STXSS	(TSSXS)
TXPTPVVV	(PTTVPXVV)[a]	VTVTTXX	(TXXTTVV)
VXXTPSST	(TSXXTVPS)	VTPV	(PTVV)
TXTVVSSX	(TSSXXTVV)	VXTTTSXP	(TXXTTVPS)
PXVTPVPS	(PTVPXVPS)[b]	VXXSSTPS	(TSSXXVPS)
TTTXXVSV	(TSXXTTVV)	XPVTVVP	(PVPXTVV)
TTPTTVPS	(PTTTTVPS)	TTTPVVT	(PTTTTVV)
SPXTVX	(TXXVPS)	TTXSVVX	(TSXXTVV)
VXVTX	(TXXVV)	XVTVTXTT	(TXXTTTVV)
TVSPXTX	(TXXTVPS)	PPSV	(PVPS)
TTVXVX	(TXXTVV)	VTTPV	(PTTVV)

Note: [a]These items have three correct solutions: PVPXTTVV, PTTVPXVV, and PTVPXTVV.
[b]Another correct solution is: PVPXTVPS.
Grammatical solutions are in parentheses.

In this experiment you have to learn (strings from 3 to 8 letters in length / pairs of letters). The composing letters are P, S, T, V, X. These (strings / pairs) are displayed on the next page. After the signal, you will have 10 minutes to study them.

For both groups, the front page began with the sentence in upper case:

Do not turn the page before the signal.

When the signal was given, subjects studied the items (the grammatical strings of letters or the bigrams) displayed on the second page of the booklet for 10 minutes. The first booklets were then collected by the experimenter.

Subjects were orally requested to take the second booklet out of the folder, and to read the instructions on the front page, which again began with the sentence:

Do not turn the page before the signal.

The subjects were asked to arrange the letter strings in the "right" order. Subjects in the string group were informed that the study items were generated by a set of rules, and correctness was defined as conformity with these rules. The bigram group was told that a correct letter string was defined by the fact that each pair of contiguous letters had to belong to the previously studied list. When the signal was given, subjects began to write their responses. They were urged to work quickly, but no time limit was imposed.

Results

Probability of Correct Solutions. In order to obtain a basis for assessment of the observed performance, the probability of correct responding was computed for each letter string as the ratio of the number of correct solutions (usually one, in a few cases two or three) over the number of possible arrangements. The mean probability was 0.013. The values for the letter strings of length 4 to 8 are reported in Table 4, first column.

For the string group, the mean proportion of correct anagram solutions was 0.113 ($s=0.137$). As shown in Table 4, second column, there was a marked decrease in correct anagram solutions with lengthening of the letter

TABLE 4
Proportion of Correct Anagram Solution by Chance, and Observed in the String Group and the Bigram Group

Length	*Chance Value*	*String Group*	*Bigram Group*
4 (2)	0.083	0.230	0.095
5 (3)	0.039	0.307	0.063
6 (4)	0.006	0.117	0.010
7 (7)	0·006	0·100	0·010
8 (12)	0·001	0·052	0·007
Weighted mean	0·013	0·113	0·020

Note: Proportions were calculated for each length of anagrams. Mean values were computed after weighting by the number of strings representative of each length, as shown in parentheses.

strings. This effect is not surprising, as long letter strings provide more opportunities for error.

As expected, the bigram group performed considerably less well than the string group (Table 4, third column). The mean proportion of correct solutions was 0.02 ($s = 0.043$). Performance was again closely dependent on the length of the letter strings. However, performance was better than chance for all lengths, a result which has a binomial probability of 0.031 ($1/2^5$) of occurring.[1]

Correlational Pattern. The proportion of occurrence of each of the 25 possible pairs of letters was assessed for each subject, then corrected for the likelihood of assembling letters by chance using the Reber and Lewis' correction formula (reproduced in the note to Table 5). Table 5 presents the raw and adjusted proportions averaged over subjects, along with frequencies of study and virtual bigrams.

The product–moment correlations for these values were computed using two different methods. Correlations were first computed for the 16 valid bigrams after averaging over subjects (Reber and Lewis computed their correlations in this way, although they only used rank-order information. We used product–moment correlations in order to test for statistical significance of differences). The coefficient values appear in Table 6, left panel. An unpredicted result was that correlations were higher for the bigram group than for the string group. *T*-tests (McNemar, 1969, p. 158) show that the differences were significant, for both the observed bigrams, $t(13) = 6.15$, $p < 0.001$, and the virtual bigrams, $t(13) = 3.47$, $p = 0.004$. However, within each group, the correlations were higher for virtual bigrams than for observed bigrams. Although fairly large, the differences failed to reach significance [string group: $t(13) = 1.67$; bigram group: $t(13) = 1.48$].

Correlations were also computed on a within-subject basis, then averaged over subjects (after r to z transformation). Correlations were higher for the bigram group than for the string group for observed bigrams only,

---

[1] An anonymous reviewer of an earlier version of this paper noted that anagram solutions may be affected by linguistic preferences, which may coincide with the grammatical ordering. In order to test this hypothesis, an additional group of 45 students in psychology were asked to arrange the test strings as they saw fit, without prior exposure to the study material. The proportion of grammatical responses was near chance (0.010 vs. 0.013, respectively). This result ensures that above-chance performance of subjects from the string or bigram groups is imputable to their exposure to material in the study phase.

This does not entail that linguistic preferences have no effect. Positive and negative effects on specific strings may compensate one another, leaving mean performance unaffected. Consider the two 4-letter strings to illustrate this point. Seven out of the 45 subjects generated the grammatical string P T V V, whereas none generated P V P S, which suggests a marked preference effect. However, the mean proportion of grammatical responses for the 4-letter strings (7/90, i.e. 0.078) is close to chance level (0.083, as indicated in Table 4).

TABLE 5
Numbers of Occurrences of Each Bigram in the Study Strings and in the Whole Set of Strings Generated by the Grammar, and Raw and Corrected Proportion of Occurrences of Each Bigram

	Occurrences of each Bigram		*Scores*			
	Material		*String Group*		*Bigram Group*	
	Study	*Virtual*	P	P′	P	P′
SS	7	15	0.056	0.022	0.045	0.011
XX	3	17	0.052	0.027	0.041	0.016
TT	4	23	0.123	0.049	0.099	0.023
VV	6	23	0.086	0.050	0.071	0.035
XS	4	6	0.031	−0.016	0.033	−0.014
SX	5	14	0.031	0.008	0.037	0.015
PS	5	14	0.028	0.007	0.026	0.005
XV	6	14	0.049	0.006	0.048	0.005
PX	5	10	0.014	0.002	0.014	0.003
XT	2	13	0.030	−0.020	0.046	−0.003
VP	10	24	0.037	0.004	0.056	0.024
TX	2	9	0.057	−0.002	0.044	−0.016
PV	4	7	0.044	−0.002	0.046	0.000
TV	6	26	0.072	0.001	0.084	0.014
TS	5	14	0.027	−0.029	0.030	−0.027
PT	4	13	0.043	0.007	0.044	0.008

P = Raw proportion of occurrences.
P' = Corrected proportion of occurrences.
Note: Correction for guessing was computed using the Reber and Lewis formula (1977, p. 342): $P' = (P - Pg) / (1 - Pg)$, where P is the proportion of actual occurrences in subjects' solutions, and Pg the proportion of occurrences resulting from guessing.

TABLE 6
Product–Moment Correlations of the Frequency of Bigrams Composing Subjects' Solutions with the Frequency of Study Bigrams of One Hand, and the Frequency of the Bigrams Composing the Whole Set of Strings the Grammar May Generate on the Other Hand

	r *on Mean Scores*		*Means of Within-Subject* r	
	String Group	*Bigram Group*	*String Group*	*Bigram Group*
Study bigrams	0.175	0.458	0.14	0.26
Virtual bigrams	0.547	0.730	0.43	0.43

Note: Coefficients were computed using two different methods (see text) for string group and for bigram group.

$t(51) = 2.52, p = 0.014$. Otherwise, the results (Table 6, right panel) display the same general pattern as above. Strikingly, the correlations were significantly higher with virtual bigrams than with observed bigrams, for both the string group, $t(25) = 10.61, p < 0.001$, and the bigram group, $t(26) = 3.40, p = 0.002$.

The fact that correlations of interest were generally higher for the bigram group than for the string group, especially when the data involving observed strings were examined, was not anticipated. This difference clearly indicates that the performance of the string group was less dependent on bigram frequency than was bigram group performance. The abstractionist interpretation would be that performance in the string group was also affected by the internal representation of higher-level rules. However, this result is also congruent with a general framework positing that implicit learning only gives access to specific features of the stimuli. In the string group, frequency of occurrence may compete with other factors in determining the strength of memory traces of the bigrams, in particular with the position of the bigrams within the strings. A case-by-case analysis of results displayed in Table 5 supports this interpretation. For instance, bigram VV had the best score in the string group, although it was initially displayed less often than SS and VP. This may be accounted for by positional factors: all VV bigrams are terminal, and thus highly salient, whereas all SS and VP bigrams are internal.

Discussion

As in Reber and Lewis (1977), the frequency of bigrams generated by subjects trained on grammatical letter strings correlated better with the frequency of the bigrams making up the whole set of strings the grammar can generate than with the frequency of the bigrams initially displayed. However, an identical pattern was observed when subjects were first exposed to individual pairs of letters constitutive of grammatical strings, a condition precluding the abstraction of high-level rules. This latter finding rules out Reber's (1989) contention that obtaining this pattern of correlations in subjects trained in standard conditions provides evidence for the abstraction of the deep structure of the grammar used to generate the strings.

One potential problem with our results is that the pattern of correlations, although reflecting the same general trend as the one reported by Reber and Lewis (1977), exhibits much less contrast. Correlations of performance of subjects studying grammatical letter strings with observed and virtual bigram frequencies were 0.17 and 0.55, respectively, in our experiment, whereas the corresponding values in the Reber and Lewis study were 0.04 and 0.72. There is empirical evidence for attributing these differences to the differential amount of subject training. In the Reber and Lewis study, in which subjects had four learning sessions, correlations of performance with virtual bigram frequencies were 0.44, 0.58, 0.68, and 0.77, respectively, for Sessions 1 to 4 (Reber & Lewis, 1977, Table 7. Corresponding data for observed bigrams were not reported). Our subjects had only one learning session, and their mean performance level (0.11 correct solutions) was notably lower than the one reported by Reber and Lewis (overall mean: 0.46 correct solutions), even on their first session (0.31 correct solutions). On this basis, it could be argued that our results only pertain to the early phase of learning, and that abstraction of deep structure occurs at a later stage.

This objection calls for several comments. (1) Differences in level of learning could be less marked than the proportion of correct solutions suggests. In the Reber and Lewis procedure, subjects had to reorder a set of shuffled cards, each containing one letter; thus they could not make any mistakes regarding the number or nature of the letters in their solution. In addition, a letter was placed in its proper location on three-quarters of the trials. Thus part of the difference in performance may be due to the fact that the Reber and Lewis subjects had their responses constrained to a greater extent than ours. (2) Our learning conditions in fact closely paralleled those used in earlier studies, which claim to provide evidence for implicit abstraction. Subjects studied the letter strings for ten minutes; as a case in point, Reber, Kassin, Lewis, and Cantor (1980, experiment 1) displayed their study list, which was generated by the same grammar as the present one, for only seven minutes. (3) As argued in the introduction, extending the duration of learning has several negative consequences. Thus the data collected after intensive training in the Reber and Lewis experiment may not be relevant to implicit learning. Further and more damaging to their arguments, it is doubtful that the correlational pattern obtained after intensive training has any significance as regards processes underlying performance, as correlations approach fixed values as subjects solve more and more anagrams. Thus, overall, differences in the amount of training between Reber and Lewis' studies and ours do not seriously undermine our conclusions.

Another possible reservation concerns the type of evidence provided in the present study. Showing that the Reber and Lewis correlational results are flawed does *not* entail that stronger correlations with bigrams composing study than virtual strings would be obtained in more valid conditions; hence, our study does not indicate that subjects are more sensitive to study than to the virtual strings. We fully agree that we provide only negative, and hence limited, evidence. However, as stated in the introduction, most of the biases to which we draw attention are intrinsic components of the Reber and Lewis procedure. Thus carrying out new experiments patterned after the Reber and Lewis procedure after eliminating its biases constitutes a dead end. We are not challenging the value of anagram technique as a means of investigating knowledge acquired in the artificial grammar setting. Rather, our point is that the technique is ill suited for the specific use that was the main focus of the present article.

SUMMARY AND CONCLUSION

Earlier work on artificial grammar learning has shown that specific and conscious knowledge regarding, for instance, valid bigrams or trigrams could be sufficient to account for grammaticality judgements (Druhan & Mathews, 1989; Dulany et al., 1984; Perruchet & Pacteau, 1990). These findings do not exclude the fact, however, that people can acquire some unconscious representation of the underlying structure of the grammar. In his recent theoretical overview on implicit learning, Reber (1989) puts forward two arguments supporting this assumption. The present article shows that both arguments are grounded on experimental data that can be accounted for in other ways.

What is challenged in this and earlier works (Perruchet et al., 1990; Perruchet & Pacteau, 1990, 1991) may be defined as implicit abstraction, that is, the unconscious formation of new abstract structures. Our position is that in current experimental settings, implicit learning conditions may only generate specific knowledge. We take it for granted that people may abstract rules from factual information. However, abstraction may be linked to the use of logical reasoning and analytic modes of processing.

This contention does not minimize the importance of implicit learning in adaptive behaviour. Restricting implicit learning to the acquisition of specific knowledge at the expense of abstract structures or rules should not be construed as a major limitation on its explanatory power. There is growing evidence in the fields of categorization and problem-solving for models positing the primacy of the specific knowledge in the formation of abstract structures. Most of the data that seemed at one time to be straightforward support for rule, schema, or prototype abstraction have been reinterpreted without assuming that abstractive abilities are called into play. For instance, it has been repeatedly shown that subjects can classify previously unseen prototypes more accurately than exemplars on which they were trained (e.g. Posner & Keele, 1968; parenthetically, this bears striking resemblance to the correlational data in this article, as Reber, 1989, noted). Medin and Shaffer (1978) show that a mathematical model based on their exemplar-based theory of categorization can predict the Posner and Keele results as well as a prototype theory. Other related works aimed at illustrating that a large amount of adaptive behaviour derives from the acquisition and use of specific knowledge are described in the Medin and Ross (1990) review.

REFERENCES

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

Broadbent, D.E. (in press). Recall, recognition, and implicit knowledge. In W. Kessen, A. Ortony, & F.I.M. Craik (Eds.), *Festschrift for George Mandler* (provisional title) (pp. 125–134). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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, D.N. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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

Dulany, D.E., Carlson, 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.

Estes, W.K. (1986). Array models for category learning. *Cognitive Psychology, 18*, 500–549.

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

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

Jacoby, L.L. & Brooks, L.R. (1984). Non-analytic cognition: Memory, perception, and concept learning. In G. Bower (Ed.), *The psychology of learning and motivation, Vol. 18* (pp. 1–47). New York, Academic Press.

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.

Mathews, R.C. (1990). Abstractness of implicit grammar knowledge: Comments on Perruchet and Pacteau's analysis of synthetic grammar learning. *Journal of Experimental Psychology: General, 119*, 412–416.

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

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

McNemar, Q. (1969). *Psychological statistics*. New York: J. Wiley and Sons (4th edition).

McKelvie, S.J. (1987). Learning and awareness in the Hebb digits task. *The Journal of General Psychology, 114*, 75–88.

Medin, D.L. & Shaffer, M.M. (1990). A context theory of classification learning. *Psychological Review, 85*, 207–238.

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

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

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

Perruchet, P., Gallego, J., & Savy, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. *Cognitive Psychology*, *22*, 493–516.

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

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

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

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

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

Reber, A.S., Allen, R., & 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. (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). Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition*, *5*, 331–361.

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

Stanley, W.B., Mathews, R.C., Buss, R.R., & Kotler-Cope, S. (1989). Insight without awareness: on the interaction of verbalization, instruction and practice in a simulated process control task. *The Quarterly Journal of Experimental Psychology*, *41A*, 553–577.

Revised manuscript received 7 January 1991